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

A model of cognitive and operational memory of organizations in changing worlds

Giovanni Dosi [a]
Luigi Marengo [b]
Evita Paraskevopoulou [c]
Marco Valente [d]

[a] Institute of Economics and LEM, Scuola Superiore Sant'Anna, Pisa, Italy
[b] Dipartimento di Impresa e Management, Universita' LUISS, Rome, Italy
[c] Departemento de Economia de la Empresa, Universidad Carlos III, Madrid, Spain
[d] School of Economics, Universita' de L'Aquila, Italy

2015/01 November 2015
ISSN(ONLINE) 2284-0400

A model of cognitive and operational memory of organizations in changing worlds

Giovanni Dosi[1], Luigi Marengo[2], Evita Paraskevopoulou[3], Marco Valente[4]

1 LEM, Istituto di Economia, Scuola Superiore Sant' Anna, Pisa, g.dosi@sssup.it
2 Dipartimento di Impresa e Management, Università LUISS, Roma, lmarengo@luiss.it
3 Departamento de Economia de la Empresa, Universidad Carlos III de Madrid, eparaske@emp.uc3m.es
4 School of Economics, Università de L'Aquila, marco.valente@univaq.it

Abstract

This work analyzes and models the nature and dynamics of organizational memory, as such an essential ingredient of organizational capabilities. There are two sides to it, namely a cognitive side, involving the beliefs and interpretative frameworks by which the organization categorizes the states of the world and its own internal states, and an operational one, including routines and procedures that store the knowledge of how to do things. We formalize both types of memory by means of evolving systems of condition-action rules and investigate their performance in different environments characterized by varying degrees of complexity and non-stationarity. Broadly speaking, in simple and stable environments memory does not matter, provided it satisfies some minimal requirements. In more complex and gradually changing ones 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. Moreover, above some (modest) complexity threshold stable and robust cognitive categorizations and routinized behavior emerge.

Keywords: organizational memory, routines, cognitive categories, condition-action rules.

Acknowledgements: we thank for the many useful comments and suggestions Dan Levinthal, the participants to the Organization Science Winter Conference, Steamboat Springs, February 2011, and in particular Bruce Kogut ; the "2nd International Conference on Path-Dependence", Freie Universität Berlin, March 2011; the DIME Final conference, Maastricht, April 2011 ; the Conference "Dynamic Capabilities and the Sustainable Competitiveness of Firms and Nations", St. Petersburg Graduate School of Management, October 2012; and the workshop "New Frontiers in the Economics and Management of Innovation", KITeS, Bocconi University, Milan, March 2012, and in particular Sarah Kaplan. Financial support from the Institute for New Economic Thinking (INET) grant IN 01100022. We also acknowledge financial support from European Union's 7th FP for research, technological development and demonstration under G.A. No 603416 - Project IMPRESSIONS (Impacts and risks from high-end scenarios: Strategies for innovative solutions).

1. Introduction

This work analyzes and models the nature and dynamics of organizational memory, as such an essential ingredient of organizational capabilities. Indeed, the notion of organizational memory stands for an elusive albeit crucial feature of the organizational reproduction of knowledge as distinct from the memory of individuals, namely the ability of organizations to elicit stored information from an organization's history that can be retrieved to bear on present decisions (Walsh and Ungson, 1991). The property of memory of being "organizational" means that, first, it may well be distributed within the organization in ways such that no individual agent or subunit embodies the full representation or the full behavioral repertoires contained in the memory itself. Second, the organizational character of the memory also implies that it is resilient to environmental shocks.

Organizations "remember" because they entail explicit norms and, together, more tacit practices addressed to collectively solve practical and cognitive problems, ranging from the production of a car, all the way to e.g. the identification of a malaria-curing molecule. This is another way of saying that organizations learn, store, elicit and modify over time routines and other "quasi genetic action patterns " (Cohen et al., 1996).

Organizational memory concerns, first, the structure of beliefs, interpretative frameworks, codes, cultures by which the organization interprets the state of the environment and its own "internal states" (Levitt and March, 1988): in brief, call all this the *cognitive memory* of the organization. Second, organizational memory includes routines, comprising standard operating procedures, rules and other patterned actions: call this the *operational memory* of the organization. In short, the two types of memory concern the organizational capabilities to "understand" the characteristics of the environment, on the one hand, and to coordinate particular sequences of actions on the other.

Both cognitive models and operational repertoires are the outcomes of learning processes and thus evolve over time in response to experimentation and feedbacks from the environment. However, they might often entail quite high degrees of inertia and path-dependent reproduction. As a consequence, a major question we shall address below concerns the role of memory in changing environments.

Cognitive and operational memories entail an "if….then" structure. Signals from the environment, as well as from other parts of the organization, elicit particular cognitive responses, conditional upon the "collective mental models" that the organization holds, which are in turn conditional upon the structure of its cognitive memory. Cognitive memory maps signals from an otherwise unknown world into "cognitive states" ("…this year the conditions of the market are such that demand for X is high..."). Conversely, the operational memory elicits operating routines in response to cognitive states ("…produce X…"), internal states of the organization ("…prepare the machine M to start producing piece P…") and also environmental feedbacks ("…after all X is not selling too well…"). In turn, the organizational memory embodies the specific features of what an organization "thinks" and does, and what is "good at", that is its distinct capabilities.[1]

A promising candidate to model both types of memory finds its roots into the formalism of Classifier Systems (CS's) (Holland, 1975, Holland et al., 1986). In a nutshell, a CS is a system of interlinked condition/action rules which evolve according to the revealed environmental payoffs. The model that we shall propose below finds its ascendancy there, and in their application in Marengo (1992), albeit with significant modifications.

In our paper we present a model which links Classifiers Systems and NK fitness landscape models (Kauffman, 1993). The former provide a model of a memory system that accounts for both cognitive and operational memory, while we use the latter to represent an environment in which exogenous environmental traits and organizational actions or policies interact in a complex way to determine the organization fitness or payoff. While in standard NK models (e.g. Levinthal, 1997), cognition, actions and resulting payoffs are folded together in a mapping between "traits" and their "fitness", here we unfold such relationships cognition/action/environmental feedbacks and explicitly model their (evolving) coupling. This is, we believe, a first major advancement with respect to the incumbent literature. Our organization explores a complex and possibly changing landscape in which some dimensions are outside (the environmental traits) and some are within (the action traits) its

[1] Within a very large literature, cf. for instance Helfat et al. (2006) and the critical survey in Dosi et al. (2008).

control. Since the former contribute to determine the payoff of the latter, the organization must base its search over the action landscape on an internal representation (its cognition) of the environmental landscape. When the landscape is complex enough and the organization has cognitive and memory bounds, such an internal representation can only be partial, imperfect and possibly wrong. However through the accumulation of experience organizations can develop better representations that enable them to act successfully in such a complex environment. This is a way to say that organizations painstakingly and imperfectly learn and develop models of their environment. However, there is an exogenous world "out there" which is indeed the object of learning, and which of course is not controlled by the organization. Rather the organization has to learn what to do – the know-how – conditional on (what it believes to be) the characteristics of the landscape mapping the combinations of state-of-the-world and actions into payoffs. This is also another major difference *vis-à-vis* the NK modeling style wherein the "blackboxing" renders all the landscape notionally under the control of the agent. Moreover, the CS formalism allows a straightforward study of learning via *non-local search*, which if undertaken at all in NK frameworks, turns out to be quite arbitrary.

The model that we propose offers a straightforward (and, to our knowledge, novel) formalization of the link between memory and organizational routines. And it is also a promising instrument to explore the double-edged role of memory, conditional on different characteristics of the environment in terms of its complexity and (types of) its dynamics. Memory may crystallize and reproduce the advantages from learning about "good representations" and "good routines" but may also entail "competence traps" (Levinthal and March, 1993), harmful in changing environment. The analysis of the contrasting roles memory plays in different environments is indeed a major task of this work.

We shall proceed as follows. In Section 2 we attempt a broad even if necessarily concise assessment of the state-of-the-art in the incumbent knowledge concerning organizational memory in changing environments. Section 3 presents the structure of the simulation model which addresses the interpretative questions stemming from the foregoing pieces of evidence and explores the dynamics of collective cognition, behavior and ensuing environmental payoff feedbacks. Section 4 discusses the major results we obtain by running the model. Finally, in section 5 we draw the main conclusions.

2. Organizational Memory, Cognition and Routines in Changing Environments: A bird-eye view

The existence and importance of organizational memory is associated with the very ability of organizations to interpret their environment, learn how to solve operational problems and, by doing that, built constructs of knowledge that can be stored and re-used (Argote and Ingram, 2000, Kaplan and Tripsas, 2008).

As already mentioned, one side of the story is in a broad sense *cognitive*. The view of organizations as fragmented and multidimensional interpretation systems is grounded on the importance of collective information processing mechanisms that yield shared understandings (Daft and Weick, 1984), or “cognitive theories” (Argyris and Schon, 1978), of the environment in which they operate, and assist organizations to bear uncertainty, and, as we shall see, environmental and problem-solving complexity. If one subscribes to the notion that organizational learning is a process of refinement of shared cognitive frames involving action-outcome relationships (Duncan and Weiss, 1979), and that this knowledge is retained –at least for some time- and can be recalled upon, this is like saying that organizational learning is in fact the process of building an organizational memory. This cognitive part of the memory is made of “mental artifacts” embodying shared beliefs, interpretative frameworks, codes and cultures by which the organization interprets the state of the environment and its own “internal states” (Levitt and March, 1988).

Together, there is an *operational* side to the organizational memory involving the coupling between stimuli (events and signals, both external and internal ones) with responses (actions), making up a set of rules that remain available to guide the orientation of the organization and execute its operations. In this domain the memory largely relates to the ensemble of organizational routines - patterned actions that are employed as responses to environmental or internal stimuli, possibly filtered and elaborated via the elements of cognitive memory (much more on routines in Nelson and Winter, 1982; Cohen et al, 1996; Becker et al., 2005; Becker, 2005 and the literature reviewed here). As Cohen and Bacdayan (1994) put it, this procedural side is the “memory of how things are done”, bearing a close resemblance with individual skills and habits, often with relatively automatic and unarticulated features (p.554)

In fact, the characteristics and evolution of organizational memory mirrors the characteristics and evolution of organizational routines. In the case of routines, the memory elicits a "relatively complex pattern of behavior triggered by a relatively small number of initiating signals or choices" (Cohen et al. (1996)). How small or big is the initiating set of signals in itself is an important interpretative question, which has to do with the ways the organization categorizes environmental and intra-organizational information. And likewise the behavioral patterns are likely to display different degrees of conditionality upon particular sets of signals. So, at one extreme the action pattern might be totally unconditional and "robust": "perform a given sequence of actions irrespectively of the perceived state of the world". At the opposite extreme actions might be very contingent on the fine structure of their "if" part.

As we shall explore below, it might well be that the coarseness of the "if" and the "robustness" of the "then" parts might well depend on the nature of the environment and its dynamics. A conjecture in this respect is that the more complex and unpredictably changing is the environment, the less contingent is the behavior (Heiner, 1983; Dosi et al., 1999). After all, routines can be seen as an uncertainty reducing device (Becker and Knudsen, 2005; Dosi and Egidi, 1991): robust and largely not contingent routines might be those memorized under highly complex and changing environments. In turn, inertia and path dependence are an almost inevitable corollary of the very existence of organizational memory. The organization is able to recall specific cognitive frames and behavioral repertoires precisely because they are stored and inertially reproduced (possibly with slight modifications) over time. Organizations path-dependently carry with them their birthmarks and what they have subsequently learned throughout their history. It is true that firms typically live in selective environments which tend to "weed out" the most dysfunctional traits and behaviors. However, typically their overall "fitness" (say, their revealed competitiveness) depends upon multiple inter-related traits: in such cases, selection occurs on a fitness landscape with multiple local maxima and with adaptation starting with (random) initial conditions. This holds under NK landscapes (Levinthal, 1997; Rivkin and Siggelkow, 2003, Castaldi and Dosi, 2006) and plausibly even more so in the environments we try to represent here. Indeed, organizations typically compete on such complex landscapes and interrelated technological and behavioral traits are responsible for path dependent reproduction of organizational arrangements (Marengo, 1996, Levinthal, 2000, Siggelkow and Levinthal, 2005).

The ruggedness of the landscape is compounded by the fact that the link between what firms do and the way they are selectively rewarded in the market is utterly opaque for at least three reasons: (i) the complexity of the environments where they operate; (ii) the multiple "epistatic correlations"2 amongst behavioral and technological traits; and (iii) significant lags between organizational actions and performance-revealing feedbacks. In such circumstances, path dependence can also be fuelled by behavioral/procedural and "cognitive" forms of inertia (Tripsas and Gavetti, 2000).

In the model which follows we shall address isomorphic issues by means of simulation exercises and will explore the relationships between the "depth" and inertia of memory and path-dependencies in organizational behaviors.

Organizational memory carries over time what the organization has learned, directly through its past experiences, vicariously by observation of the experiences of other entities, or has been so to speak "brought in" by members of the organization.

Note that learning, on the cognitive side, - unlike most economists' account – does not concern estimations of the value of some variable or even more far-fetchedly updating Bayesian priors, but rather the fuzzy and mistake-ridden categorizations of the signals stemming from the environment. Indeed in recent years the economic literature has tried to extend the standard decision making model in directions which depart from the assumption that decision makers hold a perfectly rational and perfectly knowledgeable model of the world or at least converge to that by acquiring and processing information. Models can now account for non-partitional information structures (e.g. Geanakoplos, 1989), coarse partitions (Mullainathan et al., 2008), sparse representations (Gabaix, 2014), unawareness of some states of the world (e.g. Grant and Quinggin, 2013), but in general try to accommodate such "imperfect" information structures within the rational decision making paradigm.

[2] More on the application of this notion to the economic domain in Levinthal (1997) and Marengo and Dosi (2005); see also below.

Conversely, on the operational side, the very elicitation of some content of the memory tends to improve the latter: learning by doing is a general instance of that phenomenon (on the relation between memory, forgetting, and the learning curve cf. Bailey,1989 and Benkard,2000).

A crucial question regards the usefulness over time of the outcomes of such learning activities as carried by the organizational memory, an issue that boils down to both the characteristics of learning and the depth of environmental changes the organization faces.

Ultimately, if the world is relatively simple and stable, confirming the old beliefs and repeating the "good" routines is an effective organizational behavior. However, the world is rarely simple enough to make experience an infallible teacher (March, 1981; March and Olsen, 1976).

Moreover, quite independently from any possible cognitive bias, the environment may well change in ways that decrease the "fitness" of cognitive and behavioral patterns which were well suited to the "old" environment or even make them detrimental. This is indeed what competence traps are essentially about (Levitt and March, 1988, Gavetti and Levinthal, 2000). Note that competence traps may refer primarily to the cognitive domain or alternatively to the operational one. In the former case the "trap" concerns primarily the reproduction beyond their times of usefulness of previously successful strategic orientations and heuristics. Call them cognitive traps. Conversely, the "operational trap" might concern the "way of doing things" – that is the ensemble of routines and other recurrent action patterns. In these circumstances the remedy is likely to involve also procedural and organizational changes. In actual fact, cognitive and operational lock-ins are likely to come often together.

Bresnahan, Greenstein and Henderson (2011) present an excellent illustration of this point[3] in two cases of "Schumpeterian transitions" across different technological trajectories and of the vicissitudes of the firms which were market leaders under the old regime - in their examples, IBM facing the emergence of personal computers and Microsoft vis-à-vis the arrival of the browser. Take the IBM case. Strong technological capabilities match a commitment to incrementalism in product architectures, cumulative learning, vertical integration, proprietary standards, coordinated strategic

[3] Even if, admittedly, the authors are inclined to offer a somewhat different interpretation of the evidence in terms of economies and diseconomies of scope in presence of jointly shared assets

governance and, on the market side, a reputation for post-sale service. This “IBM model”, as Bresnahan, Greenstein and Henderson (2011) insightfully show, is well aligned to market requirements under the mainframe/mini computer trajectories, but becomes misaligned to the requirements of effective production and marketing of personal computers. It is not that the “raw” capabilities are not there. They are. And in fact IBM even proceeds to a rather successful exploration of a new combinatorics between elements of technological capabilities, organizational set-ups and market orientation well suited to the personal computer world. However, that very success accelerates the clash between the “PC organizational model” and the incumbent “IBM (mainframe) model”. This latter wins and by doing that IBM ultimately kills its PC line of business. It is a story vividly illustrating the path-dependent reproduction of capabilities, shared strategic models, specific organizational arrangements and the ensuing traps. To repeat, it is not that IBM lacked any of the elements underlying successful “PC-fit” combinations. It is just that capabilities, “visions” and organizational set-ups and their specific combinations are better described at least in the short term as *state variables* rather than *control variables*, - in Winter (1987) characterization. Of course, also state variables can and are indeed influenced by purposeful discretionary strategies, that is, by the explicit manipulation of control variables. However this takes time and is tainted by initial birthmarks and subsequent historical paths the organization has taken with respect to both operational repertories and higher level collective visions concerning the very identity of the organization itself.

In fact, technological and market discontinuities, –quite a few analyses suggest-, demand forgetting and unlearning (Hedberg, 1981; Huber, 1991; Nystrom and Starbuck, 1984; Walsh, 1995; Klein, 1989) involving also changes in the organizational structures and the erasing of at least parts of the cognitive and procedural memory of the organization. A revealing example regards the “unlearning” activities involved in the Merger-and-Acquisition processes. Kunisch, Wolf and Quodt (2010) distinguish three domains of possible “misfit” between the two merging organizations - at the level of artifacts, behaviors and corporate cultures. On the ground of a large database on M&A in Germany, they find that cultural misfits are particularly conducive to a lower subsequent performance, while - irrespectively of the sources of misfit – more unlearning is associated with easier absorption of new knowledge and better post-merger performances.

Clearly, unlearning comes at the cost of the loss of a good deal of the experiential wisdom of the organization itself (Gavetti and Levinthal, 2000): however, whether this is actually a cost, in terms of organizational "fitness" is likely to depend upon the depth of the changes in the appropriate technological capabilities and in the market environments. The general intuition stemming from the empirical literature is in fact that the value, or the cost, of cognitive changes and procedural forgetting is a function of the changes in the fitness landscape which the organization faces. Indeed, in the following, we shall explore more formally this conjecture by explicitly modeling shocks on such landscapes and studying the ensuing impacts upon organizational performances under different degrees of cognitive and procedural inertia, or conversely "forgetfulness" of organizations.

3: A model of cognitive and operational memory of organizations

3.1 Formalizing firms as problem-solving organizations

Broadly speaking, the roots of the formalization we present in this paper rest on two complementary classes of models, surveyed at much greater detail in Dosi et al. (2011). The first class includes models mainly addressing learning in complex and changing environments, and focusing on the relationship between learning patterns and ensuing rational performances. Agents are adaptive learners who adjust their knowledge of the environment in which they operate and their behavior (often conflated together into "organizational traits") through local trial-and-error procedures. For this mode of analysis, see Levinthal (1997), Dosi et al., (1999), Ethiraj and Levinthal 2004, Gavetti and Levinthal (2000) Rivkin and Siggelkow (2003), Siggelkow and Levinthal (2005), among others.

The second class includes models focusing upon the relationship between the division of cognitive labor and search process in some problem-solving space, analyzing more directly organizations as repositories of problem-solving knowledge. Here the focus is on the problem-solving procedures which the organization embodies. Indeed, managing an organization, designing and producing cars or software packages, discovering a new drug, etc. can been seen as complicated problems whose "solutions" comprise of a large number of cognitive and physical acts. These kinds of activities imply the coordination of large combinatorial spaces of components. Models addressing such dynamics of problem solving knowledge include Marengo and Dosi (2005), Marengo (1992) and (1996), Denrell, Fang and Levinthal (2004), Valente (2014), Baumann and Siggelkow (2014).

On the output side, components making up an artifact can take a number of alternative states: so, for example, in the case of the production of a car, one combines different characteristics of the engine, alternative designs, different materials, etc. At the same time, innovative search may be straightforwardly represented in form of combination of multiple "cognitive acts" eventually yielding the solution of the problem at hand, e.g. the discovery of a new molecule with the required characteristics, a reasonable and coherent software package, etc. Note that in both examples the existence of strong interdependencies among the components – which often are only partially understood by all agents involved - implies that the effect on the system performance of a change in the state of a single component depends on the values assumed by the other ones. An implication of such interdependencies in this kind of problems is that it impossible to optimize the system by optimizing each single component.

Let us start by considering those (still few) models whereby information-processing and problem-solving activities are represented by ensembles of condition-action (that is, "if…then...") rules.

Marengo (1992) and Marengo (1996) present models which focus upon the modification of such information processing capabilities of individuals or subunits within the organization, i.e. a process of "structural" learning. Agents are imperfect adaptive learners, as they adjust their information processing capabilities through local trial-and-error. This adaptive learning is (at least partly) driven by the information coming from the environment and/or from other members of the organization.

Using a condition-action rule as the basic building-block of this learning system means that the execution of a certain action is conditional upon the agent's perception that the present state of the world falls within one of the categories the agent has defined in its mental model.

Moreover the system must be able not only to select the most successful rules, but also discover new ones. This is ensured, in the above cited models, by applying genetic operators which, by recombining and mutating elements of the already existing and most successful rules, introduce new ones which might or might not improve the performance of the system.

A germane family of models, of somewhat more reduced form but also more elegant and living in a lower dimensional space involves precisely some “black-boxing”, in particular concerning the relationship between organizational traits (including of course behavioral rules) and their actual expressions. Such modeling *genre* prominently includes a family of evolutionary models of organizations inspired by S. Kauffman's so-called “NK model” (Kauffman, 1993). This model of selection and adaptation in complex environments represents evolving entities characterized by nonlinear interactions among their elements, with N the number of elements and K the degrees of interaction among them (their “epistatic correlations”). In Kauffman (1993) the “NK-model” primarily deals with the evolution of populations of biological entities described by a string of "genes" evolving over a fitness landscape, wherein a fitness function is defined assigning a value to each possible string as a measure of its relative performance. One of the pioneering applications of the “NK” approach to organizational analysis is Levinthal (1997). In that simulation model, populations of randomly generated structures (organizations) evolve on a fitness landscape, whereby the evolution is driven by variation, selection and retention processes.

In complex environments the diversity of organizational forms robustly emerges: Levinthal (1997) shows that random local search induces mutations in different directions over the landscape. Moreover, the case of environmental changes can be modeled by re-drawing the fitness contributions of some features after the population has evolved and stabilized over previous optima. If the complexity of the landscape is high, even the modification of the fitness contribution of just one attribute can cause a large alteration of its shape.

Levinthal´s analysis has been expanded and broadened by quite a few works which have further studied the relationship between organizational design and environmental complexity and turbulence. Rivkin and Siggelkow (2002) (cf. also Siggelkow and Rivkin 2006) tackle the issue of multilevel organizational search by introducing an explicit representation of organizational structures in NK-type models. Decisions over the N policies (bits of the string) are allocated among different departments and a superordinate CEO has the function of coordinating departmental decisions.

Gavetti and Levinthal (2000) add a further perspective to the analysis of search processes and look at the relations between forward-looking and backward-looking search and their effects on performance. The roots of the distinction between the two search processes go back to Simon (1955): the former involves cognition-ridden, forward-looking choices based on off-line evaluation of alternatives, even very distant from current behavior; the latter entails experiential choice based on on-line evaluation of a limited set of alternatives which are close to current behaviors. In Gavetti and Levithal's model, the organization chooses a policy on the basis of a simplified and incomplete "cognitive model" of its environment, entailing "templates" which cannot directly prescribe actions. In this context, existing practices function as defaults for elements not specified by the cognitive representation and allow the identification of a specific course of action. Thus, it may happen that actors with the same cognitive template may engage in different behaviors.

An organization which chooses according to its cognitive representation explores regions, and not single points, of the landscape, while the width of these regions depends on the crudeness of the representation. The role of experimental search becomes more and more important as the crudeness of the cognitive representation increases.

Gavetti and Levinthal show that in a context of competitive ecologies in which low performance organizations are selected out, organizations which adopt a joint cognitive and experiential search dominate the population. This becomes particularly evident under rugged landscapes, in which organizations which use purely experiential search are trapped into local optima. In this framework, changes in the representation can enhance organizations' performance when the landscape itself changes as the new representation may identify more effectively new (superior) basins of attraction, and this can compensate for the loss of experiential wisdom.

The model that we present in the following refines upon the first family of models and explicitly addresses the co-evolutionary dynamics between a cognitive domain (the "if's" stemming from the "interpretation" of environmental signals) and an operational one (the "then's"). At the same time such a learning (or unlearning) dynamics is nested upon, and ultimately driven-by, fitness landscapes of the NK type –characterized by different degrees of ruggedness and generally changing overtime.

3. 2 The Model

3.2.1 An informal description

Coherently with the introductory remarks, we build a simple model of organizational cognition and action, where past experience is stored in a repertoire of condition-action rules, broadly inspired by John Holland's Classifiers Systems (Holland et al. 1986). Such a repertoire, together with an indicator of each rule's past usefulness constitutes the organizational memory. Rules embed a "know-what" component in the condition part, i.e. the capability to make some sense of the environment and distinguish different situations, and a "know-how" component in the action part, i.e. the capability to perform an appropriate action once a situation has been detected. Such a distinction is, with different nuances, common in the literatures on organizational cognition as well as on organizational routines and is germane to the distinction between "declarative" and "procedural" memory (Anderson, 1983), Cohen and Bacdayan, 1984, Miller et al., 2012).

Each rule takes the form of "**if** a given set of conditions is detected – **then** a certain action pattern is performed" and can be therefore characterized by their degree of generality vs. specificity, according to the size of the set of conditions to which they apply. General rules prescribe the same course of action for a broad range of environmental conditions, while specific rule apply only to one or very few situations. General rules may reflect different phenomena: a) ignorance, i.e. the organization does not know what to do in different situations and therefore applies the same rule to a wide range of conditions; b) inability to discriminate environmental conditions, which leads the organization to consider as equivalent situations which differ; c) routinization, i.e. a conscious or unconscious definition of relatively invariant rules which apply to ensembles of environmental conditions, either because the organization is not capable of producing more specific rules for sub-ensembles or because the cost of finding such more specific rules is higher than the potential benefit they could deliver; d) conscious generalization, i.e. the organization deliberately reckons that a broad range of situation must be treated as equivalent for action purposes.

In our simulations we will suppose that the organization starts with one fully general rule, whose condition part indicates it may be applied to any possible environmental condition, and a random action part. An adaptive mechanism, based on the feedback received when a rule acts on the environment, generates new rules as local modifications of the existing ones. Such local adaptive changes may involve both the condition part (increasing or decreasing its specificity) or the action part (by mutating of its bits). Each rule is assigned and "strength" parameter, roughly measuring its past effectiveness. The strength parameter is used to decide which rule to apply when more than one rule satisfy the environmental conditions. In this case, the strongest rule, among those which satisfy the current environmental conditions, i.e. the one that has proven more successful in the past, will be preferred for action.

Rule strength also governs the novelty generation mechanism, as stronger rules will be preferably chosen for the generation of "offspring" variant rules (i.e. new rule which are copies of the stronger ones but with some small mutations in the condition and or action part).

In our framework, the repertoire of condition-action rules represents the memory of the organization, and its size is given by the number of different rules held in this repertoire. The strength of a rule is updated every time the rule is chosen to act on the environment of the organization, and it is updated according to the payoff received by the action. Rules that are active because they satisfy the current environmental conditions, but have not been chosen for action, have their strength reduced. The system records for each rule an indicator of "inactivity", which keeps track of how frequently a potentially active rule has not been chosen for action. A rule is removed (and therefore "forgotten") when the inactivity indicator reaches a given threshold. Thus, by tuning this threshold we control for the trade-off between remembering and forgetting. Other things been equal, the higher this threshold (indicating higher tolerance for inaction) the larger the size of memory, since rules will be deleted less frequently, while a lower the threshold will reduce the size of memory as fewer rules will manage to survive the selection on inactivity. Notice that the size of memory is endogenous because, as we will show below, it depends on the all features of the environment, which influence how many rules are compatible with the environment at each time.

The overall dynamics of the memory is therefore history-driven (among all the rules which apply to the current situation the one which has been more successful in the past will tend to be preferred), but also cognition driven (only rules whose condition applies to the current situation can be used, in other words rules define a set of categories in which environmental states are classified), and also variation driven (novelty is constantly introduced as variation on existing rules).

We test the behavior of such a system in different environments, characterized by varying degrees of complexity and volatility. We assume that both environment and actions are multidimensional objects and that the complexity of the problem the organization faces is determined by the interdependencies among the elements forming the environment, among the elements composing the action and across the two elements, environment and action. In other words, the organization is placed in a "NK landscape" *à la* Kauffman (1993), but the N dimensions of the landscape belong to two different categories. $N_e<N$ dimensions are environmental features, which the organization cannot control or modify but can only observe and (try to) categorize according it its set of conditions. The remaining $N_a=N-N_e$ are instead dimensions (policies) pertaining the action of the organization and chosen by the latter according to it repertoire of rules. The payoff for the organization will, in principle, be determined by the current configuration of all the N dimensions, thus in principle each configuration of the N_e environmental dimensions determines a different landscape for the N_a action dimensions. However we can, and we will in the simulations below, test the behavior of our organization in landscapes characterized by specific structures of interdependencies, using the methodology presented in e.g., in Frenken et al., 1999. To simplify, we will suppose that, while indeed all the environmental elements contribute to determine the payoff of an action, only a subset of them modify the shape of the action landscape, while the other environmental dimensions only determine a shift of the payoff values, but no change in their relative value. We call the dimensions which together modify the shape of the action sub-landscape **core** dimensions (or **core bits** in our simulations where all dimensions take only binary values). Thus, the ranking of actions (from the most to the least fit) does not change when the core dimensions remain constant and the non-core ones change, although their fitness value does change. When instead a core dimension change, in general also the ranking of action will undergo random changes.

The organization must therefore learn to discriminate between core and non core bits, in spite of the fact that all of them cause changes of the fitness values of actions, and, possibly develop specific rules for each configuration of the core dimensions, prescribing a different action to each of them.

Also the action part of the landscape may be more or less complex. In an action landscape of complexity K_a the payoff contribution of each action bit depends upon K_a-1 other action bits (besides depending on the environmental bits as described above). Thus when K_a=1 we have a simple action landscape (for any configuration of the core bits) where the payoff contribution of each action bit is independent from the current value of the other bit; while as K_a grows the action landscape becomes more and more complex and uncorrelated.

It is worth stressing once more this fundamental difference between our model and the usual NK fitness landscape model which is, we believe, one of the significant original contributions of this paper. We assume that the landscape is made of both exogenous and endogenous components. Both contribute to determining the fitness of the organization, in tune with familiar representations, depending on the complexity structure, but only the latter are under the control of the organization while the former are exogenously determined by what we call “the environment”, which of course may well include other organizations or past actions of the same organization itself. Thus exogenous components modify the landscape of the endogenous ones and the search process on the latter must be based on some cognition of the former. Our if-then rules are a simple (and already widely used in the adaptive learning literature) way to model an adaptive system that conditions its action upon a categorization of the exogenous states. In short, our organizations must discover both the correct categorization of the environmental events and, for each relevant class of events, the appropriate action. The only available information is the feedback received by the action actually performed at each step (i.e. “on-line” learning), and we explore the results under different settings allowing for different sizes of memory for the organizations.

Concerning the environmental dimensions, we simulate stationary and non stationary environments. In the former case the environment-action landscapes are generated (with the desired complexity structure) at the beginning of the simulation and never changes. The state of the environment (the configuration of the N_e environmental bits fed to the organization) changes at each moment time, but the mapping between each environmental state, actions and payoff remains constant. On the contrary, in a non stationary environment the payoff values are subject to change, and therefore the shape of the landscape is modified, although the structure of the interdependency links remains constant so that the relevant categories do not change.

In the next subsections we will present the details of our model more precisely and formally and then, in section 4, we will examine its behavior in different environments.

3.2.2Environment, states of the world and payoffs

The **environment** is fully described by a set of n elementary "states" $E = \{e_1, e_2, \ldots, e_n\}$. For simplicity we assume that each environmental feature may take only two values $e_i \in \{0,1\}$.

Organizational behavior is characterized by an action vector made of m elementary acts $A = \{a_1, a_2, \ldots, a_m\}$. Again, for the sake of simplicity, we assume $a_i \in \{0,1\}$.

Payoffs or fitness: in general, the payoff for or the fitness (we will use both expressions indifferently) of the organization depends upon the entire profiles of organizational acts and environmental states. The payoff function is described as $\pi : E \times A \mapsto [0,1]$. We explore different *complexity structures* concerning the mapping from the $E \times A$ space to the pay-off. There are potentially three sources of complexity, namely those due to (i) interdependencies among environmental states, (ii) interdependencies among elementary acts and (iii) interdependencies among environment conditions and action patterns.

Task complexity structure: we assume an environment where some environmental features interact with the organizations' actions to determine the payoff, while others do not. More precisely, we define as core environmental components those which influence the payoffs of different ensemble of actions and also the ranking of different action profiles. Conversely, non core traits are those environmental components which influence the payoffs received by different actions, but not their ranking. Hence, suppose that the vector $[a_1,a_2,..a_m]$ is the optimal action when the environment is described by vector $[e_1,e_2,..e_n]$. Then if e_i is a non core bit, a change of its value will affect the overall payoff, but will not change the corresponding optimal action (nor of the ranking of all other actions), while if e_j is a core bit a change of its value will in general determine a change of the corresponding optimal action as well as of the ranking of all other actions.

The crucial task for a learning organization is therefore to discriminate "action-relevant" or "core" environmental signals and "understand" how they interact with the set of elementary acts which make up the action. Of course, the number of "core states" is a measure of environmental complexity. Note that learning is only driven by payoff, which is the only signal organizations receive on how good their actions, coupling the conditions, are. The complexity of the environment is revealed by a NK-like fitness function, based on correlations among state yielding specific correlations between "cognitive" and behavioral traits of the ensuing fitness.

3.2.3 Organizational cognition and action

The task of the organization is to develop the capability of correctly detecting states of the world and choosing the appropriate behavior. In order to do that, the organization stores a set of cognition–action rules that together constitute the organizational memory and action repertoire. These rules constitute a Classifiers System (Holland et al., 1986) that performs the two interrelated tasks of detecting and memorizing environmental regularities (i.e. partitioning environmental states into categories) and applying the appropriate course of action to each of them. Condition-action rules are "*if... then*" rules that map detected environmental profiles into action. Each rule takes the form:

$$c_1, c_2 \dots c_n \Rightarrow a_1, a_2 \dots a_m \text{ with } c_i \in \{0,1,\#\} \text{and } a_j \in \{0,1\},$$

where # stands for "do not care".

Each rule is characterized by its specificity σ_i, i.e. the number of its condition bits which are different from $\#$, and is assigned a strength, $S_{t,i}$, which is an indicator of the payoff it has cumulated, minus possibly some "tax" (the details will be given below). If the current environmental state matches the condition part of a rule, i.e. if either $e_i = c_i$ or $c_i = \#$, then the rule is considered as active.

We experiment with an on-line set-up of rule selection, whereby, if more than one rule is active, they bid for action. The bid of a rule i that is active at time t is denoted $B_{t,i}$ and computed according to the formula:

$$B_{t,i} = \left[\beta + (1-\beta)\frac{\sigma_i}{Ne}\right] S_{t-1,i} \tag{1}$$

Where β is a parameter representing the relative importance of specificity in bidding, σ_i is the specificity of the condition, and N_e is the number of environmental bits. In summary, a bid is higher the higher is its strength and, weighed by the β parameter, its specificity. This reflects the principle of *default hierarchies* (Holland et al., 1986), that is, more specific rules, other things being equal, should be preferred to more general ones.

The bid is computed for all the active rules, then one and only one of the latter is chosen for action with probabilities proportional to the individual bids.

After the selected rule has acted and the corresponding payoff has been observed, the strengths of all rules are updated, with different formulas, depending on whether the rule has been chosen for action or not.

For inactive rules (not chosen) the strength of the rule is updated as:

$$S_{t,i} = S_{t-1,t}(1-\tau_1) \tag{2}$$

where τ_1 is an "inaction" tax (i.e. a depreciation).

The selected rule pays an additional "tax" proportional to the proposed bid, and receives as reward the payoff:

$$S_{t,i} = \gamma\left[S_{t-1,t}(1-\tau_1) - \tau_2 B_{t,i}\right] + (1-\gamma)\, Payoff_{t,i} \tag{4}$$

τ_1 captures the cost of "maintenance" that any rule has to bear, while τ_2 is a scale parameter for the cost of bidding. Finally, γ allows to tune the "speed" by which a payoff affects the strength of rules.

At the start of each simulation run we assume no knowledge of the environment: the organizational memory contains only one rule, whose condition part is formed only by $\#$'s (reflecting a state of total ignorance), while the action part is a randomly drawn binary string.

3.2.4 Generation of new rules

New rules are normally generated as variation of existing successful ones. New rules can be obtained by either specification or generalization of the condition part and/or by mutation of the action part.

Specification means that a rule with some #'s in the condition part generates an offspring rule whose condition part is a copy of the parent's except for the mutation of some #'s into either a 0 or a 1. Such mutation is controlled by two rule-specific indicators defined as follows: when a rule is selected for action, for each un-specified bit h of the condition part (i.e. each $c_h = \#$), one defines two indicators $r^0_{t,h}$ or $r^1_{t,h}$ measuring the payoffs obtained conditional on the environmental states. If the environmental string has 0 in the h-th position, i.e. $e_h = 0$ then $r^0_{t,h} = r^0_{t-1,h} + Payoff_t$, where $Payoff_t$ is the payoff received by the rule. If instead $e_h = 1$, it is $r^1_{t,h}$ that gets accordingly updated: $r^1_{t,h} = r^1_{t-1,h} + Payoff_t$. In short, the two indicators collect the sum of the payoffs received by the rule when the environmental bits corresponding to each of its # were 0 or 1, respectively. In order to decide if the new rule is worth to generate via specification, the system computes the following indicator:

$$\iota_{t,h} = \frac{| r^0_{t,h} - r^1_{t,h} |}{r^0_{t,h} + r^1_{t,h}} \tag{5}$$

This is an indicator of concentration, ranging from 0 (minimal concentration) to 1 (maximal). The system will generate a new rule via specification if at least one $\iota_{t,h} > K^{\iota}$, where K^{ι} is a threshold parameter which can be varied across experiments.

If the rule has to generate a new offspring via specification, the system chooses which bit must be turned from a $\#$ to either a 0 or 1 by assigning each generic bit the following probability:

$$prob(t,h) = \frac{\iota^{\theta}_{t,h}}{\sum_{j=\#} \iota^{\theta}_j} \tag{6}$$

where $\sum_{j=\#}$ indicates the sum over all and only the condition bits equal to $\#$ and θ is a parameter affecting the concentration of probabilities. The chosen bit is set to 0 if $r^0_{t,h} > r^1_{t,h}$ and to 1 otherwise. If there is no bit fulfilling specification conditions, the system generates a new rule either via generalization or via mutation, choosing randomly with equal probability between the two alternatives. The action part of the specified offspring rule is instead a perfect copy of the parent's action part.

On the contrary, generalization consists of mutating one specified bit $c_i \neq \#$ into $c_i = \#$. Also in the case of generalization, the action part of the offspring rule is a perfect copy of the parent's action part.

Conversely, strings produced via action mutation maintain the same condition part c_i and switch one, randomly chosen, bit of the action part from 0 to 1 or vice versa. Furthermore, a more general rule can be introduced in the system when no existing rule can be applied because no condition part matches the current state of the environment. In this case, for all existing rules a "mismatch ratio" is computed. This is the ratio of the bits that are specific (i.e. equal to 0 or 1) and do not match the environment (i.e. $c_i \neq e_i$) divided by the number of specific bits. Thus, for instance, a rule with only one specific bit (and $\#$ everywhere else in the condition), when this bit does not match the corresponding environmental one has a mismatch ratio 1. A rule with one mismatching bit, but having other 9 specific and matching bits has a mismatch ratio 0.1. The rule with the lowest mismatching index is generalized by turning the mismatching bits into $\#$.

3.2.5 Forgetting rules

Each rule is associated to an indicator reporting how frequently the rule is chosen when its condition part is compatible with the environment, i.e. it has the chance to be chosen. This indicator $A_{t,i}$ is not modified when the rule is not applicable because its condition part does not match the environmental bits. Otherwise, it is updated as follows:

$$A_{t,i} = \lambda A_{t-1,i} + (1-\lambda)\, C(t,i)$$

where λ is a smoothing coefficient and $C(t,i)$ equals 1 if the rule i is chosen at time t and 0 otherwise. Thus, $A_{t,i}$ will approach 0 when the rule is rarely used, and 1 when, on the contrary, it is frequently used. At every time step the organization reviews all the rules in its repertoire and removes those whose indicator $A_{t,i}$ is below a given threshold. Consequently, the lower the threshold the more rules will remain in the memory of the organization, while the higher the threshold the more selective is the organization, retaining only rules used frequently.

4. Some Results[4]

We implemented the foregoing model and explored its results under different configurations. As in most agent-based models, there is a relatively large number of parameters affecting the results, and we can obviously explore only a small fraction of the parameter space: our experiments have been focused on the analysis of the revealed fitness values achieved by our artificial organizations under different memory and learning conditions (e.g. possibility or not of forgetting), conditional on different degrees of environmental complexity and patterns of environmental dynamics.

In all our experiment we consider an overall landscape made of N=14 elements, divided into N_e=9 environmental bits and 5 action bits "operationally" controlled by the organization. In each simulation run we define a given environmental setup (e.g. number of core bits, type of complexity, etc.), and one or more "populations" of organizations. Each population is a group of independent organizations with identical initial conditions and learning set-ups. At each step of a simulation run the environment is determined by its stochastic "law of motion" (if any), and organizations have to choose one rule from their repertoire. Next, given the revealed payoff, organizations update the strength of their rules, and, possibly, generate and/or remove rules from the repertoire.

For each organization we compute some statistics on its performance and properties of the repertoire of rules. The performance of an organization is measured on the basis of the fitness produced by combination of the current environmental state (common to all organizations in the model) and the action part of the rule selected by the organization for that time step. In order to allow for comparison between landscapes with different levels of complexity we report the results on the values of the relative fitness, which is the ratio of the organization's fitness divided by the highest fitness attainable with the current environment, i.e. the fitness pertaining to the optimal action. It is well known in fact that the maximum fitness value of a NK landscape depends on the value of K, thus

[4] The simulations presented here were implemented in the "Laboratory for Simulation Development"(LSD), a simulation platform developed by one of us (Valente, 2008). Code of the model and configurations used are available from the authors upon request.

if we want to compare the performance across landscapes with a different K value we must use relative fitness values instead of absolute one.

Together with the performance, we compute also the number of rules in the organizations' repertoires - as a proxy for the size of memory -, and the specificity of each rule, i.e. the number of bits in the condition part of rules whose value is not #.

4.1 Learning in stationary environment.

Let us begin by analyzing the behavior of our rule system for varying degrees of environmental complexity, in stationary environments, i.e. characterized by an stable unchanging landscape. We consider four levels of environmental complexity, measured by the number of core bits, from simple (1 core), to intermediate (3 cores), complex (6 cores) and maximally complex (9 cores) environments. For each of these settings we simulate three populations made of 100 organizations, each defined as having low, intermediate and high memory levels respectively, determined by the threshold used to remove the less often used rules.[5]

Simple environment (1 core)

Let us start with the simplest environmental set-up, with only one core environmental bit. Figure 1, 2 and 3 show the average fitness, specificity and number of rules across time for the three sizes of memory, respectively.

[5] For reasons of space we report here only a brief summary of the results. Detailed statistical analysis, along with the computer program are available from the authors upon request.

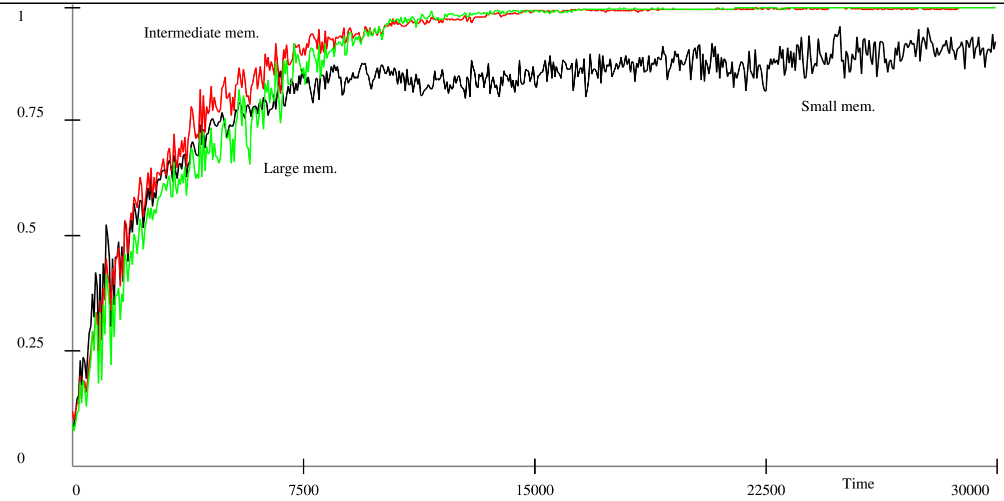

Figure 1: Average fitness over 100 organizations; simple (1 core) environment.

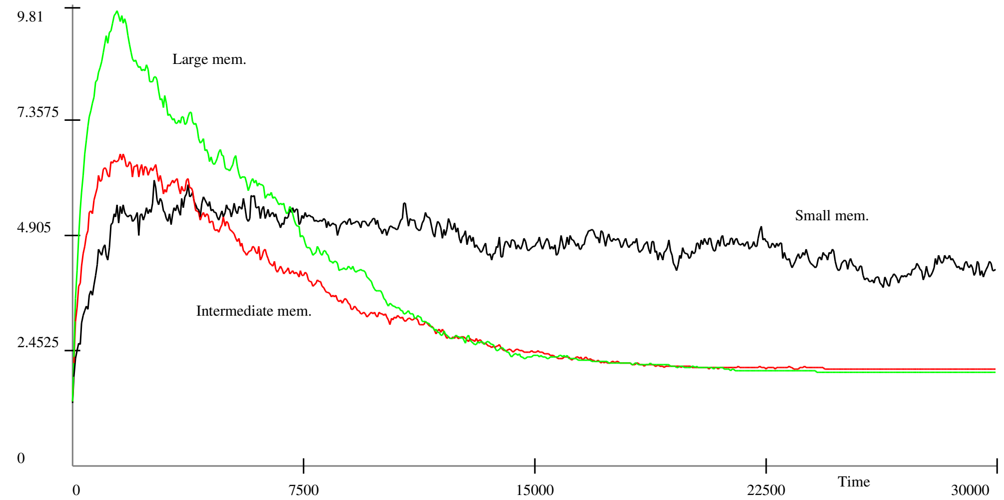

Figure 2: Average number of rules over 100 organizations; simple (1 core) environment.

Figure 3: Average specificity of rules over 100 organizations; simple (1 core) environment.

In this setting the “optimal” repertoire consists of two rules made of all #’s but for the core bit (hence specificity 1), each applying to either state of the core bit, and containing the appropriate action. The figures show that all the three populations manage to sensibly increase fitness through time, although only the organizations with large or intermediate memory size consistently reach the maximum fitness levels. Figure 2 shows that the pattern to learning, in this setting, consists in generating a lot of rules (almost 10, for the large memory population), which are then selected over, until the two populations with large and intermediate memory sizes reduce them to the only 2 required. Organizations with small memory size remain instead stuck with a larger number of rules than that strictly required for optimality. Figure 3 offers the intuition on the reasons underlying this pattern. All organizations start their learning by rough and imprecise over-categorizations[6] of the environment, quickly reaching levels of specificity much higher than 1, in search of the bits of the environment that matter in terms of action. While refining the action part, they also merge categories, therefore reducing the average specificity of rules. Organizations with too small a memory, however, cannot perform this winnowing part of learning because more general but imperfect rules are discarded too quickly, and therefore continue to generate over-specific, and sub-optimal ones.

Already this simple case highlights the tension intrinsic in the role of the memory. Less memory implicitly demands tighter selection - and thus, in a naïve reading, a sharper learning -, but in turn tighter selection entails deeper tradeoffs between exploitation and exploration, well in tune with March (1991).

Intermediate complexity (3 cores)

Let us increase the complexity of the landscape to 3 cores. In this setting we observe results broadly similar to the previous case. However, with 3 core bits the optimal repertoire generally entails a different rule for each combination of states, that is $8=2^3$ rules. Though requiring longer time span, we obtain the same outcomes as in the previous exercise: organizations with too small memory fail to reach the highest fitness, while the other two populations follow the same pattern of initial over-specification, discovery of the right categories, and identification of optimal actions for each category.

[6] Incidentally, note that these "categories" are not partitions: their intersection is not the empty set.

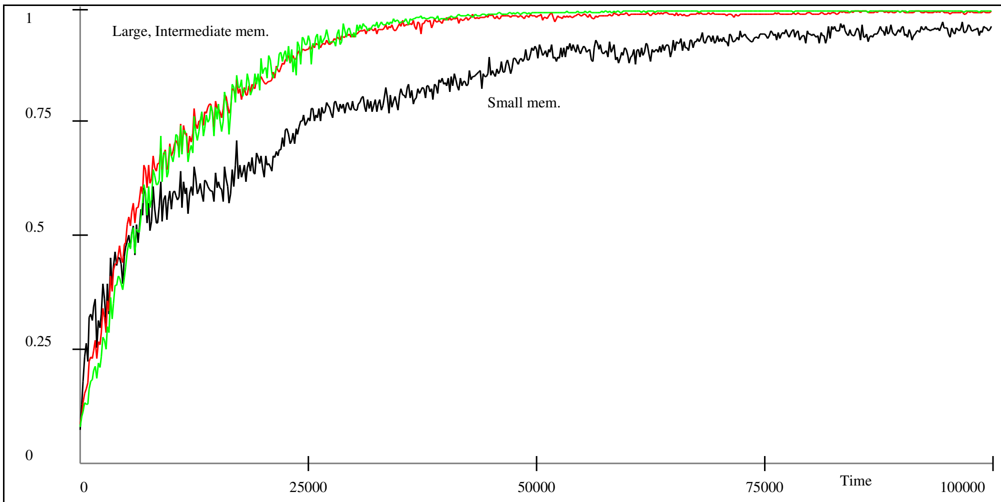

Figure 4: Average fitness over 100 organizations; intermediate (3 core) environment.

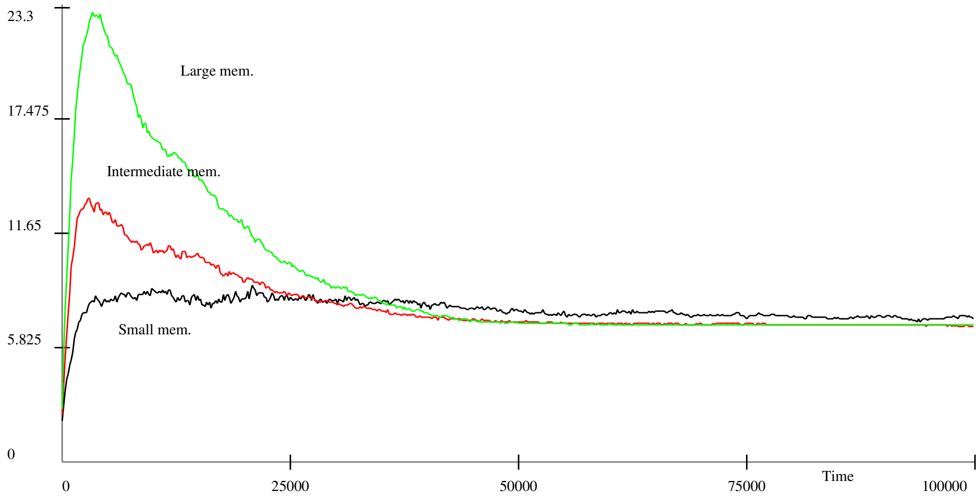

Figure 5: Average number of rules over 100 s; intermediate (3 core) environment.

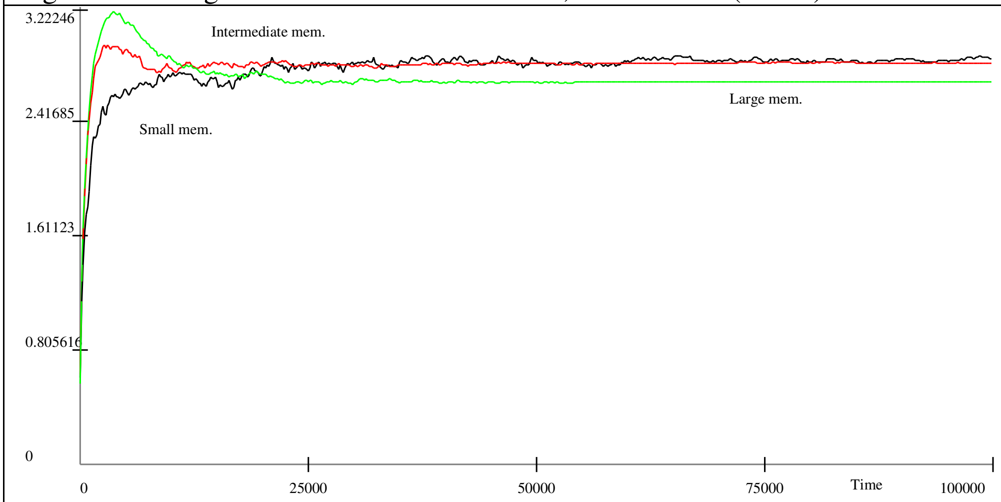

Figure 6: Average specificity of rules over 100 organizations; intermediate (3 core) environment.

The most notable difference with respect to the previous simpler environment is that the gap between the organization with smaller and larger memory is narrower here than in the simpler case, contrary to what one might expect. It seems that errors due to “premature learning” under very small memory are more costly when (nearly) optimal behaviors entail fewer actions to be fired more frequently. On the contrary, in more complex environments there are more chances to get it roughly right in at least some of the (many) possible relevant environmental conditions, when apparent “disfunctionality” -at least as defined on the grounds of short term reinforcement- is kept alive.

Complex environment (6 cores)

With landscapes of higher complexity, our artificial organizations fail to reach systematically the maximum fitness. The specificity statistics show that organizations fail to map correctly the relevant “true” environmental categories but rather develop cross-cutting categorization/action routines. Moreover, for a long initial period (notice the difference in time scale with respect to the previously discussed results), intermediate memory sizes show a consistently better performance than larger ones, hinting that excessively large memory starts to be a liability, rather than an asset.

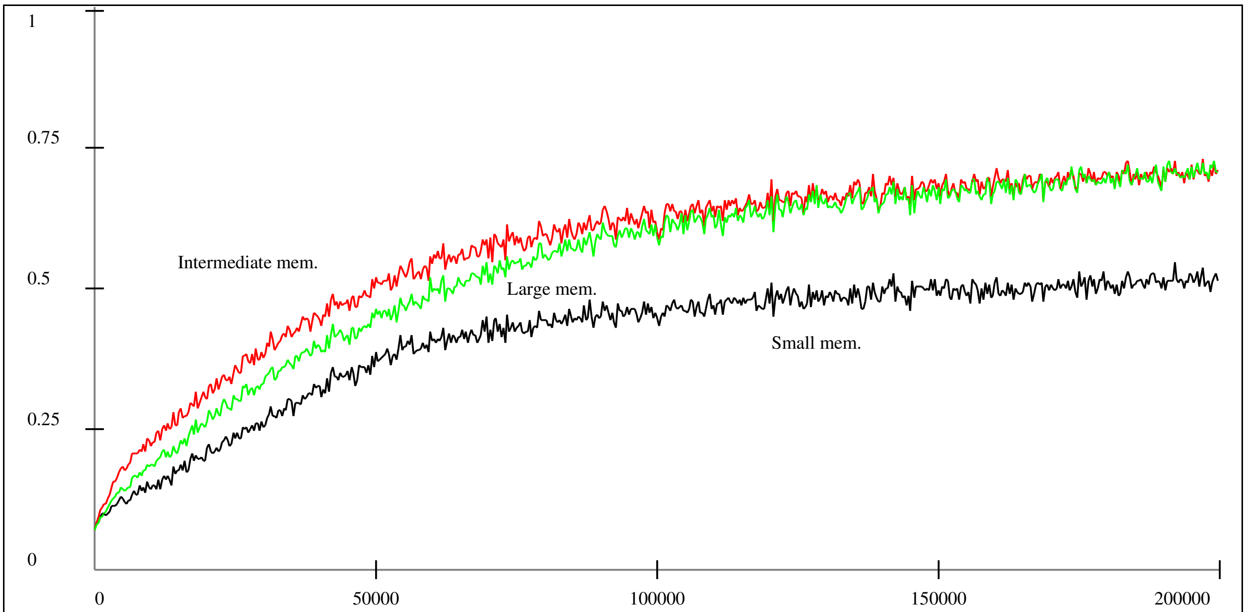

Figure 7: Average fitness over 100 organizations; complex (6 core) environment.

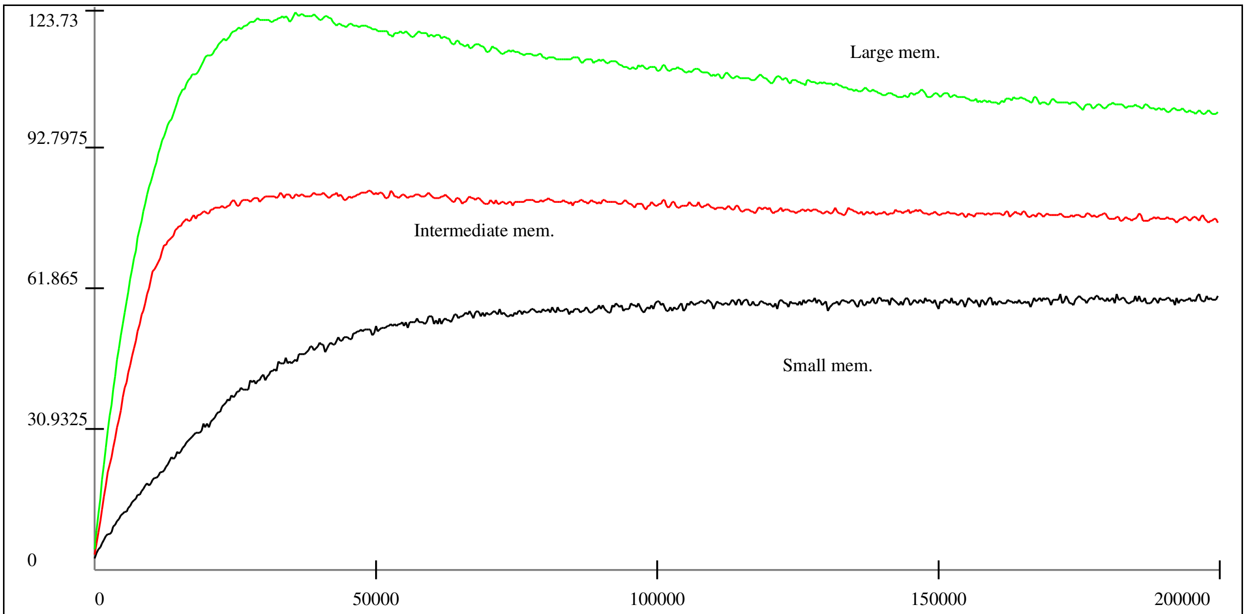

Figure 8: Average number of rules over 100 organizations; complex (6 core) environment.

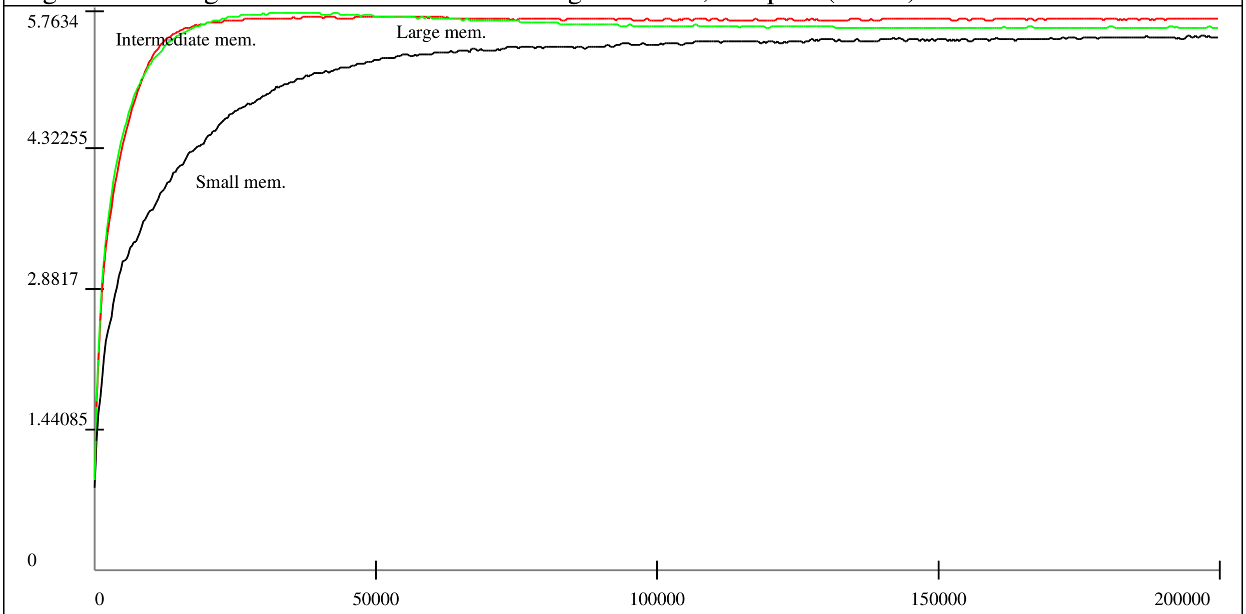

Figure 9: Average specificity of rules over 100 organizations; complex (6 core) environment.

Maximally complex environment (9 cores)

Pushing still higher the environment complexity (all the 9 environmental bits are core bits) yields somewhat different dynamics. Even under such very high levels of complexity, organizations learn and manage to improve their performance. However, the cost of an excessively large memory size becomes all the more apparent: the average fitness of the intermediate memory size is persistently and increasingly higher than that achieved by organizations with larger memory.

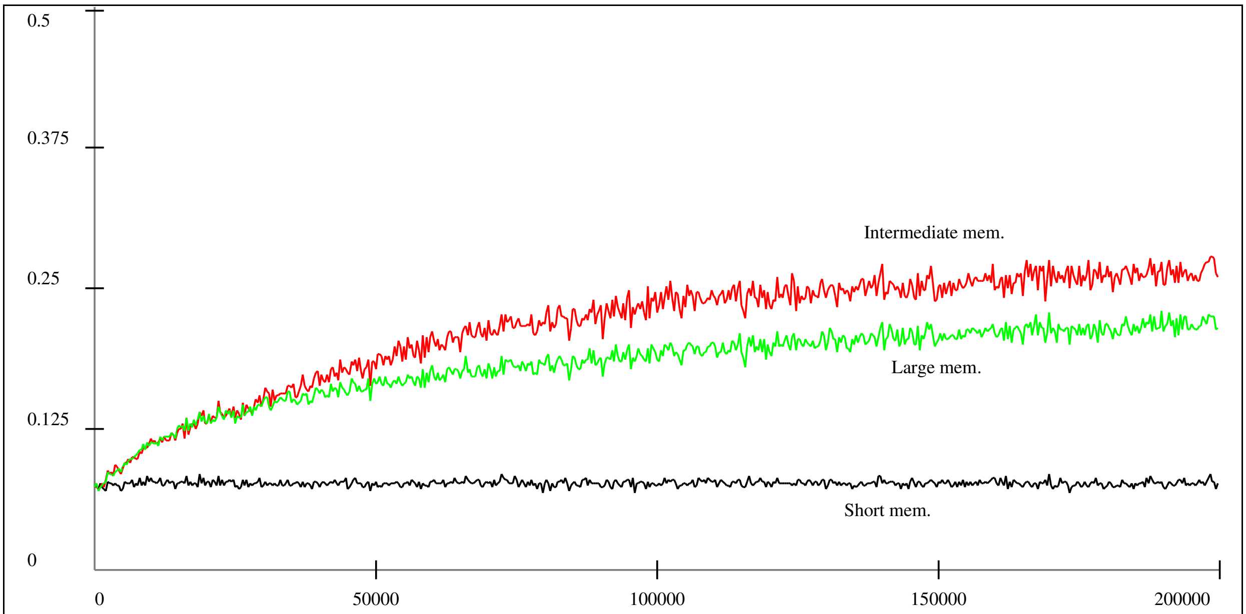

Figure 10: Average fitness over 100 organizations; highly complex (9 core) environment.

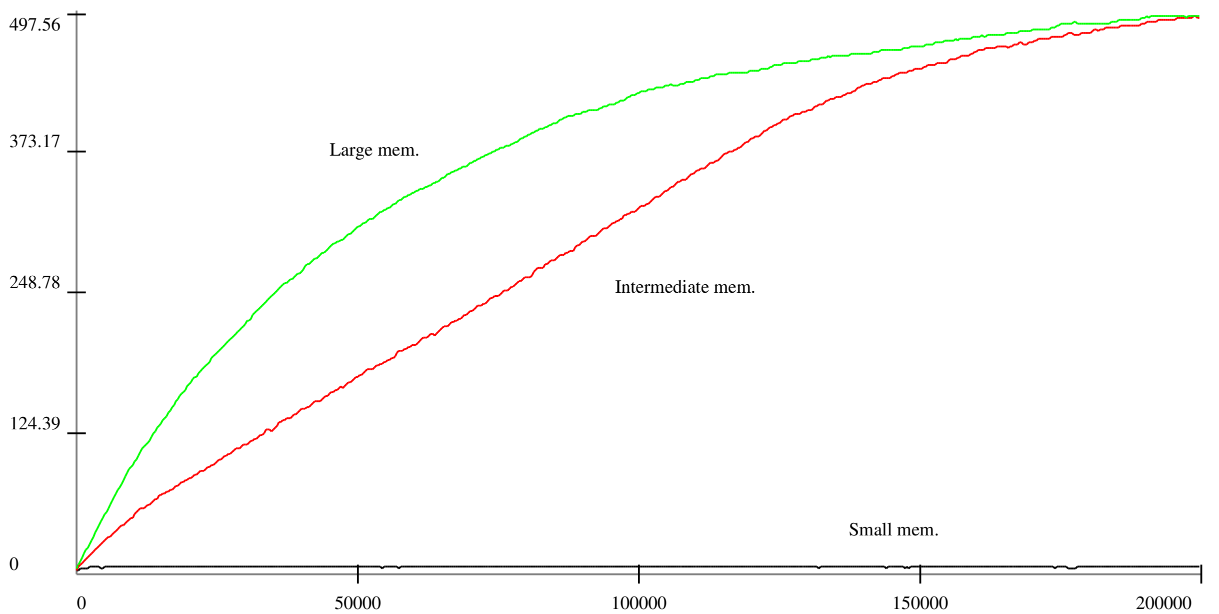

Figure 11: Average number of rules over 100 organizations; highly complex (9 core) environment.

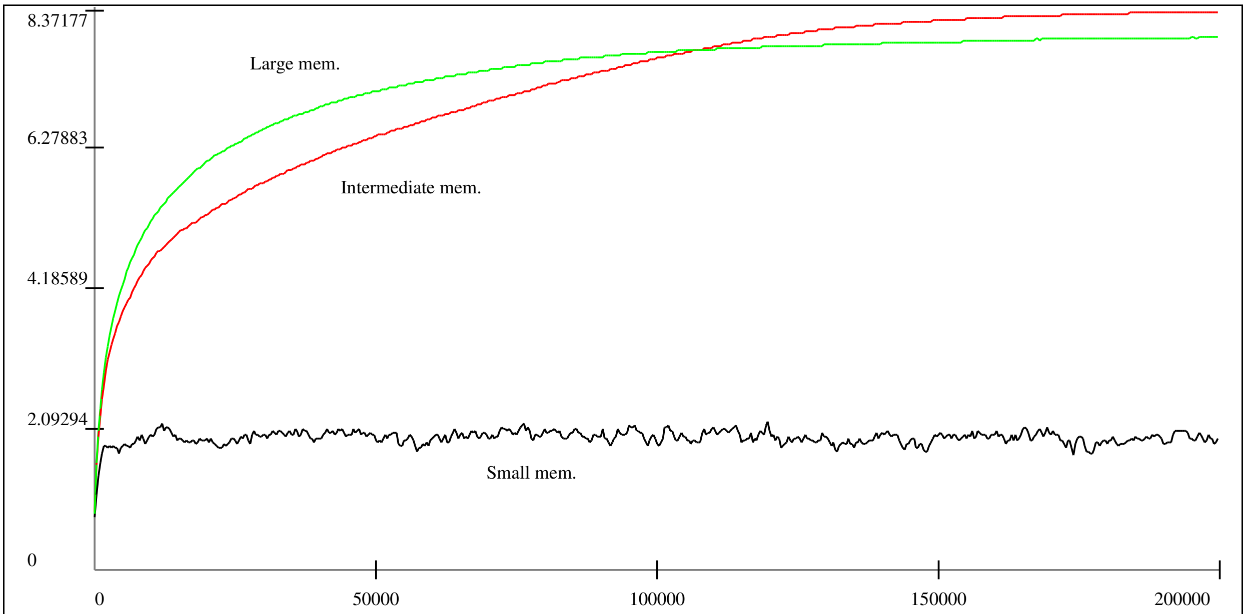

Figure 12: Average specificity of rules over 100 organizations; highly complex (9 core) environment.

4.2 Non stationary environments

So far we considered stationary environments, in which a good rule, if discovered, remains good forever. Let us now consider environments which persistently undergo regular shocks generating each time a new landscape, i.e. a novel mappings between states of the world, actions and payoffs. Shocks occur at regular intervals (we test for different frequencies) and imply that all the fitness values of the landscape are re-drawn, though the complexity structure of the landscape is kept constant. Therefore the performance of incumbent rules, associating a given action to some conditions of the environment, is suddenly and abruptly modified.

In these circumstances, one is able to explore also whether forgetting current (and suddenly at least partly obsolete) knowledge may be desirable in response to a radical shock. In order to do that, we compare the performance of two alternative learning strategies for each memory size. In the first setting organizations maintain their entire set of rules: in fact they do not "know" of the shock, but painstakingly amend the existing knowledge facing the new landscapes. Conversely, in the second setting organizations "reset" completely their memory in response to shocks, starting the learning process from scratch (i.e. from a single rule with randomly chosen action). In fact, this is equivalent to evaluate comparatively the performance of new organizations (without any memory from the past) at each shock. In the tables below we report the results for the average fitness generated by organizations exploring a increasingly complex landscapes: 1, 3 and 6 cores[7]. The rows correspond to different frequencies of shocks (as reported in the first column). The columns report the size of memory and whether the organizations reset their memory or not. The last raw, that we call ”Never”, is the equivalent of the previous case for stationary environments (with the small technical difference that for sake of comparability we “force” exploration for new rules to continue forever, thus undermining a bit the fitness).

Simple environment (1 core)

	Small	Medium	Large	SmallReset	MediumRes.	LargeRes.
300	0.12566	0.0989861	0.0919705	0.10344	0.0837442	0.0606114
600	0.189397	0.139157	0.115108	0.165856	0.132492	0.120285
1000	0.243851	0.206673	0.157258	0.228764	0.187792	0.145979
1500	0.314622	0.253612	0.237312	0.292384	0.265688	0.229637
2000	0.368529	0.345406	0.308086	0.359418	0.315149	0.26155
2500	0.410609	0.4001	0.351153	0.403276	0.368321	0.33263
3000	0.484363	0.457313	0.451933	0.47254	0.442825	0.378531
5000	0.580595	0.573387	0.538362	0.518848	0.561429	0.496277
10000	0.69922	0.770904	0.763659	0.690147	0.761173	0.676752
20000	0.748382	0.770636	0.717614	0.752883	0.80075	0.731601
Never	0.803684	0.942152	0.931624			

Tab.1 Simple landscape (1 core). Average fitness over 50,000 time steps starting from t=150,000.

Results show that in highly volatile environments, with high frequency of shocks, small memory size provides higher fitness. In these cases a large memory allows the survival of inefficient rules that spoil both learning patterns and average performance. The advantage disappears as the shocks become less frequent providing the opportunity for organizations with larger memory to better deploy their learning potential. In any case, organizations with the largest memory size still pay a price in terms of performance, supporting the hypothesis that in simple but volatile environments a large memory is not only redundant, but is effectively negative on performance. The reason is that a larger

[7] We skip the analysis of maximally complex landscapes (9 cores) because the results are perfectly in line with the other cases.

stock of experience takes more time to be replaced than a smaller one, allowing obsolete knowledge to linger longer within the organization dragging down its performance.

The average performance of organizations which reset their memory after a shock is no better, and in many cases markedly worse, than the one of organizations with equivalent memory size which maintain their (obsolete) experience after a shock. This result hints to the reasons for the advantage of smaller memory sizes. Intense competition for scarce memory speeds up learning, favoring recycling of useful chunks of existing knowledge, as opposed to generating new knowledge from scratch in a context with abundant remembering capacity. This is why the first column (small memory, keep obsolete rules after a shock) shows persistently better performance than the last column (largest memory, delete all rules after a shock). The advantage is clear at all frequencies, but for the most rare, in which performance value are very similar.

Intermediate complexity (3 cores)

With intermediate complexity we obtain similar results (see Table 2), though, of course, performance is generally lower due to the increased complexity of the landscape.

	Small	Medium	Large	SmallReset	MediumRes.	LargeRes.
300	0.107307	0.0862687	0.0880694	0.0949113	0.0872966	0.082382
600	0.129952	0.0975477	0.0956609	0.11328	0.0980064	0.0867317
1000	0.153389	0.103191	0.106947	0.137083	0.121099	0.109123
1500	0.197236	0.142404	0.128115	0.161011	0.155791	0.132149
2000	0.218118	0.161657	0.138115	0.175368	0.154216	0.106668
2500	0.241845	0.183643	0.160308	0.201468	0.183583	0.145434
3000	0.264002	0.222403	0.198537	0.219785	0.208797	0.160834
5000	0.314652	0.285477	0.272994	0.260625	0.272635	0.215972
10000	0.444146	0.487149	0.506658	0.407526	0.425191	0.363751
20000	0.552035	0.58209	0.607224	0.510752	0.608566	0.536618
Never	0.671532	0.920983	0.902416			

Tab.2 Intermediate complexity landscape (3 core). Average fitness over 50,000 time steps starting from t=150,000.

It may be worth noting that the distribution of fitness values across the whole landscape is highly biased, with few, high values and the vast majority of low values. The performance distribution is a power law produced by the function x^{30}, with x being the sum of 14 independent uniform random variables (the individual fitness contributions) taking values in the range [0,1]). Such distribution has an expected value of about 0.032 (=1/31), that is, picking randomly a performance value from the whole landscape would generate values in the range [0,1] with most of them being very close to 0 (on average 0.032). The fact that all our simulations report far higher performance levels, which are shown by only a tiny portion of the points in the landscape, means that our learning mechanism is effective in locating, even in the worse cases, high fitness local peaks.

Complex environment (6 cores)

In this case of higher complexity, we find, even more so, that smaller memory provides a fitness advantage. Recall that memory plays two roles, storing so to speak "established and old" learning, as well as novel candidate rules to be assessed. In that, above some level of environmental complexity, the resetting of memory after a shock provides an advantage to keeping memory across radical modifications of the environment. But, remarkably, this seems to occur if shocks are rare enough to allow a relatively thorough learning. Otherwise, cross-cutting robust and relatively blind routines seem to better perform, as shown by Table 3 below.

	Small	Medium	Large	SmallReset	MediumRes.	LargeRes.
300	0.081146	0.078099	0.0792957	0.077693	0.0775796	0.0758342
600	0.0876982	0.0785803	0.0829455	0.0804502	0.0790845	0.0748973
1000	0.0932931	0.079472	0.0817679	0.0840607	0.0830139	0.0788371
1500	0.0991952	0.0815934	0.0815307	0.0887702	0.0887166	0.0810441
2000	0.109788	0.0868158	0.0869898	0.0972125	0.103213	0.0948689
2500	0.116153	0.0855875	0.0873492	0.096938	0.104273	0.0945893
3000	0.122875	0.0875066	0.091828	0.0978429	0.11075	0.100329
5000	0.151306	0.0992009	0.10047	0.110426	0.140738	0.122169
10000	0.20531	0.13491	0.135438	0.126998	0.178356	0.134209
20000	0.272416	0.219173	0.20577	0.147881	0.234087	0.175662
Never	0.514746	0.702701	0.693051			

Tab.3 High complexity landscape (6 core). Average fitness over 50,000 time steps starting from t=150,000.

4.3 Path-dependency

We showed so far that our artificial organizations adaptively learn, but do they learn the same things? That is, do they converge to the same cognition/action patterns facing the same stable environment and the same "objective structure of incentives" stemming from the revealed payoffs? Or is learning path-dependent in the sense that organizations facing the same environment but starting from different initial conditions (rules with the same condition part but different randomly generated actions) and undergoing different adaptation will produce different rules?

In this section we address this question for stationary environments of intermediate complexity.

The behavior of our organizations depends on the interaction of the whole set of rules and of their relative strength. Hence, measuring directly the similitude among organizations is hard. Potentially, two identical set of rules with only a slight difference in strength may produce highly different results, while, on the contrary, very different sets of rule may produce very similar results. Hence, we measure the differentiation of organizations using the indirect measure of the variance across the population of the fitness received through time. Figure 13 shows the variance of the relative fitness across the population of 10 organizations in a complex (6 core) but stationary environment. In more simple settings (1 and 3 cores), the time series of variances show an initial increase but a subsequent fall to zero as all organizations reach the optimal point: the transients converge. In the foregoing case, instead, variance grows and then stabilizes, indicating a persistent dispersion of observed fitness, suggesting that organizations differentiate along path-dependent trajectories leading them to different areas of the landscape. An even stronger pattern emerges for the maximally complex setting (9 cores).

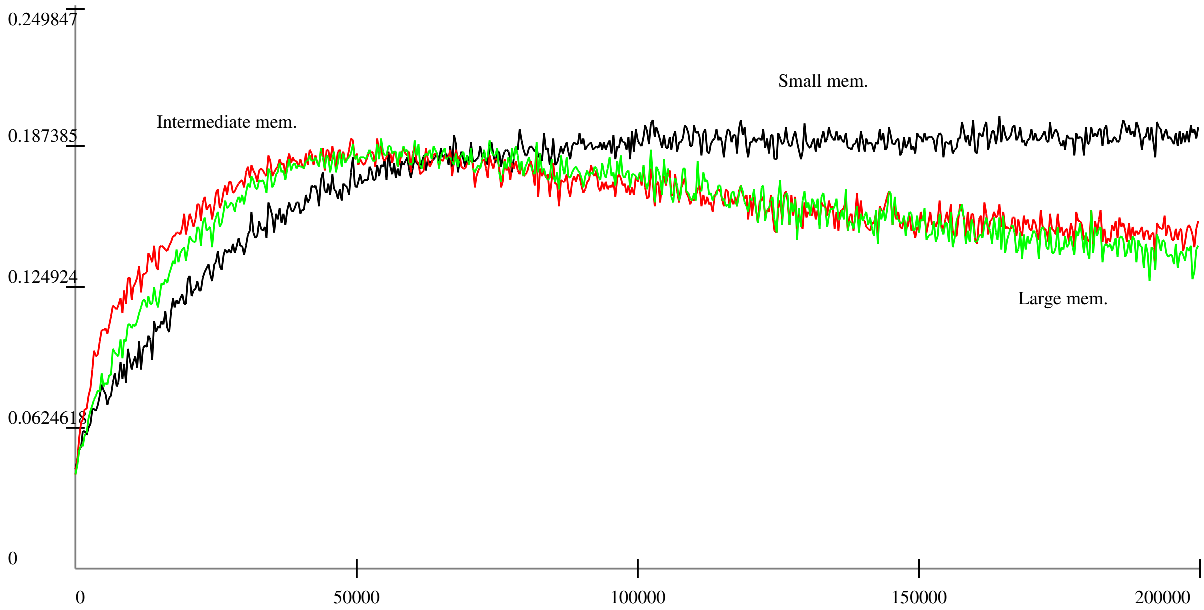

Fig. 13. Variance of relative fitness across a population of 100 organizations. Complex environment (6 cores), stationary environment.

This general path-dependency property is corroborated under the non-stationary setting showing the relative[8] variance for the case of 1 core. When the rates of change of the environment are of an order of magnitude similar to the rates of learning, the dispersion is very high. When the change is more sedate, on the contrary, organizations tend to converge to the global maximum, consequently reducing their differentiation.

[8] To normalize the variance we divided the absolute variance for the square of average values, removing the bias due to the unit of measurement for the indicator of dispersion.

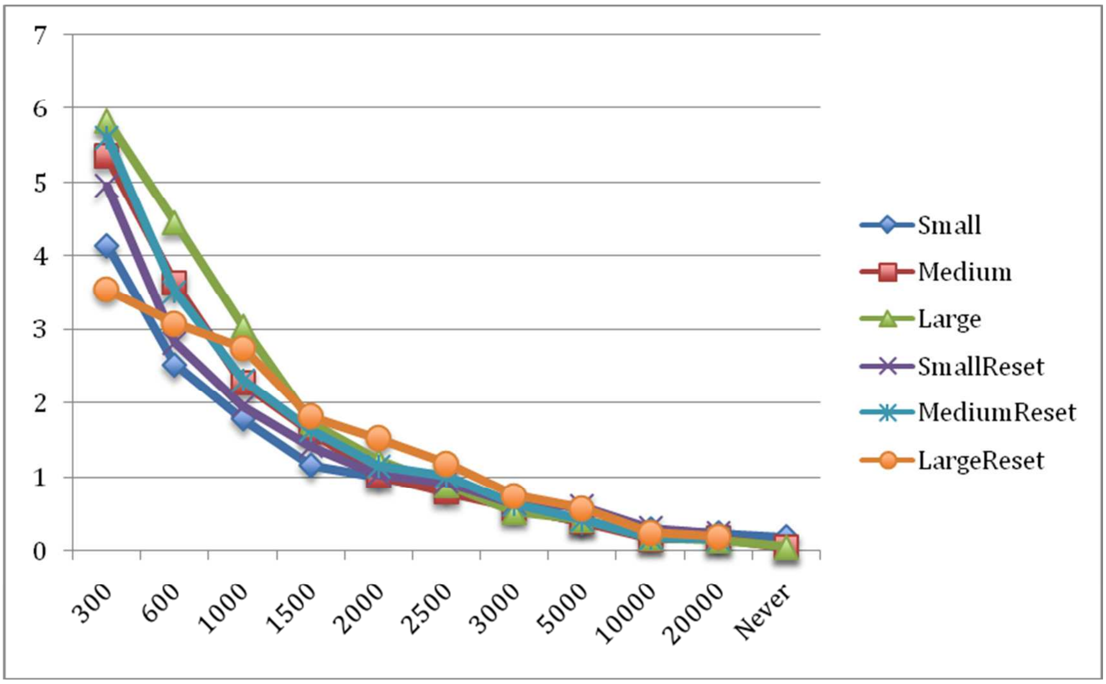

Fig. 14. Relative variance in non stationary environments for different shock frequencies.

5. Conclusions

In this work we proposed a formalization of the notion of organizational memory which straightforwardly accounts also for its being stored in organizational cognitive frames and operational routines. The model allows the analysis of the features of such cognitive and action patterns under different degrees of complexity of the environment and of the problem-solving tasks and under diverse environmental dynamics.

Complexity and non stationarity play a crucial role in determining the evolution of our system of rules.

First, except in the most simple and stationary environments, what organizations learn and "remember" thereafter are not fine-tuned detections of the precise states of the environment and equally fine-tuned behavioral responses, but rather cognitive states ("categories") which capture ensembles of environmental states and correspondingly patterned behavioral responses, indeed, *routines*. If the set of relevant environmental states is relatively large and interdependent, the organization limited learning system is not able to discover all the relevant specific rules and hold them into memory. Consequently, organizations produce general rules each of which applies the same action to sets of states that would each require a different action for optimization, while the organization settles for a single action producing a good ("satisficing"), though suboptimal, result. Highly specific rules will emerge and survive only if they apply to often experienced environmental situations or if they generate very high payoffs. In all other cases, specific optimal rules, even when developed, will be forgotten because they are applied too rarely, and provide too small an advantage, in respect of more general and (slightly) suboptimal competitors.

Second, path-dependence is ubiquitous. With the only exception of the simplest environments, different organizations exposed to the same environmental signals and living on the same fitness landscapes develop distinct interpretative frames and distinct action patterns stemming from idiosyncratic learning processes. In our experiments we study a population of different organizations each facing exactly the same sequence of environmental conditions, and each starting with a single rule with the most general condition part and a randomly generated action part. In very simple environments, with one or very few core bits, all organizations converge to the same repertoire of "optimal" rules which correctly define the optimal action for each state of the core bit(s). Conversely, as the complexity increases and the set of core bits becomes larger or the environment becomes non stationary, organization develop different sets of rules. This is due to three different factors. First, the complexity of the overall landscape (the one combining together environment and action dimensions) generates multiple peaks and different organizations will move in the basin of attraction of different local optima depending upon random mutations of their rules. Second, there are many rules which generate the same or very similar behaviors and therefore have the same or very similar fitness. In other words, there is a good amount of neutrality in the selection landscape for rules and there is a lot of neutrality in memory systems (see Jain and Kogut, 2014 and Marengo, 2014 for two recent papers that develop this argument). Third, and related, there is also a good amount of redundancy in a memory system: if the memory size constraint is not too binding a good number of rules are kept in memory (see below for details). Neutrality and redundancy are fundamental for the evolvability of the system i.e. its capacity to produce variation and novelty and therefore to adapt to environmental changes. If, at each moment in time, memory contained only the specific rules optimally adapted to all and only the environmental situations experiences so far, then adaptation to new environmental conditions would be more difficult as the organization would find itself in a competency trap (Levinthal and March, 1993).

Third, we explored the impact of different memory sizes. Broadly speaking, "more memory is better" for organizational performances, as long as the fitness landscape does not change or changes gradually. This does not hold under frequent environmental shocks: in these circumstances, memory is associated with the competence traps highlighted by Levitt and March (1988) and Gavetti and Levinthal (2000). A more effective evolutionary strategy is to *unlearn,* that is to erase the memory of cognitive frames and routines which were successful in the past but tend to hinder adaptation under the new landscape.

Fourth, somewhat counter-intuitively, above a certain (quite high) threshold of environmental and problem-solving complexity and under repeated and massive environmental shocks, an effective evolutionary strategy returns to the remembrance of what we call *robust interpretative categories* and *robust routines* which yield satisficing outcomes across an array of (imperfectly understood) and changing environments.

Throughout this work we have focused on business organizations. However, the foregoing analytical framework is equally apt to represent other institutional forms, including public institutions. In such case a good deal of the "if's" are of course ideologies and "political visions" while a good part of the "then's" are executive routines. Indeed, it is plausible that the features of path-dependence and inefficient lock-in identified above are likely to be strengthened in such political contexts, generally characterized by high environmental complexity, opaque causal links and blurred landscapes over which policies and administrative behaviours are tested.[9]

Needless to say, there are plenty of exploratory directions ahead. A straightforward one entails the "opening up" of the organization in sub-units which learn partly independently from each other

[9] As known, "administrative behaviour" is also the title of the seminal work by Herbert Simon (cf. Simon, 1997) in whose tradition the resent work is firmly rooted (see also March and Simon 1993, March et al. 2000).

and store different pieces of the overall "memory of the organization". This is generally the case of business organizations, and even more so of political institutions and public policies whereby the "if's" and "then's" are generally controlled and activated by different groups of agents, so that policy makers need to affect in coordinated manners both the if's and the then's.

A second but related venue of research regards the explicit introduction of authority and power. We begin doing that in Dosi and Marengo (2015), but there is a long way to go in order to fully model the power dimension in all types of economic and political organizations.

Consider the foregoing model as a first exploratory but, in our view, promising template of a family of models trying to capture the co-evolution of "organizational cognition" and organizational routines in changing environments.

References:

Anderson, J. R. (1983). *The Architecture of Cognition*. Cambridge, MA: Harvard University Press.

Argyris, C., D. Schon 1978. *Organizational Learning: A Theory of Action Perspective*, Addison-Wesley Publishing Co., Reading, MA.

Argote, L., P Ingram 2000. Knowledge Transfer: A Basis for Competitive Advantage in Firms, Organizational Behavior and Human Decision Processes, 82(1) 150-169

Baumann, O. and N. Siggelkow 2014, Dealing with Complexity: Integrated vs. Chunky Search Processes, forthcoming in ***Organization Science***

Becker, M. 2005. A Framework for applying organization routines in empirical research: Linking antecedents, characteristics and performance outcomes of recurrent interaction patterns, *Industrial and Corporate Change*, 14 (5) 817-846

Becker, M., T. Knudsen 2005. The role of routines in reducing uncertainty. *Journal of Business Research*, 58 (6) 746-757

Becker, M., Lazaric, N., Nelson, R. R., S. Winter 2005. Applying organizational routines in understanding organizational change, *Industrial and Corporate Change*, 14(5) 775-791

Bresnahan, T., S. Greenstein, R. Henderson. 2011. Schumpeterian competition and diseconomies of scope: illustrations from the histories of Microsoft and IBM, in J. Lerner and S. Stern (eds.), *The rate and direction of inventive activity revisited,* NBER, Chicago

Carroll, G.R., M.T Hannan. 2004. *The demography of corporations and industries*, Princeton University Press, Princeton

Castaldi, C., G. Dosi. 2006. The grip of history and the scope for novelty: some results and open questions on path dependency in economic processes", A. Wimmer, R. Koessler (eds), *Understanding Change Models, Methodologies and Metaphors*, Palgrave, London

Cohen, M., P. Bacdayan. 1994. Organizational Routines Are Stored As Procedural Memory: Evidence from a Laboratory Study, *Organization Science*, 5 (4), 554-568

Cohen, M., R., Burkhart, G. Dosi, M. Egidi, L. Marengo, M. Warglien, S. Winter. 1996. Routines and Other Recurring Action Patterns of Organizations: Contemporary Research Issues, *Industrial and Corporate Change*, 5 (3), 653-698

Daft, R., K.E. Weick. 1984. Towards a model of organizations as interpretation systems, *Academy of Management Review*, 9, 284-295.

Denrell, J., C. Fang, D.A. Levinthal. 2004. Form T-mazes to labyrinths: learning from model-based feedback, *Management Science*, 50, 1366-1378.

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

Dosi, G., M. Faillo and L. Marengo, (2008), "Organizational Capabilities, Patterns of Knowledge Accumulation and Governance Structures in Business Firms: An Introduction", *Organization Studies*, 29, 1165-1185

Dosi, G., M. Faillo, L. Marengo, 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, 247-286.

Dosi G., L. Marengo. 2015. The dynamics of organizational structures and performances under diverging distributions of knowledge and different power structures, *Journal of Institutional Economics*, forthcoming.

Dosi G, Marengo, L., A. Bassanini, M. Valente. 1999. Norms as emergent properties of adaptive learning: the case of economic routines, *Journal of Evolutionary Economics,* 9, 5-26.

Duncan, R., A. Weiss. 1979. Organizational learning: implications for organizational design, in *Research in Organizational Behavior*, B. M. Staw ed., 1:75-123., JAI Press, Greenwich

Ethiraj, S., D.A. Levinthal, 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

Feldman, M., J. March. 1981. Information in Organizations as Signal and Symbol, *Administrative Science Quarterly*, 26 (2), 171-186

Fischhoff, B. 1975. Hindsight or foresight: The effect of outcome knowledge on judgment under uncertainty, *Journal of Experimental Psychology*, 1, 288-99

Gabaix, X. 2014. A Sparsity-Based Model of Bounded Rationality, *Quarterly Journal of Economics*, forthcoming

Garund, R., and Rappa, M., (1994), "A Socio-Cognitive Model of Technology Evolution: The Case of Cochlear Implants", *Organization Science*, 5 (3), 344-362

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

Geanakplos, J. 1989. Game Theory Without Partitions, and Applications to Speculation and Consensus, *Cowles Foundation Discussion Paper no. 914*, Cowles Foundation, New Haven, CT

Grant, S. and J. Quiggin 2013. Inductive reasoning about unawareness, *Economic Theory*, 54(3), 717-755

Hedberg, B. 1981. How organizations learn and unlearn?, P. C. Nystrom, W. H. Starbuck (Eds.), *Handbook of organizational design.* 8-27, Oxford University Press, London

Heiner, R. 1983. On the origins of predictable behavior, *American Economic Review*, 73, 560-595

Helfat, C., S. Filkenstein, W. Mitchell, M. Peteraf, H. Sing, D. Teece and S Winter (2006). "*Dynamic Capabilities: Understanding strategic change in organizations*", Oxford: Blackwell

Holland, J.H.1975. *Adaptation in Natural and Artificial Systems*, Ann Arbor, University of Michigan Press

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

Huber, G. 1991. Organizational Learning: The Contributing Processes and the Literatures, *Organization Science*, 2(1), 88-115.

Jain, A., B. Kogut. 2014. Memory and Organizational Evolvabilityina Neutral Landscape. *Organization Science*, 25, 479-493.

Kaplan, S., M. Tripsas. 2008. Thinking about technology: Applying a cognitive lens to technical change, *Research Policy*, 37, 790-805.

Klein, J. I. 1989. Parenthetic Learning in Organizations: Toward the Unlearning of the Unlearning Model, *Journal of Management Studies*, 26, 291-308.

Kauffman., S. A. 1993.*The origins of order*. Oxford University Press, Oxford

Kunisch S., C.Wolf, J. Quodt. 2010. When forgetting is the key: the value of unlearning activities during post-acquisition integration, *Performance*, 3, 5-12, Ernst & Young

Levinthal, D. 1997. Adaptation on Rugged Landscapes, *Management Science, 43*, 934-950

Levinthal, D. 2000.Organizational Capabilities in Complex Worlds. G. Dosi, R. Nelson, S. Winter (eds) *The Nature and Dynamics of Organizational Capabilities*, Oxford University Press.

Levinthal, D., J.G. March. 1993. The myopia of learning, *Strategic Management Journal*, 14, 95-112.

Levitt, B., J.G. March. 1988. Organizational Learning, *Annual Review of Sociology*, 14, 319-340

March, J.G., 1981. Footnotes to Organizational Change, *Administrative Science Quarterly*, 26, 563-577

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

March J., J.P. Olsen. 1976. *Ambiguity and Choice in Organizations*, (2nd Ed.) Universitetsforlaget Bergen

March J., M. Schulz, X. Zhou, 2000. *The Dynamics of Rules*, Stanford University Press, Stanford CA

March J., H. Simon 1993, *Organizations*, 2nd ed., Blackwell Publ., Oxford

Marengo, L. 1992. Coordination and organizational learning in the firm, *Journal of Evolutionary Economics*, 2, 213-226

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

Marengo, L., 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. 2014.Representation, Search and theEvolution of Routines in ProblemSolving, *Industrial and Corporate Change*, forthcoming.

Miller, K.D., B.T. Pentland and S. Choi (2012), Dynamics of Performing and RememberingOrganizational Routines, *Journal of Management Studies*, vol. 49(8), pp.1536-1558.

Mullainathan, S., J. Schwartzstein, and A. Shleifer. 2008. Coarse Thinking and Persuasion, *Quarterly Journal of Economics* 123(2), 577-619

Nelson R., S. Winter, 1982. *An evolutionary theory of economic change*, Belknap Press/Harvard University Press, Cambridge MA

Nystrom, P.C., W.H. Starbuck, 1984. To Avoid Organizational Crises, Unlearn, *Organizational dynamics*, 12, 53-65

Pettigrew, A. M. 1985.*The Awakening Giant: Continuity and Change in Imperial Chemical Industries*, Blackwell Oxford

Rivkin, J.R., N., Siggelkow. 2003. Balancing search and stability: interdependencies among elements of organizational design, *Management Science*, 49, 290-311

Rumelt, R. 1995. Inertia and Transformation, C.A. Montgomery (ed.), *Resource-based and Evolutionary Theories of the Firm*, Kluwer, Boston

Siggelkow, N., D.A., Levinthal. 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., J.W., Rivkin. 2006. When exploration backfires: unintended consequences of multilevel organizational search, *Academy of Management Journal*, 49(4) 779-795

Simon, H.A. 1955. A behavioral model of rational choice, *Quarterly Journal of Economics*, 69, 99-118.

Simon, H.A. (1997). *Administrative Behavior: a Study of Decision-Making Processes in Administrative Organizations* (4th ed.). New York: Free Press

Tripsas, M., G. Gavetti. 2000. Capabilities, cognition, and inertia: Evidence from digital imaging, *Strategic Management Journal*, 1 (10-11), 1147-1161,

Valente, M. 2008.*Laboratory for Simulation Development - LSD*, LEM working papers 2008/12, Pisa, http://www.lem.sssup.it/WPLem/files/2008-12.pdf.

Valente, M. 2014."An NK-like model for complexity", *Journal of Evolutionary Economics,* 24(1), 107-34

Walsh, J.P. 1995. Managerial and Organizational Cognition: Notes from a Trip Down Memory Lane, *Organization Science*, 6(3), 280-321

Walsh, J.P., G.R. Ungson. 1991. Organizational Memory, *The Academy of Management Review*, 16(1), 57-91

Winter, S. 1987. Knowledge and competence as strategic assets, in D. Teece (Ed.) *The competitive challeng*e: 159-184. Ballinger, Cambridge, MA