

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

Giovanni Dosi, Luigi Marengo, Evita Paraskevopoulou and Marco Valente*

This work analyzes and models the nature and dynamics of organizational memory, as such an essential ingredient of organizational capabilities that determine strategic choices in different competitive environments. 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, having more memory provides an advantage. 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 behaviour emerge.

Key words: Organizational memory, Routines, Cognitive categories, Condition-action rules

Manuscript received 30 January 2015; final version received 31 October 2016.

Address for correspondence: Luigi Marengo, Dipartimento di Impresa e Management, Università LUISS, Roma; e-mail: lmarengo@luiss.it

* LEM, Istituto di Economia, Scuola Superiore Sant'Anna, Pisa; Dipartimento di Impresa e Management, Università LUISS, Roma; Departamento de Economía de la Empresa, Universidad Carlos III de Madrid; and School of Economics, Università de L'Aquila, respectively. We thank for the many useful comments and suggestions Dan Levinthal; the participants at 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 'Dynamic Capabilities and the Sustainable Competitiveness of Firms and Nations' Conference, St. Petersburg Graduate School of Management, October 2012; the 'New Frontiers in the Economics and Management of Innovation' workshop, KITeS, Bocconi University, Milan, March 2012, in particular Sarah Kaplan; and, last but not least, the editor and two anonymous referees for this special issue. Financial support from the Institute for New Economic Thinking (INET) grant IN 01100022 is acknowledged. 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); Horizon 2020 Research and Innovation, grant no. 649186; project ECO2013-45864-P financed by the Spanish Ministry of Economics and Competitiveness; and project S2015/HUM-3417 (INNCOMCON-CM) co-financed by the Community of Madrid and the European Social Fund.

1. Introduction

This work analyzes and models the nature and dynamics of organizational memory, an essential ingredient of organizational capabilities that determine strategic choices in different competitive environments. 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 behavioural repertoires contained in the memory itself. Second, the organizational character of the memory also implies that it is resilient to environmental shocks as well as to the replacement of individual members of the organization.

Organizations "remember" because they entail explicit norms and, together, more tacit practices addressed to collectively solve practical and cognitive problems, ranging from, say, the production of a car, all the way to, e.g. the identification of a malaria-treating 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 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 and its bearing in shaping strategies.

Here we set ourselves the ambitious goal to root such research endeavour in three complementary and partly overlapping pillars. The first one is as 'foundational' as it can be, namely, providing at least some basic hints, especially from other disciplines such as cognitive and social psychology and 'applied' epistemology, about what agents (both individuals and organizations) do in terms of cognition and behaviours in complex and changing worlds, that is in the worlds we typically live in. The second pillar regards what economic theories do, or do not do, in order to capture and represent the features mentioned above. Finally, the third pillar regards what organizations do in terms of their (somewhat metaphorical) cognition, their actions and their learning processes about both. The latter is also the ultimate topic of this paper.

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 environments. The analysis of the contrasting roles that 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 current state-of-the-art concerning the very foundations of knowledge in terms of both individual and organizational cognition and operational repertoires in changing environments. Section 3 presents the structure of the simulation model we propose to address the interpretative questions stemming from the foregoing pieces of evidence and to explore the dynamics of collective cognition, behaviour 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 and implications of this work.

2. Cognition, routines and memory in changing environments: a bird’s-eye assessment

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, build constructs of knowledge that can be stored and reused (Argote and Ingram, 2000; Kaplan and Tripsas, 2008).

2.1 *Learning in opaque and changing worlds*

Organizational cognition, behavioural repertoires, and, over time, memory are of course an issue only insofar as one departs from a representation of organizations and of the underlying agents as maximizing entities with ‘Olympic rationality’. The foundational critiques of this model include the contributions of Simon (1955), Cyert and March (1963), March and Simon (1958), Nelson and Winter (1982) and Winter (1964). They are devastating: to any reasonable scientist they would sound like the end of any debate and thus resurrecting the controversies would be like beating a dead horse. But the dead horse is instead alive and kicking, with economists openly neglecting in their models all the evidence and business scholars neglecting such foundational debates. Thus it might be useful to recall telegraphically the basics of the story.

In our view, the fundamental point of departure is the acknowledgement that the world is not transparent. That is, our understanding of the ‘world out there’, in its inner causal links, is at best imprecise and, most often, utterly wrong. Conversely, mainstream economics, alone among social sciences, holds that our understanding of the world is perfect, except possibly for some noise in the perception of the signals coming from it, essentially measurement errors.

A crucial aspect of learning is with regard to *cognition*—that is, the process by which decision-makers form and modify representations in order to make some sense of a reality that is generally too complex and uncertain to be fully understood. Hence the necessity to acknowledge the existence (and persistence) of a systematic gap between the agent’s cognitive abilities and ‘reality’ (were there an omniscient observer able to grasp it fully). Such a gap can take at least two, often interrelated, forms: first, a *knowledge gap*, involving incomplete, fuzzy or simply wrong representations of the environment; and second, a *problem-solving gap* between the complexity of the tasks agents face and their capabilities with respect to accomplishing them.

Regarding both, evolutionary theories of learning might significantly benefit from that branch of cognitive studies concerned with the nature and changes of *categories and mental models* (for different perspectives, see Johnson-Laird, 1983, 2006; Lakoff, 1987; Holland *et al.*, 1986; Margolis, 1987; Byrne, 2005. For the presentation of a few alternative theories, see Mayer, 1992). It is crucial to notice that if one accepts any ‘mental model’ view, learning cannot be reduced to information acquisition (possibly including Bayesian processing of it) but rather is centred on the construction of new cognitive categories and ‘models of the world’.

Cognitive categories, it has been repeatedly shown, go together with various mechanisms of *framing*, by which information is interpreted and also rendered operationally meaningful to the decision-makers (see Kahneman *et al.*, 1982; Borcherding *et al.*, 1990; and March, 1994).

Frames have long been recognized in the sociological and anthropological literature (whatever name is used to refer to them) as being grounded in the collective experience of the actors and in the history of the institutions in which the agency is nested.¹

Indeed, embeddedness seems to go a strikingly long way and affect even the understanding and use of cognitively basic categories, such as that of causality and the very processes by which humans undertake basic operations such as inferences, generalizations, deductions, etc. (Lakoff, 1987; Luria, 1976).

2.2 *The ways the economic discipline did (or did not) take this up...*

Given all the foregoing evidence, how did economic theory accommodate it, or not? A first defence of ‘rationality’ as a descriptive tool (as opposed to a normative, prescriptive one) has been to consider it as a sort of yardstick against which to assess actual behaviours and their deviations vis-à-vis Olympic rationality. In turn, such deviations can be considered as biases.² Or, conversely, deviations can be shown not to be biases at all but straight utility maximizing behaviours once informational imperfections have been taken into account. There is a long tradition in this vein of rationalization of behavioural rules, from the early Baumol and Quandt (1964) to the more recent Love (2013).

Indeed, there are two major pitfalls with either version of this defence of rationality. One is that even in moderately complex environments, let alone evolving ones, it is impossible even for the theorist to determine what the Olympically rational behaviour ought to be (more in Dosi and Egidi, 1991, and Dosi *et al.*, 2005). The other is the epistemological sloppiness of the ex-post rationalization invoking an ensemble of unobservable variables. The reasoning goes more or less like this: since behaviours must be optimal, let us find the constraints on information, computational costs, etc., by which seemingly non-optimal behaviours turn out to be optimal. It is clear that this line of reasoning renders the rationality hypothesis shielded from any empirical refutation.³

¹ Within an enormous volume of literature, a good deal of the sociological tradition has been influenced by the works of Talcott Parson or of the classic Pierre Bourdieu (1977); in anthropology, among others, see the discussions of ‘embeddedness’ by Karl Polanyi (1944, 1957) and Clifford Geertz (1963); see also Robert Edgerton (1985).

² Incidentally, this is the perspective under which some version of behaviour has become respectable for part of the mainstream.

³ A similar approach has been taken recently to accommodate in the self-regarding utility maximizing paradigm the vast evidence on other-regarding preferences, social norms and the like. By adding arguments to the utility function the modeller is tautologically able to state that any observed behaviour can be explained by the maximization of some ad hoc utility function (cf., for instance, Rabin, 2013).

A second defensive trench of orthodoxy, despite the evidence, involves a continuing commitment to ‘rational’ micro-foundations of economic interactions, together with a radical redefinition of the status of rationality assumptions themselves.

‘Rationality’ (however defined), rather than being an approximation to the empirical behaviour of purposeful and cognitively quite sophisticated agents, is assumed to be—so to speak—an ‘objective’ property of behaviours in equilibrium. Add the presumption that (most) observed behaviours are indeed *equilibrium* ones. And, finally, postulate some dynamics of individual adaptation or intra-population selection leading there. What one gets is some version of the famous ‘*as if*’ hypothesis, suggested by Milton Friedman (1953) and rejuvenated in different fashions by more recent efforts to formalize learning/adaptation processes, the outcome of which is precisely the ‘rationality’ assumed from the start (archetypical examples of this faith can be found in Sargent, 1993, and Marimon, 1997).

A thorough and critical discussion of the ‘as ... if’ epistemology has been put forward by Sidney Winter in various essays (see Winter, 1964, 1971). For our purposes here, let us just note the following:

- (i) Any ‘*as if*’ hypothesis on rationality is bound to involve quite a few restrictions similar to those briefly overviewed earlier with reference to more ‘constructive’ notions of rational behaviours, simply transposed into a more ‘ecological’ dimension. That is, canonical rationality, *stricto sensu*, postulates that one decides and acts by purposefully using the appropriate procedures, or by learning them in purposeful, procedurally coherent ways. ‘*As if*’ hypotheses of any kind apparently relax the demands on what agents must consciously know, but at the same time must assume some background mechanism generating the available alternatives—which must include the ‘correct’ ones.
- (ii) While ‘realistic’ interpretations of rationality put most of the burden of explanation upon the power of inbuilt cognition, ‘*as if*’ shifts this burden to the selection dynamics. But, then, supporters of the view ought to show, at the very least, that robust convergence properties are demonstrated by some *empirically justifiable* selection dynamics. In our view, as it stands, nothing like that is in sight. On the contrary, except for very special set-ups, negative results are abundant in, for example, evolutionary games or other forms of decentralized interactions. No matter whether applied to biology or economics, path dependency cannot easily be disposed of; cyclical limit behaviours might occur (see Posch, 1994; and Kaniovski *et al.*, 1996), etc.

2.3 From individuals to organizations

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 that assist organizations to bear uncertainty, besides, as we shall see, manage 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 necessity, 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).

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.

2.4 Modelling routines, memory and learning

For a long time all the way to the present, organizational models have run far behind the qualitative interpretations briefly discussed above. Some catching-up has occurred, however, especially in the field of modelling learning processes in high-dimensional spaces with relatively limited adaptation mechanisms. A promising candidate to model routines and 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 that partly evolves according to the revealed environmental payoffs. Aiming to balance the rather unsynchronized research efforts and respective results between empirical and theoretical research, we build a model that finds its ascendancy there, and in their application in Marengo (1992), albeit with significant modifications.

Below, we present a model which links Classifier Systems and NK fitness landscape models (Kauffman, 1993). The former provides 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 a map defining explicitly the cognition/action/environmental feedbacks and modelling their (evolving) coupling. This is, we believe, a first major advancement with respect to the existing literature. Our organization explores a complex and possibly changing landscape in which some dimensions are outside its control (the environmental traits) and some are within (the action traits). 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, in practice, 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.

In fact, the characteristics and evolution of organizational memory mirror the characteristics and evolution of organizational routines. In the case of routines, the memory elicits a '*relatively complex pattern of behaviour triggered by a relatively small number of initiating signals or choices*' (Cohen *et al.*, 1996).

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 the environment, the less contingent the behaviour (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.⁴ The organization is able to recall specific cognitive frames and behavioural 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 behaviours. 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. Indeed, organizations typically compete on such complex landscapes, where interrelated technological and behavioural traits are responsible for path-dependent reproduction of organizational arrangements (Marengo, 1996; Levinthal, 1997, 2000; Rivkin and Siggelkow, 2003; Siggelkow and Levinthal, 2005).

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 behaviours.

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 behaviour (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

⁴ The idea of path-dependence in economics is generally associated with David (1985) and Arthur (1989), but of course is a general property of non-linear, non-ergodic systems. We discuss different forms and degrees of influence of the past upon the future in Castaldi and Dosi (2006).

Levinthal (2000), Rivkin and Siggelkow (2003) and Siggelkow and Levinthal (2005), among others.

The second class includes models focusing upon the relationship between the division of cognitive labour 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 embodied in the organization. Indeed, managing an organization, designing and producing cars or software packages, discovering a new drug, etc., can be 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, 1996), Denrell *et al.* (2004), Valente (2014) and Baumann and Siggelkow (2013).

On the output side, components making up an artefact 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 the form of a 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.

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, 1996) presents models focused 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 with traits similar to those in existing successful rules, though being completely novel, giving hope to further improvement of the performance.

A germane family of models, of somewhat more reduced form but also more elegant and related to a lower-dimensional space, involves precisely some ‘black-boxing’, in particular concerning the relationship between organizational traits (including, of course, behavioural rules) and their actual expressions. Such a modelling *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 non-linear 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 modelled 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 behaviour; the latter entails experiential choice based on on-line evaluation of a limited set of alternatives which are close to current behaviours. In Gavetti and Levinthal'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 behaviours.

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 that adopt a joint cognitive and experiential search dominate the population. This becomes particularly evident under rugged landscapes, in which organizations using 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 over time.

3.2 *The model*

3.2.1 An informal description. Operationalizing these conceptual paths, 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, 1994; 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 its degree of generality vs. specificity, according to the size of the set of environmental conditions to which it can be applied. General rules (with low specificity) prescribe the same course of action for a broad range of environmental conditions, while specific rules 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 generic action 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; and 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 a 'strength' variable, whose value is updated on the basis of the performance provided when applied. The strengths of the rules whose conditional part fits current conditions are used to decide which one to apply. The strongest rule, among those which satisfy the current environmental conditions, will be preferred for action, under the assumption that high performance in the past should guarantee good results.

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 rules 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 difference between the expected performance (estimated on the basis of past rewards) and the payoff received by its action. Rules that are active because they satisfy the current environmental conditions, but have not been chosen for action, have their strength reduced by a small tax punishing their failure to be selected. 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 being 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 threshold will reduce the size of memory, as fewer rules will manage to survive the selection on inactivity. Notice that, contrary to most of the literature on CS, the size of memory is endogenous because, as we will show below, it depends on all the features of the environment, which influence how many rules are compatible with the environment at each time, and, consequently, the frequency of activation, eventually determining the size of organizational memory expressed by the number of rules.

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 variation driven (novelty is constantly introduced as variations on existing rules).

We test the behaviour 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 an 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 to conditions of its set of rules. The remaining $N_a = N - N_e$ are instead dimensions (policies) pertaining to the action of the organization and chosen by the latter according to its repertoire of rules. The payoff for the organization will be determined, in principle, by the current configuration of all the N dimensions, though 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 behaviour of our organization in landscapes characterized by specific structures of interdependencies, using the methodology presented in, e.g. Frenken *et al.* (1999) and Valente (2014). To represent the necessity of developing a set of categories, which the organizations have to discover, 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 relative shape of the action landscape, while the remaining 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 changes, in general also the ranking of action will undergo random changes as pertaining to a different landscape.

In order to identify the best action under different environmental conditions, 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 bits⁵ (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, besides changing entirely for each configuration of the environmental dimensions.

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 and complexity of the overall task faced by the organization.

Concerning the environmental dimensions, we simulate stationary and non-stationary environments. In the former case the environment-action landscapes are generated (with the warranted complexity structure) at the beginning of the simulation and never change. The state of the environment (the configuration of the N_e environmental bits fed to the organization) changes at each moment in time, but the mapping between each environmental state, actions and payoff remains constant throughout the whole simulation

⁵ Following the cited literature, we use a different indicator of complexity from that used in the original NK models. In our text we define K as expressing not the number of epistatic relations of each dimension, but the total number of dimensions affecting the fitness contribution of each dimension. So, instead of varying from 0 (independent dimensions) to $N-1$ (maximum complexity) as in Kauffman, our indicator varies from 1 to N .

exercise. On the contrary, in a non-stationary environment the payoff values are subject to change, and therefore the shape of the landscape is modified, along with the relative convenience of deploying one action instead of another. In this case, only the structure of the interdependency links remains constant so that the relevant categories do not change.

In the next subsections we will present a detailed description of our model, and then, in Section 4, we will examine its behaviour in different environments.

3.2.2 Environment, states of the world and payoffs. The **environment** is fully described by a set of n elementary ‘states’ $E = \{e_1, e_2, \dots, e_n\}$. For simplicity we assume that each environmental feature may take only two values, $e_i \in \{0, 1\}$.

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

Payoffs or fitness: In general, the payoff 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 $\pi : E \times A \rightarrow [0, 1]$. We explore different *complexity structures* concerning the mapping from the $E \times A$ space to the payoff. 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 relative ranking. Hence, suppose that the vector $[a_1, a_2, \dots, a_m]$ is the optimal action when the environment is described by vector $[e_1, e_2, \dots, e_n]$. 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_i 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, since their number effectively determines the number of different landscapes to be explored by the action part. Note that learning is exclusively driven by the payoff, which is the only signal organizations receive on how good their actions are on that specific environmental state. The complexity of the environment is implemented by an NK-like fitness function, where the epistatic links are assigned so as to generate the desired correlation properties among the environmental and action dimensions of the landscape.

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 behaviour. 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, which constitute our Classifiers System, perform the two interrelated tasks of detecting and memorizing environmental regularities (i.e. partitioning environmental states into relevant 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:

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

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, $S_{i,i}$, which is an indicator of the payoff it has cumulated from past applications, minus possibly some ‘tax’ reducing the strength (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 active, i.e. its action part can be potentially adopted by the organization in the current environmental state. In case more than one rule is active at a given time step, the organization selects one by considering both the strength and the specificity of each rule. For details on how different rules compete for action and are rewarded, refer to the Appendix at the end of the paper.

At the start of each simulation run we assume that the organization has 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.

While the organization faces new environmental conditions and collects payoff from the application of one of its actions, new rules are regularly introduced into the memory as variations of the most successful among existing ones. In particular we use two distinct rule-generation mechanisms: specification and generalization. The former takes an existing rule and narrows down its domain of applicability, turning a # in the condition part into either 1 or 0; the latter does the opposite, basing the decisions on the collection of directly observed performance. Details on how specification and generalization work are in the [Appendix](#).

Finally, rules are erased from memory, i.e. forgotten, when they are infrequently used, under the assumption that its failure to win the competition with other rules when the environmental condition is compatible with its condition suggests worse expected performance than alternative possibilities. Some of the technical details on how we implement the forgetting mechanism can be found in the [Appendix](#).

4. Some results⁶

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. Therefore, we do not advance claims concerning universal properties of our model; rather, our aim is to replicate qualitatively a few properties of the systems relevant to the theoretical debate, and to interpret the underlining generative mechanisms as revealed by both statistical analysis on a large number of replications to

⁶ The simulations presented here were implemented in the ‘Laboratory for Simulation Development’ (LSD), a simulation platform developed by one of us ([Valente, 2008](#)). The code of the model and the configurations used are available from the authors upon request.

ensure robustness of results and detailed inspection of the state of the model at crucial times to provide reliable insight.⁷

In our experiments we focus 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 experiments 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), determining a current environmental state, and organizations have to choose one rule from their repertoire to match the observed environmental conditions. 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 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 in response to the current state of the environment. It is well known in fact that the maximum fitness value of an NK landscape depends on the value of K , thus 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 a stationary environment

Let us begin by analyzing the behaviour of our model for varying degrees of environmental complexity in stationary environments, i.e. those characterized by a stable landscape capturing for instance eras of stability of industries with established technological standards and ‘predictable’ demand patterns. 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.⁸

⁷ For a methodological discussion on the use of agent-based models see [Valente \(2016\)](#).

⁸ 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.

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

In this setting the ‘optimal’ repertoire consists of two rules made of all #’s but for the core bit (hence specificity 1, cf. Figure 3), each applying to either state of the core bit, and containing the appropriate action. The figures show that all 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 initially 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 two 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




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

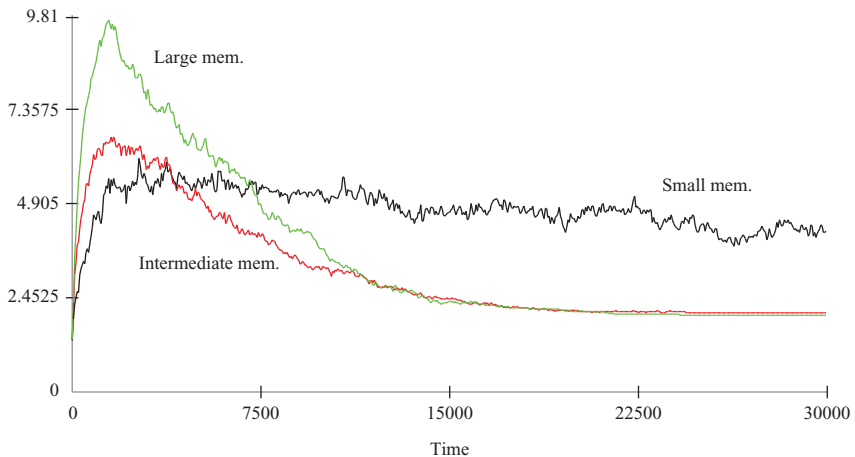


Figure 2: Average number of rules over 100 organizations; simple (1 core) environment.




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

start their learning by rough and imprecise over-categorizations⁹ 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. At this stage, while organizations are still learning both the relevant categories of the environment and the relative actions, they seem to apply a sort of superstitious behaviour, exploring different actions for state conditions that, in reality, would require the same. While refining the action part, they also realize that some categories can be merged, as they obtain the same action part for distinct categories. Doing so, the newly merged rules are more general and we observed a reduction in 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 suboptimal, 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 trade-offs between exploitation and exploration, well in tune with [March \(1991\)](#).

4.1.2 Intermediate complexity (3 cores). Let us now increase the complexity of the landscape to 3 cores. In this settings ([Figures 4, 5 and 6](#)) 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 a 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.

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

⁹ Incidentally, note that these ‘categories’ are not partitions: their intersection is not the empty set.

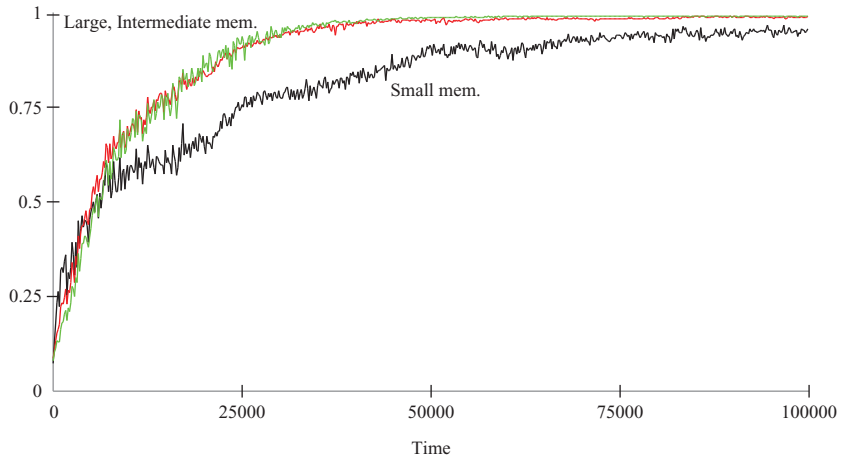


Figure 4: Average fitness over 100 organizations; intermediate (3 core) environment.

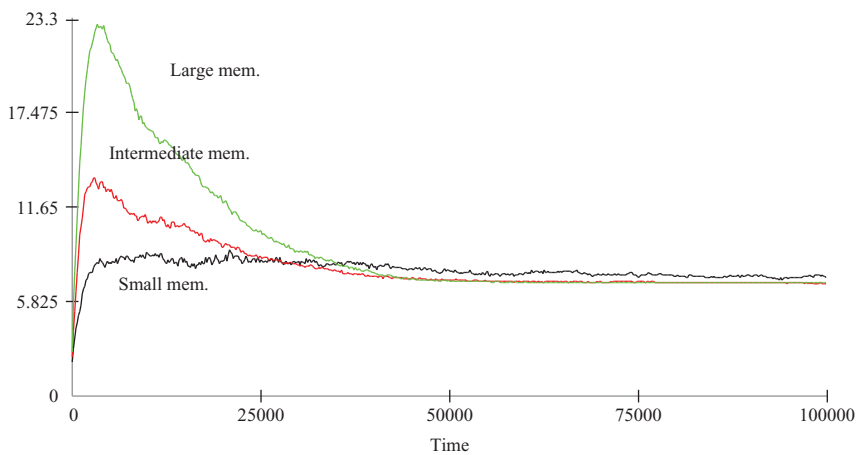


Figure 5: Average number of rules over 100 s; intermediate (3 core) environment.

‘premature learning’ under very small memory are more costly when (nearly) optimal behaviours 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.

4.1.3 Complex environment (6 cores). With landscapes of higher complexity, whose results are reported in [Figures 7, 8 and 9](#), our artificial organizations fail to reach systematically the maximum fitness. The specificity statistics ([Figure 9](#)) shows that organizations fail to map correctly the relevant ‘true’ environmental categories, that should be 6 for the correct categorization, 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

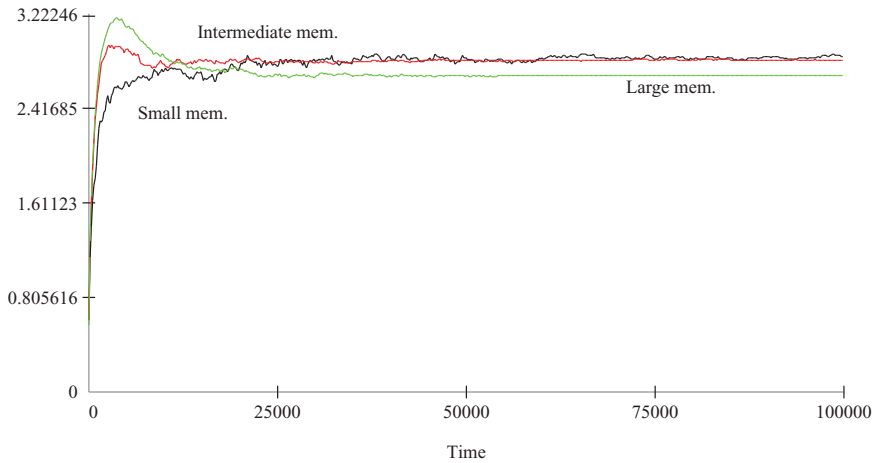


Figure 6: Average specificity of rules over 100 organizations; intermediate (3 core) environment.




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

show a consistently better performance than larger ones, hinting that excessively large memory appears, at least for a while, as a liability rather than an asset.

4.1.4 *Maximally complex environment (9 cores).* Pushing still higher the environment complexity (all the 9 environmental bits are core bits) yields somewhat different dynamics (Figures 10, 11 and 12). 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. The reasons can be understood by comparing the results on the specificity and number of rules. Organizations with large memory sizes collect and maintain a larger number of rules, but fail to extract from this a useful map between environmental states and actions; beyond a certain limit, they fail to push specialization which, evidently, is produced merely by retaining any rule irrespective of their past performance. On the contrary, organizations with average memory size




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

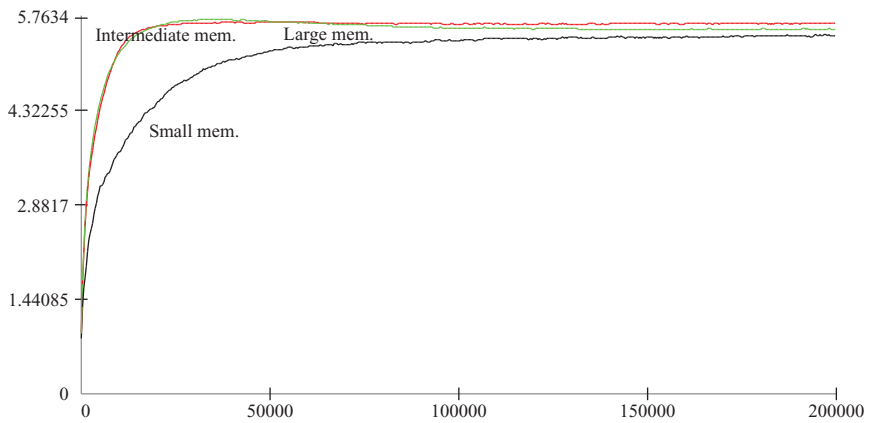


Figure 9: Average specificity of rules over 100 organizations; complex (6 core) environment.




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

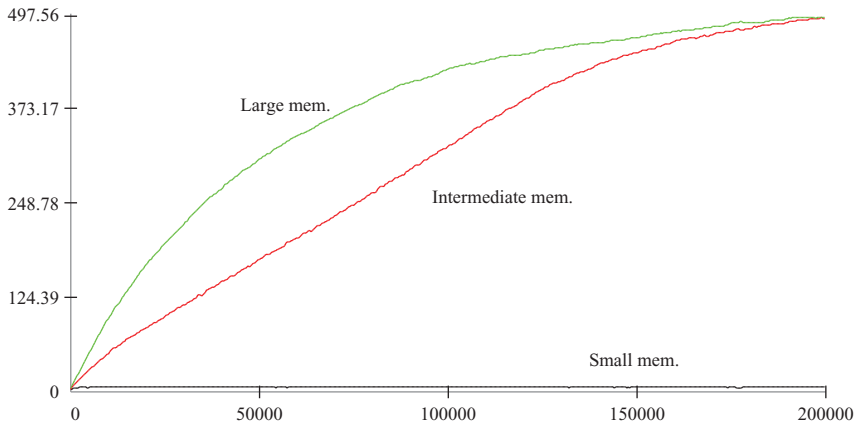


Figure 11: Average number of rules over 100 organizations; highly complex (9 core) environment.

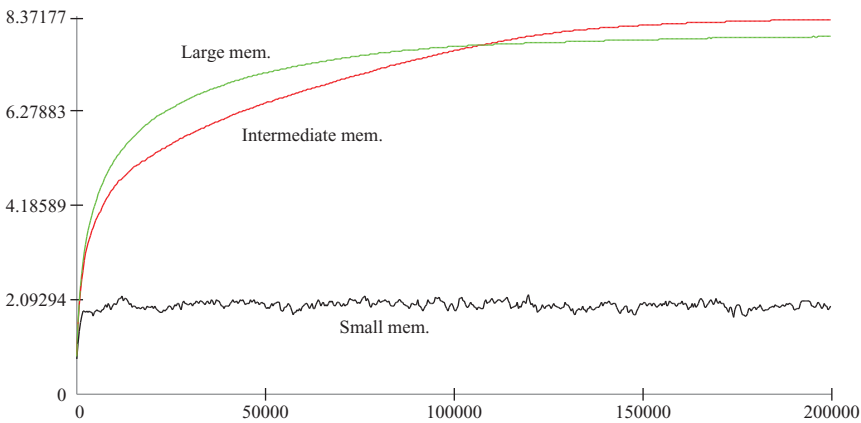


Figure 12: Average specificity of rules over 100 organizations; highly complex (9 core) environment.

need to be more disciplined, and retain rules only if proven good. They leverage on this attitude to build, eventually, a similar number of rules as the former organizations, but the set is different, showing a greater capacity to distinguish the relevant states, as indicated by the larger specificity.

4.2 Non-stationary environments

So far we considered stationary environments, in which a good rule, if discovered, remains good for a long period of time. Let us now consider environments which persistently undergo regular shocks, generating each time a new landscape, i.e. a novel mapping between states of the world, actions and payoffs. Shocks occur at regular intervals (we test for different frequencies), implying 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. Shocks might include, for instance, disruptive technological innovations, new competition and regulations policies, etc.

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 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 developed with the ‘obsolete’ landscape: in fact, they do not ‘know’ of the implications of the shock, but painstakingly need to 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 evaluating the performance of new, untested, organizations. In the tables below we report the results for the average fitness generated by organizations exploring increasingly complex landscapes: 1, 3 and 6 cores.¹⁰ The rows in the table correspond to different frequencies of shocks, reported in the first column. The columns report the size of memory and whether the organizations reset their memory or not. The last row, that we call ‘Never’, is the equivalent of the previous case for stationary environments (with the small technical difference that for the sake of comparability we ‘force’ exploration for new rules to continue forever, thus undermining the measure of the fitness as organizations keep on exploring even after having reached the maximum fitness).

4.2.1 Simple environment (1 core). Results show that in highly volatile environments, with high frequency of shocks, small memory size provides higher fitness than strategies allowed larger memory size. In these cases, reported in [Table 1](#), 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 detrimental for performance. The reason is that a larger 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, favouring reuse 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 values are very similar.

4.2.2 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.

It may be worth noting that the distribution of fitness values across the whole landscape is highly biased, with few high values and a vast majority of low values. The

¹⁰ We skip the analysis of maximally complex landscapes (9 cores) because the results are perfectly in line with the other cases.

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 a 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. 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, but the ecology yields systematically heterogeneous performances.

4.2.3 *Complex environment (6 cores)*. In this case of higher complexity we find, even more so, that smaller memory provides an advantage. Recall that memory plays two roles, storing so to speak ‘established and old’ learning, as well as novel candidate

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

Frequency	Retain memory after shocks			Reset memory after shocks		
	Small	Medium	Large	Small	Medium	Large
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			

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

Frequency	Retain memory after shocks			Reset memory after shocks		
	Small	Medium	Large	Small	Medium	Large
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			

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 perform better, as shown by [Table 3](#) below.

4.3 Path-dependency

We showed so far that our artificial organizations learn adaptively, but we may ask: 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 behaviour 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 rules may potentially 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

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

Frequency	Retain memory after shocks			Reset memory after shocks		
	Small	Medium	Large	Small	Medium	Large
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			

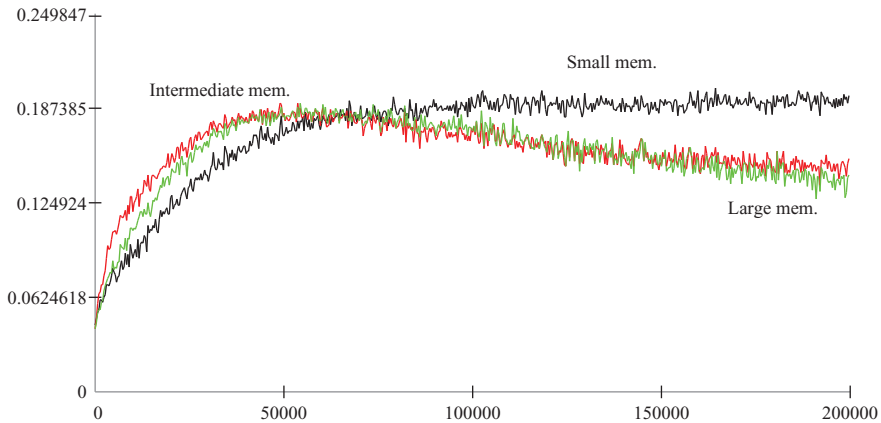


Fig. 13. Variance of relative fitness across a population of 100 organizations. Complex environment (6 cores), stationary environment.




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

them to different areas of the landscape. An even stronger pattern differentiation emerges for the maximally complex setting (9 cores).

This general path-dependency property is corroborated under the non-stationary setting, reported in Figure 14, showing the relative¹¹ 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.

5. Conclusions

This work presents a formalization of the notion of organizational memory embodying and combining organizational ‘cognitive’ frames and operational routines. This allows the analysis of both cognitive and operational patterns of organizational behaviours. We explore them under different degrees of complexity of the environment and of the

¹¹ To normalize the variance, we divided the absolute variance for the square of average values, removing the bias due to the difference on unit of measurement for the indicator of dispersion.

problem-solving tasks and under diverse environmental dynamics. We find that different degrees of complexity and non-stationarity play a crucial role in determining the evolution of the overall rules system.

First, complexity and non-stationarity determine (to a degree) the balance between specialization and generalization. We find that, 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 behavioural responses, but rather cognitive states (‘categories’) which capture ensembles of environmental states and correspondingly patterned behavioural responses, indeed, *routines*. If the set of relevant environmental states is relatively large and interdependent, as in the case of high complexity, the organization’s limited learning system is not able to discover all the relevant specific rules and hold them in 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. In practice, although such practice allows organizations to survive in complex environments it may downplay incentives for excellence related to specialization, as highly specific rules will emerge and survive only if they apply to frequently 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 comparison to more general and suboptimal competitors. While generic rules are seen as a safety policy against uncertainty, what is here at stake is the incentive for specialization, namely for the quest of optimal solutions.

Second, with the only exception of the simplest environments, the path-dependent and idiosyncratic learning processes, rooted in different organizations, result in the development of distinct interpretative frameworks even if all organizations are exposed to the same environmental signals and living on the same fitness landscapes. This is due to three different factors. The first one is that the complexity of the overall landscape (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. The second factor is that there are many rules which generate the same or very similar behaviours 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, 2015](#); for two recent contributions to the topic). The third related factor is that 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 experienced 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, a larger capacity to develop and store rules is beneficial for organizations, as long as the fitness landscape does not change or changes gradually so that adaptation to novelty is

unnecessary or can be done at a slow pace. This does not hold under frequent environmental shocks: in these circumstances flexibility and fast adaptation are prerequisites for survival and memory becomes 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, and somewhat counterintuitively, 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.

Bibliography

- Anderson, J. R. 1983. *The Architecture of Cognition*, Cambridge, MA, Harvard University Press
- Argote, L. and Ingram, P. 2000. Knowledge transfer: a basis for competitive advantage in firms, *Organizational Behaviour and Human Decision Processes*, vol. 82, no. 1, 150–69
- Argyris, C. and Schon, D. 1978. *Organizational Learning: A Theory of Action Perspective*, Reading, MA, Addison-Wesley Publishing Co.
- Arthur, W. B. 1989. Competing technologies, increasing returns, and lock-in by historical events, *Economic Journal*, vol. 99, 116–31
- Baumann, O. and Siggelkow, N. 2013. Dealing with complexity: integrated vs. chunky search processes, *Organization Science*, vol. 24, no. 1, 116–32
- Baumol, W. J. and Quandt, R. E. 1964. Rules of thumb and optimally imperfect decisions, *American Economic Review*, vol. 54, 23–46
- 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*, vol. 14, no. 5, 817–46
- Becker, M. and Knudsen, T. 2005. The role of routines in reducing uncertainty, *Journal of Business Research*, vol. 58, no. 6, 746–57
- Becker, M., Lazaric, N., Nelson, R. R., and Winter, S. 2005. Applying organizational routines in understanding organizational change, *Industrial and Corporate Change*, vol. 14, no. 5, 775–91
- Borcherding, K., Larichev, D. L., and Messick, D. M. (eds.). 1990. *Contemporary Issues in Decision Making*, New York, North-Holland
- Bourdieu, P. 1977. *Outline of a Theory of Practice*, Cambridge, Cambridge University Press
- Byrne, R. M. J. 2005. *The Rational Imagination: How People Create Counterfactual Alternatives to Reality*, Cambridge, MA, MIT Press
- Castaldi, C. and Dosi, G. 2006. The grip of history and the scope for novelty: some results and open questions on path dependency in economic processes, pp. 99–128 in Wimmer, A. and Koessler, R. (eds.), *Understanding Change Models, Methodologies and Metaphors*, London, Palgrave
- Cohen, M. and Bacdayan, P. 1994. Organizational routines are stored as procedural memory: evidence from a laboratory study, *Organization Science*, vol. 5, no. 4, 554–68
- Cohen, M., Burkhart, R., Dosi, G., Egidi, M., Marengo, L., Warglien, M., and Winter, S. 1996. Routines and other recurring action patterns of organizations: contemporary research issues, *Industrial and Corporate Change*, vol. 5, no. 3, 653–98
- Cyert, R. M. and March, J. G. 1963. *A Behavioural Theory of the Firm*, Englewood Cliffs, NJ
- Daft, R. and Weick, K. E. 1984. Towards a model of organizations as interpretation systems, *Academy of Management Review*, vol. 9, 284–95
- David, P. A. 1985. Clio and the economics of QWERTY, *American Economic Review*, vol. 75, no. 2, 332–37
- Denrell, J., Fang, C., and Levinthal, D. A. 2004. From T-mazes to labyrinths: learning from model-based feedback, *Management Science*, vol. 50, 1366–78

- Dosi, G. and Egidi, M. 1991. Substantive and procedural uncertainty: an exploration of economic behaviours in changing environments, *Journal of Evolutionary Economics*, vol. 1, 145–68
- Dosi, G., Faillo, M., Marengo, L., and Moschella, D. 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*, vol. 24, 247–86
- Dosi, G., Marengo, L., Bassanini, A., and Valente, M. 1999. Norms as emergent properties of adaptive learning: the case of economic routines, *Journal of Evolutionary Economics*, vol. 9, 5–26
- Dosi, G., Marengo, L., and Fagiolo, G. 2005. Learning in evolutionary environments, pp. 255–328 in Dopfer, K. (ed.), *The Evolutionary Foundations of Economics*, Cambridge, Cambridge University Press
- Duncan, R. and Weiss, A. 1979. Organizational learning: implications for organizational design, pp. 75–123 in Staw, B. M. (ed.), *Research in Organizational Behaviour*, vol. 1, Greenwich, JAI Press
- Edgerton, R. B. 1985. *Rules, Exceptions and Social Order*, Berkeley and Los Angeles, University of California Press
- Ethiraj, S. and Levinthal, D. A. 2004. Bounded rationality and the search for organizational architecture: an evolutionary perspective on the design of organizations and their evolvability, *Administrative Science Quarterly*, vol. 49, no. 3, 404–37
- Frenken, K., Saviotti, P. P., and Trommetter, M. 1999. Variety and niche creation in aircraft, helicopters, motorcycles and microcomputers, *Research Policy*, vol. 28, no. 5, 469–88
- Friedman, M. 1953. *Essays in Positive Economics*, Chicago, University of Chicago Press
- Gavetti, G. and Levinthal, D. 2000. Looking forward and looking backward: cognitive and experimental search, *Administrative Science Quarterly*, vol. 45, 113–37
- Geertz, C. 1963. *Peddlers and Princes*, Chicago, University of Chicago Press
- Heiner, R. 1983. On the origins of predictable behaviour, *American Economic Review*, vol. 73, no. 4, 560–95
- Holland, J. H. 1975. *Adaptation in Natural and Artificial Systems*, Ann Arbor, University of Michigan Press
- Holland, J. H., Holyoak, K. J., Nisbett, R. E., and Thagard, P. R. 1986. *Induction: Processes of Inference, Learning and Discovery*, Cambridge, MA, MIT Press
- Jain, A. and Kogut, B. 2014. Memory and organizational evolvability in a neutral landscape, *Organization Science*, vol. 25, 479–93
- Johnson-Laird, P. N. 1983. *Mental Models*, Cambridge, MA, Harvard University Press
- Johnson-Laird, P. N. 2006. *How We Reason*, Oxford, Oxford University Press
- Kahneman, D., Slovic, P., and Tversky, A. (eds.). 1982. *Judgment under Uncertainty: Heuristics and Biases*, Cambridge, Cambridge University Press
- Kaniovski, Y. M., Kryazhinskii, A. V., and Young, H. P. 1996. 'On the Robustness of Stochastic Best-Reply Dynamics in Repeated Games', Working Paper WP 96-45, International Institute of Applied Systems Analysis, Laxenburg, Austria
- Kaplan, S. and Tripsas, M. 2008. Thinking about technology: applying a cognitive lens to technical change, *Research Policy*, vol. 37, 790–805
- Kauffman, S. A. 1993. *The Origins of Order*, Oxford, Oxford University Press
- Lakoff, G. 1987. *Women, Fire and Dangerous Things: What Categories Reveal about the Mind*, Chicago, University of Chicago Press
- Levinthal, D. 1997. Adaptation on rugged landscapes, *Management Science*, vol. 43, 934–50
- Levinthal, D. 2000. Organizational capabilities in complex worlds, pp. 363–79 in Dosi, G., Nelson, R., and Winter, S. (eds.), *The Nature and Dynamics of Organizational Capabilities*, Oxford, Oxford University Press
- Levinthal, D. and March, J. G. 1993. The myopia of learning, *Strategic Management Journal*, vol. 14, 95–112
- Levitt, B. and March, J. G. 1988. Organizational learning, *Annual Review of Sociology*, vol. 14, 319–40
- Love, D. A. 2013. Optimal rules of thumb for consumption and portfolio choice, *Economic Journal*, vol. 123, 932–61
- Luria, A. R. 1976. *Cognitive Development: Its Cultural and Social Foundations*, Cambridge, MA, Harvard University Press

- March, J. G. 1991. Exploration and exploitation in organizational learning, *Organization Science*, vol. 2, 71–87
- March, J. G. 1994. *A Primer on Decision Making: How Decisions Happen*, New York, Free Press
- March, J. G. and Simon, H. A. 1958. *Organizations*, New York, Basil Blackwell
- Marengo, L. 1992. Coordination and organizational learning in the firm, *Journal of Evolutionary Economics*, vol. 2, 213–26
- Marengo, L. 1996. Structure, competence and learning in an adaptive model of the firm, pp. 124–54 in Dosi, G. and Malerba, F. (eds.), *Organization and Strategy in the Evolution of the Enterprise*, London, Macmillan
- Marengo, L. 2015. Representation, search and the evolution of routines in problem solving, *Industrial and Corporate Change*, vol. 24, 951–80
- Marengo, L. and Dosi, G. 2005. Division of labor, organizational coordination and market mechanisms in collective problem-solving, *Journal of Economic Behaviour and Organization*, vol. 58, 303–26
- Margolis, H. 1987. *Patterns, Thinking and Cognition*, Chicago, University of Chicago Press
- Marimon, R. 1997. Learning from learning in economics: towards a theory of the learnable in economics, pp. 278–315 in Kreps, D. M. and Wallis, K. F. (eds.), *Advances in Economics and Econometrics: Theory and Applications*, vol. 1, Cambridge, Cambridge University Press
- Mayer, R. E. 1992. *Thinking, Problem Solving and Cognition*, New York, W. H. Freeman
- Miller, K. D., Pentland, B. T., and Choi, S. 2012. Dynamics of performing and remembering organizational routines, *Journal of Management Studies*, vol. 49, no. 8, 1536–58
- Nelson, R. and Winter, S. 1982. *An Evolutionary Theory of Economic Change*, Cambridge, MA, Belknap Press/Harvard University Press
- Polanyi, K. 1944. *The Great Transformation*, Boston, Beacon Press
- Polanyi, K., Arnsberg, C.M. and H.W.Pearson (eds.), 1957. *Trade and Market in the Early Empires*, Glencoe, IL, Free Press
- Posch, M. 1994. ‘Cycling with a Generalized Urn Scheme and a Learning Algorithm for 2X2 Games’, Working Paper WP94-76, International Institute of Applied Systems Analysis, Laxenburg, Austria
- Rabin, M. 2013. Incorporating limited rationality into economics, *Journal of Economic Literature*, vol. 51, no. 2, 528–43
- Rivkin, J. W. and Siggelkow, N. 2002. Organizational sticking points on NK landscapes, *Complexity*, vol. 7, no. 5, 31–43
- Rivkin, J. W. and Siggelkow, N. 2003. Balancing search and stability: interdependencies among elements of organizational design, *Management Science*, vol. 49, 290–311
- Sargent, T. J. 1993. *Bounded Rationality in Economics*, Oxford, Clarendon Press
- Siggelkow, N. and Levinthal, D. A. 2005. Escaping real (non-benign) competency traps: linking the dynamics of organizational structure to the dynamics of search, *Strategic Organization*, vol. 3, no. 1, 85–115
- Siggelkow, N. and Rivkin, J. W. 2006. When exploration backfires: unintended consequences of multilevel organizational search, *Academy of Management Journal*, vol. 49, no. 4, 779–95
- Simon, H. A. 1955. A behavioural model of rational choice, *Quarterly Journal of Economics*, vol. 69, 99–118.
- 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*, vol. 24, no. 1, 107–34
- Valente, M. 2016. Agent-based simulation modeling: definitions and a methodological proposal, in Biggiero, L., Angelini, P.P., Basevi, M., Carbonara, N., Mastrogiorgio, A., Pessa, E., Sevi, E., Valente, M. (eds.), *Relational Methodologies and Epistemology in Economics and Management Sciences*, Hershey, PA, IGI Global
- Walsh, J. P. and Ungson, G. R. 1991. Organizational memory, *Academy of Management Review*, vol. 16, no. 1, 57–91
- Winter, S. G. 1964. *Economic ‘natural selection’ and the theory of the firm*, doctoral dissertation, Yale University
- Winter, S. G. 1971. Satisficing, selection and the innovative remnant, *Quarterly Journal of Economics*, vol. 85, 237–61

Appendix

In this appendix we provide some additional details on the model. We start with our system of rule selection and rule reward.

Rule selection and reward

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:

$$B_{t,i} = [\beta + (1 - \beta)\sigma_i / N_e] S_{t-1,i} \quad (1)$$

where β is a parameter representing the relative importance of specificity in bidding, σ_i is the specificity of the condition (number of bits different from #) and N_e is the number of environmental bits. In summary, a bid is higher, the higher 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 values are computed for all the active rules, those whose conditional part matches the current environmental state. A single rule is chosen with probabilities proportional to the normalized values of the bids, and its action part is played, producing a payoff.

The strengths of all rules are updated by detracting a small ‘tax’ on strength to avoid the permanence of a rarely applied rule:

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

where τ_1 is the generic ‘maintenance’ tax, acting like depreciation in case the rule is never applied.

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

$$S_{t,i} = \gamma [S_{t-1,t} (1 - \tau_1) - \tau_2 B_{t,i}] + (1 - \gamma) Payoff_{t,i} \quad (3)$$

The additional components are τ_2 , a scale parameter tuning the cost of bidding, and γ , controlling the speed by which the strength tracks the level of the payoff.

Generation of new rules

New rules are normally generated as a variation of existing successful ones, and can occur after a rule is applied, i.e. its action part has been played. A parameter determines the frequency by which an applied rule has the opportunity to generate an offspring, i.e. a new rule derived from its own structure. In simulation exercises on static environments the total number of offspring per rule is also limited in order to avoid a continuation of a search after a limit level is reached. The parameter is chosen to not affect the results, being triggered only when the performance of the organization stabilizes.

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. Two rule- and bit-specific indicators defined as follows control such mutations. Each un-specified bit h of the condition part (i.e. each $c_h = \#$) is associated to two indicators $r_{t,h}^0$ or $r_{t,h}^1$ recording the payoffs received when the rule is fired and the environmental state of the bit is 0 or 1, respectively. If the environmental string has 0 in the h th position, i.e. $e_h = 0$, then $r_{t,h}^0 = r_{t-1,h}^0 + \text{Payoff}$, where *Payoff* is the payoff received by the rule.

If instead $e_h = 1$, it is $r_{t,h}^1$ that gets accordingly updated: $r_{t,h}^1 = r_{t-1,h}^1 + \text{Payoff}_t$. 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.

When the rule has the opportunity to generate an offspring, the first operation is to compute the following indicator for each non-specific bit:

$$I_{t,h} = | r_{t,h}^0 - r_{t,h}^1 | / (r_{t,h}^0 + r_{t,h}^1) \tag{4}$$

This is an indicator of concentration, ranging from 0 (minimal concentration) to 1 (maximal). If the indicator is very high, above a certain threshold, it means that the rule recorded systematically higher payoff when applied to environmental states defined by a given value in that bit, and therefore there is grounds for generating a new rule by means of specification. If this is the case, and more than one bit passes the threshold, the system chooses which bit must be turned from a # to either a 0 or 1 by assigning each generic bit the following probability:

$$\text{prob}(t,h) = 1_{t,h}^\theta / \sum_{j=\#} 1_j^\theta \tag{5}$$

where $\sum_{j=\#}$ indicates the sum overall and only the condition bit equal to # and θ is a parameter affecting the concentration of probabilities. Obviously, the bit selected for specification is set to 0 if $r_{t,h}^0 > r_{t,h}^1$ and to 1 otherwise. The action part of the new rule generated by specification is a perfect copy of the parent’s action part.

If the rule cannot generate a rule by means of specification because no bit shows sufficient concentration, the system chooses randomly whether to create a new rule via generalization or mutation. Generalization consists of mutating one specified bit $c_i \neq \#$ into $c_i = \#$, selecting randomly the bit to generalize. Also in the case of generalization, the action part of the offspring rule is a perfect copy of the parent’s action part.

New rules produced via 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.

Besides the regular procedure producing new rules from existing ones, an additional procedure generates new rules to meet exceptional conditions. 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, a ‘mismatch ratio’ is computed for all existing rules. 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 of 1. A rule with one mismatching bit, but having 9 other specific bits matching the current environmental state, has a mismatch ratio of 0.1. The rule with the lowest mismatching index is generalized by turning the mismatching bits into # and has therefore been able to meet the current environmental conditions.

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:

$$A_{t,i} = \lambda A_{t-1,i} + (1 - \lambda)C(t,i) \quad (6)$$

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 the organization, retaining only rules used frequently.