Organization Science
Vol. 25, No. 2, March–April 2014, pp. 479–493
ISSN 1047-7039 (print) | ISSN 1526-5455 (online)

http://dx.doi.org/10.1287/orsc.2013.0841

Memory and Organizational Evolvability in a Neutral Landscape

Amit Jain
National University of Singapore, Singapore 117574, amit_jain@nus.edu.sg

Bruce Kogut
Columbia University, New York, New York 10027, bruce.kogut@columbia.edu

Many organizational theories are not sanguine over the chances of organizations to adapt and evolve, even if they should learn from the past through memory. Innovative search in the adaptive biology tradition leads quickly to dead ends. However, memory is useful for rendering innovative search more efficient. The concept from evolutionary biology of neutrality and drift along neutral pathways introduces the possibility that organizations are robust to local innovations and therefore potential candidates for evolvability. Through simulations in a neutral NK hypercube, our analysis shows that neutrality does not create value when future payoffs are discounted and change is costly. Here is the role for memory. Memory enables the faster development of better capabilities and reverses the negative assessment of evolvability. Even when allowing for forgetting, memory is a positive capability that improves evolvability of organizations so they can achieve better performance and better ways of doing so. Memory and neutrality are complementary for creating organizational evolvability, a finding consistent with the overwhelming evidence that organizations are more productive today than before because of innovation.

Keywords: organizational memory; organizational and radical change; evolvability and innovation; complementarities; robustness

---

Introduction

The most remarkable observation on the historical evolution of economic organizations is their pronounced increases in productivity. These gains accrue by learning through small innovations that cumulatively improve the innovative potential of the organization—that is, its evolvability. Yet borrowing from neo-Darwinian theory grounded in molecular biology, many adaptionist theories of organizations are pessimistic about the chances for organizations to adapt and to evolve because of competency traps. From these traps, organizations can only escape by taking on high-risk actions, by making errors in their myopic assessment, or by imperfectly imitating other organizations under selection pressures.[1] High-risk or lucky errors are not strategies of convenience but of dire times.

One might hope for sunnier days if organizations, unlike gene sequences, could profit from knowledge of what they have experienced before, that is, from their memories. Since memory is the storage of coded information, the provision of more information should aid decision making and practice. But academic opinions vary on the benefits of memory. Learning is a kind of memory insofar that a firm codes its past experiences into its current repertoires. A number of studies show that accumulated learning increases productive efficiency and that loss of this knowledge through forgetting lowers productivity (Argote et al. 1990, Darr et al. 1995). This positive assessment of learning is tempered by arguments indicating that memory constrains exploration and is deleterious to long-term organizational viability (March 1991). That is, memory acts as an overly conservative influence "to retrace the steps" instead of taking new and risky exploration.

We can understand why many studies have concluded that memory is bad for an organization. If an organization is myopic, even if it has perfect local knowledge, it has already made its best choice. There is little gain in perfectly revisiting its history of myopic choice. The organization is sitting contently on a little hill, thinking there are no higher vistas and oblivious to the mountain ranges in the distance. Memory endows the adaptive organization with moderate intelligence to lament its past but gives it no motivation to improve its local future.

If we can get this stuck organization to move down from its hill and look around, things might start looking brighter. Here, the adaptionist program promotes a pessimistic view that the molecular biology theory of neutrality rejects in favor of a far more optimistic view for the potential of developmental (i.e., organizational) change. The great majority of mutations to biological organisms, according to neutral theory, have no impact on their fitness or their survival chances; they are "neutral" (Kimura 1968, 1983; King and Jukes 1969). Accumulated over time, however, neutral mutations enable major change to an organism's phenotype, as manifest in its observable characteristics and traits. Consequently, the theory of neutrality directly contradicts the stark adaptive prediction found in organizational literature that replication requires exact copy of practices or else the organization suffers a severe performance detriment.

The goal of this article is to evaluate whether memory and neutrality are complements and are jointly required to improve innovation and fitness. More fundamentally, the goal is to move from an adaptionist program relying on selection and competency traps to a theory of organizational evolvability in which innovation plays a central role. To this end, we examine more fully through simulations the radical implications of neutrality and memory for organizational theory. These implications include a departure from the emphasis on adaptive search and change for a theoretical, and modeling, perspective of neutrality that allows for random drift to enhance exploration. As Lobo et al. (2004) have shown in the model that we harness in our analysis, neutrality improves the performance of an organization because it dramatically lowers the risk of selection for minor innovations.

But this sudden hope is quickly dashed once costs of innovation are considered. Any assessment of the innovative potential of organizations to evolve requires a consideration of investment costs and the time value of money. It is analytically unfair to bestow risk-free innovation on an organization and then announce its improved chances for higher performance. The biological literature has misled organizational research to ignore costs and time discounting, for nature is profligate, and the Darwinian mechanism is to explain survival of the species, in which individual organizations and their gene sequences may be sacrificed. Organizations, which have the capacity to make forward-looking decisions, not to mention recall their past, prefer not only survival but also the good life, and the sooner, the better. Here, the sensible observation is that a meaningful economic assessment of organizational evolution cares if the costs of search are worth it or if delayed satisfaction of a better outcome is worth many years in the desert. We thus introduce two central costs into our simulations: the costs of space (i.e., costs from local innovation) and of time (i.e., discounting future rewards more than current rewards).

This fuller characterization of the evaluation of search is standard in financial and business simulation modeling, such as Monte Carlo simulations of the value of an investment in real or financial assets. The dynamic problem is whether the shadow of the past offers a useful reservoir of diversity to offset the static costs of search that is detrimental to immediate fitness. This trade-off between static and dynamic choices imposes the analytical importance of evaluating the intertemporal costs and benefits.[2] We analyze this trade-off by evaluating the net present value (NPV) of the net benefits of search with and without neutrality and memory.

In our simulations, we consider neutrality and memory as two experimental treatments, and the comparison of the four resulting permutations forms the core of our analysis. Our proposed analysis centers first on evaluating whether neutrality and memory independently do not lead to optimism for organizational evolvability. In organizations, however, neutral change does not exclude making use of prior learning retained in memory. Given this, our second observation is that jointly, neutrality and memory increase the probability of evolution and innovation, even when subject to static inefficiencies, such as forgetting, delayed search, and costly exploration.

In the sections that follow, we first define the central concepts, borrowed with appropriate qualifications, from molecular biology. The second section discusses the effects of neutrality on search and prototypical change in organizations before turning to a discussion of the effects of memory on the ability of firms to evolve. In the third section, we describe graphically the key features of the NK model.[3] We show how this modeling approach can be adapted to explore organizational evolvability in regard to the presence or absence of memory and neutrality. To this end, we utilize the simulation design of Lobo et al. (2004), who used a modification of the standard NK model to explore the implications of neutrality. Since memory and neutrality are the two treatments, the simulations, which are presented in the fourth section, explore each of the four experimental conditions. In a final step, we show the robustness of our results when memory is characterized by knowledge decay using the forgetting rates estimated by Argote et al. (1990), Benkard (2000), and Thompson (2007) from empirical data.

Neutrality and Evolvability

Because the analysis below borrows and translates concepts from molecular biology, it is useful to ground our ideas and definitions first in an organizational example of the revolutionary innovation in manufacturing credited to the Ford Motor Company in 1913. Prior to the adoption of assembly line production, the manufacture of automobiles relied on a prototype of "craft production" that required days and numerous workers.[4] The slow accretion of knowledge in the design of interchangeable parts and tooling had taken essentially a century. This body of accumulated tacit knowledge led to modest productivity improvements until it was coupled with an idea that the Ford engineers observed at a nearby Chicago slaughterhouse involving the transport of material by conveyor belts. The adoption of an additional practice called the "assembly line" led to a nonlinear impact on productivity, but this causal effect was dependent on the accumulation of many other practices, such as standardization of parts, piece rates for worker effort, and Taylorist deskilling (i.e., substitution of craft with lower human capital).

These changes at Ford are easily expressed by a type of genetic, or organizational, sequence. It is traditional in both econometric and simulation modeling of organizations to code organizational practices as binary (or Boolean) genetic sequences.[5] For example, prior to 1913, the Ford manufacturing process consisted of interchangeable parts, piece labor payments, unskilled labor, and the absence of an assembly line, which we can code as $\langle 1100 \rangle$. Of course, the Ford Motor Company consisted of many other practices, and the complete (though never completely observed) genetic sequence contains additional digits. Looking back, these additional organizational dimensions were "hitchhiking genes" and not consequential in this particular historical time frame. After 1913, Ford switched from piece labor to $5 a day wages and adopted the assembly line, and the code became $\langle 1001 \rangle$. The (Hamming) distance between these two sequences is 2: it takes the "flip" of two digits for one sequence to evolve to the other. If only one practice is adopted, then the Hamming distance is 1; metaphorically, we describe Ford as having moved a single step along a pathway.

If there is no change in fitness when only one of these digits (say, the fixed pay of $5) is flipped, the change is said to be neutral. From an observational point of view, though, there are considerable differences between staying put and adopting new practices. This one-digit change etches out a neutral pathway between two points. This change is, then, similar to walking along a ridge in a fitness landscape, where higher heights denote higher fitness than lower heights. By this illustration, adopting a new pay scheme did not require Ford to descend into a valley or be stuck on a low-peak competency trap, unaware of the nearby mountain tops.

Here, we can understand how neutrality improves *evolvability*, which is the capability of an organization to innovate new and valuable functions.[6] With the subsequent adoption of the assembly line, the Ford process clearly defined a new prototype of mass production that radically improved productivity. This last change is not neutral, because the fitness of mass production was discretely superior to the previous system. Walking along a ridge permitted Ford to find a better way of producing cars without descending into the valley and incurring a loss in productivity. This is the intuition behind the theoretical concept of a neutral pathway.

Neutrality, then, is the set of organizational (genetic) sequences that result in equivalent organizational functional (phenotypic) solutions. Movements of organizations within this set etch neutral pathways, which, with some finite probability, arrive at a boundary crossing from one prototype to another over time. Neutrality is distinct from adaptive search because organizations do not know, or experience, fitness consequences from location mutations along these neutral pathways. Since selection acts on phenotypes (e.g., functional capabilities to innovate) but distinct genetic sequences produce equivalent phenotypic expressions, evolution is neutral with respect to selection prospects. This break between phenotypic selection and genetic variation gives genes a degree of protection from selection, thus permitting evolution to proceed by the neutral drift of local mutations.

On the surface, evolvability appears to contradict neutrality. In fact, the two are deeply related through the joint effects of robustness and plasticity, two opposing properties that are nevertheless jointly necessary for the evolution of organizations. An organization must be robust to mutation and preserve its core functional capabilities; otherwise, any one-digit deviation would lead to a rapid change in performance. Yet the organization also must maintain the ability to innovate and improve by changing what it does. An organization can benefit from both properties: the central function of a specific genetic (i.e., organizational) sequence remains robust, and plasticity allows the organization to benefit potentially from innovations in secondary functions that are neutral.

The joint influence of robustness and plasticity permits an organization to get the trade-off between exploitation and exploration on the cheap. Evolvability does not, consequently, easily fit the sharp distinction between exploitation and exploration as proposed by March (1991). From the perspective of robustness, the organization is *exploiting* what it has already done; from the perspective of plasticity, the organization is *exploring* neutrally its local space, preserving the potential for radical innovative change from one organizational *prototype* to another prototype.

It is no wonder that these concepts of neutrality and evolvability have been controversial in molecular biology, as they should be for adaptive theories in organizational theory. Neutrality is about how things change, as opposed to the emphasis in adaptive models on why they change.[7] Adaptive search models assign organizations "addresses" (i.e., genetic sequences in a space), a rule of search (i.e., choose higher fitness in one's local "neighborhood"), and a decision to move there if it is better. Once at a local suboptimal peak, progress stops; otherwise, the organization must descend into valleys of lower fitness to search for new and maybe better peaks. Neutrality considers, however abstractly, a genetic space that maps onto phenotypes (e.g., enzymatic functions), but this mapping is not a bijection; there are many sequences that produce the same function. As these innovations accumulate, the organization crosses along ridges of neutral pathways, potentially arriving at radical new prototypical ways of *how* things are done, such as mass production of autos.

Thus, neutrality proposes a theory about the robustness and yet plasticity of organizational structure. In this regard, it captures well, for example, the hidden finding in the seminal MacDuffie (1995) article that analyzed the impact on productivity and quality of managerial practices: in either mass production or Toyota production systems, teams made a contribution to higher performance. In the metaphor of landscapes, teams provided a ridge between two prototypes, permitting some auto plants to adopt new practices without suffering profound drops in performance as predicted by adaptive models.

The Shadow of the Past

Although memory plays no discernible role in molecular biology, it has spawned a large literature in organizational theory. Organizational memory refers to "stored information from an organization's history that can be brought to bear on present decisions" (Walsh and Ungson 1991, p. 61). This definition favors a view of memory as information storage as if knowledge is encoded without decay.

We broaden this view to include the problem that the knowledge of "how to do" something can decay.[8] After all, knowledge is frequently acquired through learning by doing, resulting in greater productivity (Argote 2013, Dutton and Thomas 1984, Jain 2013). These sources of productivity improvement may consist of many distinct practices. Argote (2013, p. 18) lists these sources as "improvements in the performance of individual workers; improvements in the technology, tooling, and layout; improvements in the organization's structure; and better understanding of who in the organization is good at what."

Knowledge acquired through learning by doing is retained in different bins in organizational memory (Reagans et al. 2005). The retrieval and sharing of differentiated individual knowledge and experience working together has been the subject of study of transactive memory systems (Liang et al. 1995, Moreland et al. 1996, Wegner 1986). The retention of information and its retrieval from memory involves search costs and exchange costs. Some of these costs are more subtle. The provider of information needs to incur an opportunity cost of transfer as she could dedicate her efforts to other productive activities (Levine and Prietula 2012).

The Ford Motor Company manufacturing example illustrates processes of acquisition, retention, and retrieval of information encoded in practices. For example, consider a hypothetical Ford ab initio with an organizational sequence of $\langle 0000 \rangle$, where the 0s represent, respectively, the absence of interchangeable parts, piece labor payments, unskilled labor, and an assembly line. Ford decides sequentially to implement manufacturing processes consisting of interchangeable parts and piece labor payment, and its new sequence is $\langle 1100 \rangle$. It learns through doing about the new states and retains this information in memory. In due course, it then implements the assembly line, achieving an organizational sequence of $\langle 1101 \rangle$, but because piece rates are hard to use with assembly line production, productivity falls perceptively. As a result, rather than wastefully reinventing past practice, it can efficiently draw from experience and retrace its steps to a prior location—say, $\langle 1000 \rangle$. It then implements a policy of paying low-skill labor at lower wages, and it notes now a perceptible improvement in productivity, resulting in $\langle 1010 \rangle$. When it adopts the assembly line, it then has the sequence of $\langle 1011 \rangle$. Ford has made history; it has invented mass production.

Memory and Exploration. The dominant view is that memory constrains exploration. By dampening the effects of "chance caught on a wing" in Jacques Monod's (1971) expression, memory can be harmful to organizations. Belief structures that develop as a result of prior learning can blind decision makers to aspects of the environment and thus compromise effectiveness (Walsh and Fahey 1986). In this vein, Nystrom and Starbuck (1984, p. 53) wrote, "Encased learning produces blindness and rigidity that may breed full-blown crises." Similarly, Kahneman and Tversky (1979) found that people "anchor" their beliefs in past experience, biasing their subsequent decisions.

Consequently, the search for novel solutions to problems at hand is compromised and organizational performance declines (Walsh and Fahey 1986). Noting the presence of these effects, Nelson and Winter (1982, p. 135) wrote that "firms may be expected to behave in the future in ways that resemble the behavior that would be produced if they simply followed their routines of the past." Finally, in his much cited paper on exploration and exploitation, March (1991) noted that "each increase in competence at an activity increases the likelihood of rewards for engaging in that activity, thereby further increasing competence and the likelihood [of further rewards]" (p. 73). However, like anchoring bias, too much competence is deleterious to subsequent decisions and novel search. In a provocative expression of this point of view, March (1972, p. 427) treated "memory as an enemy" of organizations and stated that "the ability to forget, or overlook, is also useful."

A number of arguments have a contrary view: memory facilitates problem definition, alternative generation and evaluation, and choice (Neustadt and May 1986, p. 32). To use the felicitous phrasing of Hedberg et al. (1976, p. 41), the benefits of memory derive from "footholds in time" that serve as starting points for establishing new trajectories of search. Many opportunities may be missed along the way, and recognition of these "missed opportunities" serves as "footholds." Similar observations have been made by others that "better decision making involves drawing on history to shape sharper questions" (Neustadt and May 1986, p. 32). Consistent with this, using data on product innovation practices at the design firm IDEO, Hargadon and Sutton (1997) found that the memory of prior designs retained in the memory repository facilitates innovation.

Models of memory indicate an ambiguous effect: although information retained in memory improves productivity, it may either constrain or facilitate exploration. An assessment of memory thus requires a consideration of its dynamic effect on cumulative learning. But learning comes at a cost and is subject to being lost as it is forgotten over time. These are important considerations to the evaluation of memory that we will return to in the next section.

The Experimental Design

Let us summarize the above arguments in the form of an experimental design. Consider sampling from a patient population of organizations that are pure adaptionists: they myopically search in their immediate neighborhood, and they have no memory of where they have been. Neutrality and memory are two treatments that can be given independent of each other, as well as jointly. This gives rise to four experimental conditions (see Figure 1). The first quadrant (Q1) of the figure corresponds to adaptive search when firms have information regarding their prior position only and do not have memory and neutrality. The second quadrant (Q2) corresponds to the case of search with neutrality. The third quadrant (Q3) pertains to firms searching adaptively with memory, and the final quadrant (Q4) refers to search with both memory and neutrality.

We now summarize the dynamics that characterize each quadrant.

Adaptive Search (Q1). In adaptive search, a firm moves to a new location only if the location provides greater payoffs. Given complementarities, in rugged landscapes adaptive search terminates relatively quickly in a competency trap. Better global solutions exist, but the local neighborhood offers only worse outcomes. Because search is limited and fitness is low, a firm is relatively inert to change and has low evolvability.

Figure 1 The Experimental Design: Memory and Neutrality as Treatments

	No neutrality (adaptive search)	Neutrality (neutral search)
No memory	(Q1) TRAPS	(Q2) COSTLY EVOLVABILITY
Memory	(Q3) TRAPS LIMITED EVOLVABILITY	(Q4) EVOLVABILITY

Notes. Neutrality and memory are two treatments that give rise to four experimental conditions. We argue that adaptive search (Q1) leads to competency traps and that memory (Q2) does not change this prediction substantively. In comparison, neutral search (Q3) leads to escape from competency traps, but the costs of search are greater than the realized improvements in performance. Neutral search with memory (Q4), however, leads to evolvability, as it results in better outcomes faster.

Neutrality (Q2). As discussed earlier, search along a neutral pathway leads to a greater exploration of different alternatives, increasing the likelihood of finding new solutions that have greater fitness relative to adaptive search. If search incurs costs, the fitness-enhancing effect of neutral search is mitigated by the cost of extended trial and error neutral search, resulting in costly evolvability.

Adaptive Search with Memory (Q3). Adaptive search satisfices by adopting the first better local solution rather than the optimal one. As a result, this search may miss out on better solutions because of missed opportunities along the search pathway. When adaptive search leads to a competency trap, memory guides the organization to explore these missed opportunities. As a result, it might be expected that memory leads to greater fitness. However, adaptive search with memory must satisfy the condition that a move to a future location results in greater fitness than the current location. Because of this, memory does not change the pessimistic assessment of adaptive search.

Neutrality with Memory (Q4). Search with memory in neutral landscapes enables organizations to draw from their prior experiences and avoid wasteful efforts and relearning. Because of neutrality, organizations drift within the prototypical category, and these moves do not need to meet the condition of moving to a better payoffs, as in the case of adaptive search. Memory is now useful. This combination of memory and neutrality leads to evolvability.

In sum, these arguments suggest that neither neutrality nor memory independently improves evolvability. Together, however, they complement one another and improve it substantially. The role of neutrality promotes forward-looking search into new domains of activity, but this search is inefficient, and performance deteriorates when search costs are incurred. Memory, in contrast, is backward-looking, experience-based search that improves efficiency by benefiting from experience. Looking to the future under the shadow of the past—searching neutrally with memory—leads to better solutions faster, to evolvability.

The NK Model with Neutrality and Memory

We analyze these competing effects of memory and neutrality through a simulation in an NK modeling environment. An NK model has a Boolean random graph as the underlying structure. Graphs as mathematical structures have particular spatial properties that, in low dimensionality, can be depicted visually. An NK model is a hypercube, or a Boolean lattice that in two dimensions looks like a square or in three dimensions like a cube.[9]

The appeal of a Boolean lattice to evolutionary biology is that it provides a compact structure by which to describe the graph, more familiarly called the "network," of distances between gene sequences. A Boolean lattice provides a binary address to every point. For example, a square has four points with addresses $\langle 00\rangle, \langle 01\rangle, \langle 10\rangle, \langle 11\rangle$; each point has two neighbors to which it is directly connected. The addresses serve as representations of genetic sequences, and a single-digit change in a sequence represents a move of Hamming distance of 1. The number of addresses, or unique genetic sequences, explodes exponentially. A sequence of 2 has 2^2, or 4, addresses; a sequence of length 10 has 2^{10}, or 1,024, addresses; a length of 100 has 2^{100} addresses. Yet when we compare the diameters (i.e., maximal hamming distances), they form the ratio of 2:10:100. The property of the hypercube then is to generate tremendous genetic diversity with short lengths (or path lengths).

The N in NK represents the number of dimensions and K the number of interactions (as defined below). We might want to know if using a genetic sequence of 10 is a big or small N for organizations. Most studies in organizational theory probably consist of relatively few dimensions, for example, levels of hierarchy, size, or span of control. The studies that analyze precisely the relation of the characteristics of the organization to performance generally reduce the dimensionality to fewer than 10 factors (e.g., Ichniowski et al. 1997, MacDuffie 1995). In this study, N is taken as 10, and there exist 2^{10} (or 1,024) possible locations in the landscape.[10]

The complex and interdependent nature of choices in search is represented by the degree of interaction or interdependence K among the N dimensions of the search space. The choice of K indicates the number of complements shared among two neighbors (i.e., those sequences that differ by one digit). When the value of K is 0, there is no interdependence among two neighboring points, and the landscape is smooth and correlated. When the value of K is high, then a change in one decision variable influences the payoffs of K other decision variables. This gives rise to a relatively uncorrelated and rugged landscape with a large number of local optima.

The NK landscape provides a mapping of addresses onto the domain of payoffs, which is commonly visualized as a fitness landscape. If we stay with the metaphor of a landscape, then the height of a peak represents its fitness (e.g., the productivity of a particular sequence of managerial practices). The calculation of these fitness values works by assigning randomly from a uniform distribution a fitness value to each dimension. If $K = 0$ and there are no interactions, then the fitness of an address, or sequence, is simply the normalized sum of the fitness of each dimension. In the case of interactions and complementarities, the fitness contribution of each of the N attributes is dependent on K other attributes. There then exist 2^{K+1} payoff values for each of the N attributes present in a configuration, depending on the value K other attributes take (0 or 1), and on the value of the given attribute itself (0 or 1). Each of these 2^{K+1} combinations is assigned a fitness value drawn randomly from a uniform distribution from 0 to 1. The payoff F to a variation is expressed as an average of the contributions of its N attributes and may be formally represented as $F(x_1, x_2, \ldots, x_N) = \sum_{i=1}^{N} f(x_i)/N$, where $f(x_i)$ is the fitness contribution of the ith attribute.

Search and Adoption of New Practices

In organizational simulations, the conventional terminology is to view the organization as searching in the neighborhood of current practices for better ways to do things. We now specify how this search takes place, recalling the earlier discussion on experimental design.

Adaptive Search. At each time period in the simulation, decision makers flip one of the N attributes in the organizational sequence from 0 to 1, or vice versa, and the organization moves to the first new location that is better. If no neighboring position provides a better payoff, the firm is stuck in a competency trap, and search stops.

Neutral Search. Simulation with neutral search requires a modification of fitness values mapped to nodes in an NK hypercube: a structure of neutral categories based on fitness values. Given this, consistent with prior studies of neutrality (Lobo et al. 2004), we assign all sequences of organizational attributes that have similar fitness values to a neutral set. We divide payoffs into C distinct neutral sets and compute neutrality Ne as ($\text{Ne} = 1/C$). For illustration, we partition fitness into five neutral categories, c1 to c5, corresponding to the fitness intervals [0–0.2), [0.2–0.4), [0.4–0.6), [0.6–0.8), and [0.8–1.0]. Here, the granularity of neutrality is $\text{Ne} = 1/5 = 0.20$. (In simulations, we set $\text{Ne} = 1/10 = 0.10$.) Following Kogut et al. (2004), we call these partitions "prototypes," such as customized production or mass production. During search, if there are no points in the neighborhood that belong to the higher fitness category, then the firm drifts to a location that does not decrease payoffs—that is, to a position within the current neutral category. Search stops when all neighboring positions are in lower fitness neutral categories; the firm is then caught in a competency trap.

Search with Memory. Firms retain in memory information pertaining to the configuration of N attributes and the payoffs of each prior choice in search. We assume that information retained in memory does not decay over time and that memory has a finite capacity limited to 25 locations. In search, firms check whether a new alternative has already been visited (i.e., is present in memory) and do not repeat prior searches by moving to these positions. If the firm is stuck in a competency trap, it starts from the last position retained in memory and searches for a neighboring position that provides better payoffs. If such a "missed opportunity" exists, it moves to it, or else it moves to a neutral location. If no better or neutral position is found, then the firm successively draws from memory the immediately prior position retained and repeats the search steps. If all positions in memory have been retrieved in this way and there are no positions to move to, the firm is stuck, and it stops searching.

The Adoption of New Practices in a Neutral NK Landscape: An Example

We illustrate the role of neutrality and memory by modifying our earlier Ford example by considering only three organizational practices. The workhorse of our analysis relies on describing organizations as a sequence of practices coded by binary digits; the binary (or Boolean) value indicates abstractly the presence or absence of an organizational practice—for example, interchangeable parts, piece rates, and deskilling of labor. In this representation, an organizational sequence serves as an address in a geometric space that can be represented by a (hyper)cube. Organizations move by "flipping" the binary values of a practice element in this sequence, resulting in a change in the address given by their organizational sequence.

Consider a simple search space characterized by these three Boolean organizational dimensions, or practices. The Boolean configuration ⟨000⟩ thus is an organizational sequence that represents no interchangeable parts, no assembly line, and no low skill labor (i.e., craft labor). With three different Boolean practices, there are eight (2^3) distinct sequences, represented by the vertices of a three-dimensional hypercube, which in this case is a simple cube (see Figure 2(b)). This time, we will not let Ford discover mass production, but we will illustrate visually and numerically its search.

Adaptive and Neutral Search. Adaptive search is the simple case: an organization searches to move to a better set of practices by looking at the neighboring organizational sequences to its current location and deciding to move to the first location that has a higher fitness value; if it finds one, the organization moves to that location. If not, it searches again. If no neighboring location is better, it stops at its current location. Neutrality is slightly more complicated, for here, we wish to model the mapping of many organizational sequences to the same fitness value, implying a functional equivalence along a neutral pathway or set of practices. Small changes in organizational practices (genotype) result in no change to its categories of activity (phenotype) and fitness. Using the five partitions of neutrality (i.e., $\mathrm{Ne} = 1/5$), the neutral categories that correspond to each location in the hypercube are listed in Figure 2(a).

Fitness. Each organizational sequence is assigned a fitness value, which corresponds to organizational performance. Fitness is a function of the N (i.e., three) dimensions, or practices, of an organizational sequence. The contribution of each dimension, however, depends on its interactions with K other dimensions. Without interactions, the contribution of each practice is independent of others; in the adaptive case, the organization hill climbs incrementally to the optimal solution. As the number

Figure 2 Adaptive Search with Memory ($N = 3$)

(a) Fitness of configurations in a three-dimensional Boolean space

Position	Fitness	Neutral category
000	0.47	c3
001	0.59	c3
010	0.42	c3
100	0.18	c1
110	0.79	c4
011	0.55	c3
101	0.48	c3
111	0.28	c2

(b) Representation of choices as a cube

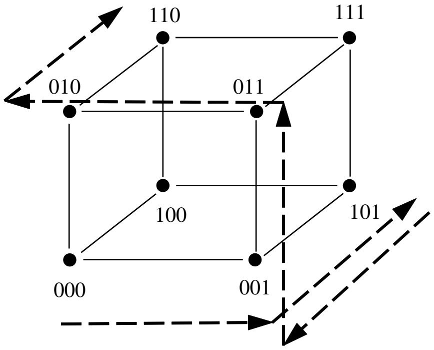

(c) Search using each of the four experimental conditions

Search type	Sequence of search	Ending fitness	NPV
Q1. Adaptive search	⟨000⟩ ⟨001⟩	0.59	137
Q2. Neutral drift	⟨000⟩ ⟨001⟩ ⟨101⟩	0.48	111
Q3. Adaptive search with memory	⟨000⟩ ⟨001⟩	0.59	137
Q4. Neutral drift with memory	⟨000⟩ ⟨001⟩ ⟨101⟩ ⟨011⟩ ⟨010⟩ ⟨110⟩	0.79	180

of interactions K increases, a number of local optima characterize the search. These local optima are competency traps, as all neighboring configurations provide inferior payoffs. At these positions, an organization cannot find a successful path toward further evolution and improvement.

For our example, we consider three different practices ($N = 3$) with high complexity ($K = 2$), indicating that each dimension interacts with the other. In this case, adaptive search will get stuck in suboptimal solutions or competency traps at a high probability.[11] One set of possible payoffs to each of these organizational sequences is listed in panel (a) of Figure 2, and the positions are represented graphically in panel (b) as a cube. It may be observed from the payoff table in Figure 2(a) that there are two optima: a local optimum at $\langle 001\rangle$ and a global optimum at $\langle 110\rangle$. All neighboring positions to these two positions provide lower payoffs.

Metrics. We compare the outcomes of the experimental conditions using two metrics: fitness and net present value. NPV accounts more accurately for the costs of exploration. To compute the NPV, we observe fitness on a scale of 0–1 (see Figure 2(a)), search costs of 0.5 per search move, and an annual discount rate of 5% (0.43% per simulation time period).[12] We illustrate the trajectory of search for the case with neutrality and memory (Q4) in Figure 2(b) and provide the fitness and NPV statistics for each of the four experimental conditions (see Figure 1) in Figure 2(c). High fitness and NPV indicate evolvability.

Q1: Adaptive search. Adaptive search looks at the local neighborhood of an organization and checks sequentially if a change to one practice provides greater fitness. If a location provides better fitness, then Ford satifices by choosing this new organizational location; it does not sample the other locations. (Recall that a neighborhood is the set of all sequences at a Hamming distance of 1.) For instance, $\langle 001\rangle$, $\langle 010\rangle$, and $\langle 100\rangle$ are neighbors of the point $\langle 000\rangle$ that may be realized by changing one practice. Starting from $\langle 000\rangle$ (fitness $=$ 0.47), Ford proceeds to $\langle 001\rangle$ (fitness $= 0.59$) by implementing the deskilling of labor, a change to the third digit. At this new sequence, moving to any of the three local neighbors—$\langle 011\rangle$, $\langle 101\rangle$, and $\langle 000\rangle$—does not improve payoffs, and Ford is stuck in a competency trap. Thus, the final fitness is 0.59, and the NPV is 137.[13] The business press report is that Ford has overoptimized its operations.

Q2: Neutral search. Like adaptive search, neutral search also takes place through a change of one element in a sequence, but now change can proceed within any of the neutral categories c1–c5. Thus, although adaptive search stopped at the position $\langle 001\rangle$, neutral search continues to $\langle 101\rangle$ along the neutral pathway c3. At this point, the neighboring solutions $\langle 100\rangle$ and $\langle 111\rangle$ are located in lower fitness categories (c2), and search stops. In this case, the final fitness is 0.48, and the NPV of search is 111. Thus, in this case, neutrality led to lower fitness, to greater search costs, and to low evolvability. The press reports that Ford's strategy is (neutrally) "drifting."

Q3: Adaptive search with memory. Adaptive search (Q1) proceeded from $\langle 000\rangle$ to $\langle 001\rangle$ and then stopped at $\langle 001\rangle$, since it was a competency trap. Memory enables retracing steps back to the prior location $\langle 000\rangle$, but neighboring locations ($\langle 010\rangle$ and $\langle 100\rangle$) have lower fitness than $\langle 001\rangle$. Memory does not help in this situation, since there were no missed opportunities along the way that provided payoffs greater than the competency trap $\langle 001\rangle$. As a result, the final fitness and NPV of adaptive search with memory (Q3) is identical to that of adaptive search (Q1) in this example. Ford is stuck in its past.

Q4: Neutral search with memory (and no forgetting). Neutral search (Q2) stopped at $\langle 101\rangle$. With memory, Ford can retrace its steps back to $\langle 001\rangle$ and investigate missed opportunities for search such as $\langle 011\rangle$. Since $\langle 011\rangle$ lies in the same neutral category c3 as $\langle 101\rangle$, Ford can move to it. In this case, neutral search with memory has an option value, as it opens up new areas of the search space. Accordingly, search can proceed to $\langle 010\rangle$ within the same category (c3) and then attain the global maxima at $\langle 110\rangle$ (c4), a new prototype. This search trajectory is given in Figure 2(b) and 2(c). With memory, neutral search attains a final fitness of 0.79, and the NPV of search is 180. Ford's stock is on the rise.

Comparing the evolvability metrics for the four experimental conditions (see Figure 2(c)), it may be observed that adaptive search, with or without memory, leads to little exploration and to low fitness and low evolvability. The firm is stuck in search space close to its initial configuration—inert. Furthermore even though neutral search results in more exploration, it may actually result in low evolvability. Finally, with memory, neutral search results in extensive exploration and the highest possible fitness (the global optima). This also results in the highest NPV; in sum, it results in evolvability.

Results of the Simulations

The road map to the simulations given in Figure 1 describes the four experimental combinations of our two treatments: neutrality and memory. Our goal is to ascertain if memory and neutrality are complements and whether this complementarity is sustained when adoption is costly. In a first step, we study the relative performance of two benchmark memoryless populations, the first searching adaptively (Q1) and the second searching with neutral categorization (Q2). Following this, we compare adaptive search with memory (Q3) to the memoryless case (Q1). In a final step, we study the relative performance of a population searching with both neutrality and memory (Q4) relative to the memoryless case (Q2).

In the simulations, each population is made up of 100 organizations. The first population is at the start randomly located in the search space, and firms in the second population are allocated identical initial positions. Search takes place in a landscape with $N = 10$ and complexity K ranging from 2 to 8. We do not simulate the case of $K = 0$, as it has little empirical interest: adaptive search always converges to the global optimum. Each simulation consists of 100 runs for 250 time periods. The results reported are population averages across these runs. In the case of neutral search, we assume 10 neutral categories (Ne = 0.10). What this means is that organizations drift randomly if a new position is, on average, less than 5% better or worse relative to a given position. In robustness tests we relax this assumption. Consistent with our example of the adoption of new organizational practices in the previous section, search costs are 0.5 per move, and the annual discount rate is 5% (0.43% per month or simulation time period). We also consider memory to have a finite size of 25 positions. The results reported are population averages across these runs.

We compute our two evolvability metrics (i.e., fitness and NPV) as ratios across the two populations. If the two populations under consideration engage in neutral (Q2) and adaptive (Q1) searches, respectively, then the fitness ratio is the average fitness of firms in the neutral population (Q2) divided by the average fitness of firms in the adaptive population (Q1) at the same point of time. A fitness ratio greater (lower) than 1 implies that firms searching with neutral categorization have greater (lower) fitness relative to the adaptive population.

Figure 3 **Costly Evolvability: Neutral Drift (Q2) Leads to Exploration and Fitness, But at a High Cost**

(a) Fitness ratio

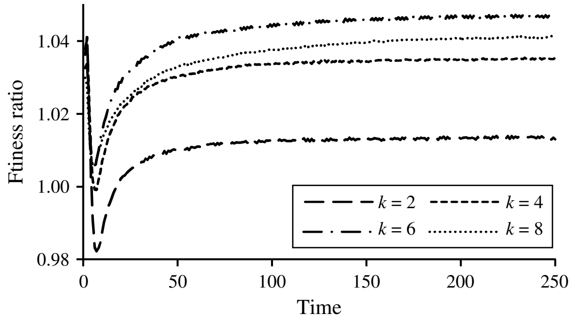

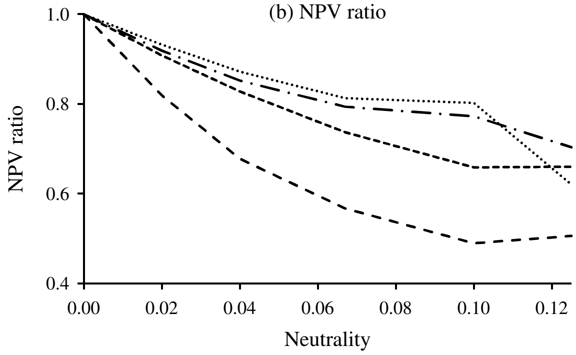

Notes. Simulations are done using a neutrality value of 0.10 and k values from 2 to 8. The statistics plotted are ratios across the two populations at a point of time (i.e., fitness ratio = (Average fitness of Q2 firms)/(Average fitness of Q1 firms)). Panel (a) provides evidence that extended search along a neutral pathway results in greater fitness for all values of complexity K greater than 0. Panel (b) presents the NPV ratio computed by taking into consideration normalized fitness, search costs (0.5 per move), and a discount rate of 0.43% per period for a range of neutrality values. The NPV ratio is less than 1 for all values of complexity K, indicating that neutrality destroys value. Thus, neutrality results in costly evolvability.

Costly Evolvability: Adaptive Search (Q1) vs. Neutral Search (Q2)

We start by comparing adaptive and neutral search. We plot in Figure 3 the two evolvability statistics, fitness and NPV ratios, based on our simulations of adaptive search (Q1) and neutral search (Q2). We argued earlier that neutral search leads to more exploration of the search space and to greater fitness relative to adaptive search, and this would suggest high evolvability. Figure 3(a) shows that for all values of complexity, the fitness ratio increases over time and ranges from 1.01 to 1.05 at the end of the simulations. Adaptive organizations (Q1) move to a new position in search space only if it provides better payoffs. Given complementarities among K practices, adaptive search in rugged landscapes terminates relatively quickly in a competency trap. Here, we see the workings of the pessimistic argument for adaptive search.

In comparison, firms travel along a neutral pathway (Q2) by random drift until they come to the boundary of a new prototypic category, indicating a radical innovation. As a result, they do not get stuck as often in competency traps and are more likely to engage in extended exploration. This extended search increases the likelihood of arriving at a new position that has substantially better payoffs, and thus the fitness ratio is greater than 1 (see Figure 3(a)). It is of note that adaptive search by myopic organizations leads to a quicker initial increase in fitness relative to those drifting firms engaged in neutral search; the fitness ratio initially falls below 1 and remains less than 1 until about the 25th time period. Greater fitness for the drifters comes through extended exploration. These results confirm those of Lobo et al. (2004).

Our analytical results differ from those of Lobo et al. (2004) once search costs and the time value of money are included. Since search along a neutral pathway requires considerable time to result in greater fitness, it incurs higher costs when costly adoption is added to the simulation. Figure 3(b) shows that relative to adaptive search, costly neutral drift destroys value; the NPV ratio is *less* than 1 for all values of complexity K and for values of neutrality greater than 0, and the NPV ratio decreases as neutrality increases. That is, neutrality does not lead to evolvability when costs are considered. Better fitness outcomes are found, but the overtime costs erode the incremental value. We show later in additional robustness tests in the online supplement (available as supplemental material at http://dx.doi.org/10.1287/orsc.2013.0841) that this result is valid over a broad range of search costs (from 0 to 0.6 per move) and discount rates (from 0.43% to 25%) per simulation time period.

Traps: Adaptive Search With Memory (Q3) and Without Memory (Q1)

Figure 4 plots the results of simulations of adaptive search in two populations of organizations, one searching with memory (Q3) and the other without it (Q1). In both populations, adaptive search leads rapidly to competency traps as search stops by the 17th time period. Memory does not significantly help search. Firms draw from memory when stuck in a competency trap, but they need to find a position starting from memory that is better than their prior position (the competency trap). This is a stringent condition to satisfy. Nevertheless, memory does capitalize on missed opportunities, resulting in a modest increase of 1%–2% in fitness (per Figure 4(a)) relative to the case without memory.

Because the duration of search is similar and memory leads to greater fitness, it also results in greater NPV. Figure 4(b) plots the NPV ratio, taking into consideration all simulation time periods. The NPV ratio is greater than 1 over all simulated values of complexity; memory leads to modest increases in NPV ranging from 2% to 4%. All in all, even though memory leads to small improvements in fitness and NPV ratios, it does not change the earlier negative assessment that adaptive organizations are subject to competency traps.

Evolvability: Neutral Drift With Memory (Q4) vs. Without Memory (Q2)

To complete the analysis, we compare neutral drift with memory (Q4) to that of neutrality without memory (Q2) in respect to fitness and NPV ratios; the granularity of neutrality is set to 0.10, K varies from 2 to 8; as before, search costs are 0.5 per move, and the annual discount rate is 5% (0.43% per time period). The simulation results are presented in Figure 5. Under neutrality, memory *speeds up* exploration and results in an increase in a quick initial increase in the fitness ratio (panel (a)). Neutrality with memory (Q4) outperforms neutrality without memory (Q2) in the short run. This occurs as firms do not engage in wasteful relearning and reinvention—they do not visit (neutral) configurations already visited. Furthermore, firms stuck in competency traps draw from memory to start a new search trajectory. To illustrate this effect, we plot in Figure 5(b) the proportion of firms in a population that become stuck in competency traps and utilize memory to benefit to revisit missed opportunities. Even at the very end of the simulation at $t = 250$, a number of firms still benefit each period from revisiting missed opportunities.

Figure 4 **Traps: Adaptive Search with Memory (Q3) Does Not Change the Negative Assessment that Adaptive Search (Q1) Leads to Competency Traps**

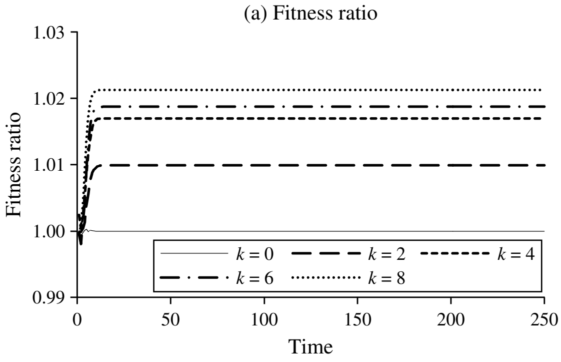

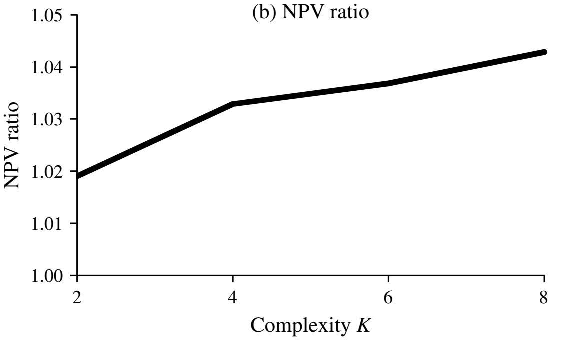

Notes. Simulations are done using complexity K values ranging from 2 to 8, and Q3 firms retain information of their last 25 locations in memory. The fitness ratio = (Average fitness of Q3 firms searching adaptively with memory)/(Average fitness of Q1 firms searching without memory). NPV ratio is calculated in a similar manner. Memory leads to a small increase in fitness (panel (a)) and in NPV (panel (b)) for all values of complexity. Since these increases are small, memory does not revise the negative assessment of the tendency of adaptive search to become stuck in local traps.

In Figure 5(c), we plot the NPV ratio. Memory leads to an NPV ratio greater than 1 for all values of neutrality and complexity, indicating that it improves the performance of neutral search and compensates for one of its shortcomings, namely, that neutral evolvability is costly, as it takes a considerable period of time. In all, memory leads to evolvability in neutral search and thus creates additional value.

Shadows of the Past: Forgetting

To this point, we have shown that memory and neutrality are complementary and that together they result in evolvability even when search is costly. In this analysis, we assume that memory is a perfect repository of information and that firms retain and recall information of each combination of practices flawlessly. However, because knowledge is experiential, memory is also

Figure 5 Evolvability: Memory and Neutrality Jointly Enable Evolvability

(a) Fitness ratio

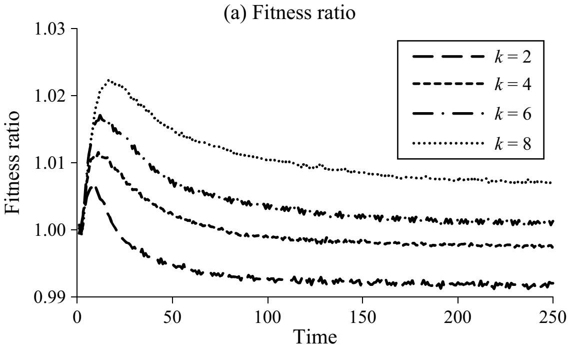

(b) Missed opportunities

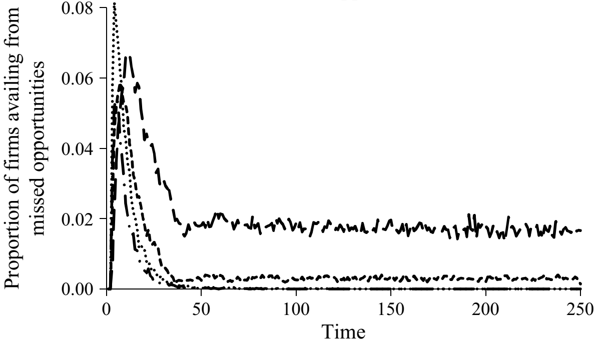

(c) NPV ratio

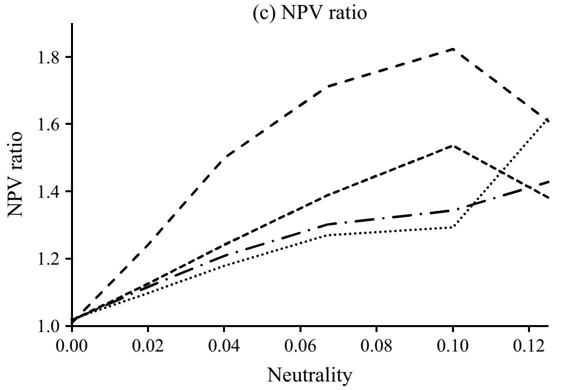

Notes. In simulations, neutrality = 0.10 and memory may retain information about 25 prior positions visited in search. Here, we plot fitness and NPV ratios for a population of firms searching with neutrality and memory (Q4) relative to another population searching neutrally only (Q2). Memory leads to quick improvements in fitness (panel (a)), as it makes neutral search more efficient and avoids relearning. These fitness increases also result from the possibility of restarting search from positions in memory when stuck in a competency trap. In panel (b), we plot the proportion of firms in a population that used memory to benefit from missed opportunities in search. To the very end of the simulation at $t = 250$, firms benefit from missed opportunities. Finally, memory leads to a greater NPV: the NPV ratio (Q4/Q2) is greater than 1 across all parameter values. Thus, memory and neutrality jointly result in evolvability.

subject to forgetting. Here, we test whether the effect of memory on evolvability is robust when memory is imperfect and finite in size, two forms of forgetting. Identical to our prior analysis, we use our baseline topological (i.e., NK) parameters with $\mathrm{Ne} = 0.10$ and $K = 4$ and the baseline search costs of 0.5 per move and an annual discount rate of 5% (0.43% per period).

Imperfect Recall. The quality of information recalled from memory may decrease over time as a result of forgetting. Here, we have empirical research on knowledge depreciation in the context of experience curves; this research reports rates of forgetting ranging from 4% to 25% *per month* (Argote et al. 1990, Benkard 2000, Darr et al. 1995, Jain 2012, Thompson 2007). Corresponding to these empirical findings, we operationalized forgetting by assuming that organizations imperfectly recall each element of a configuration retained in memory with an error rate ranging from 1% to 25% *per month*. (We have italicized "per month" to emphasize the match among parameter choices to time, anchored in empirical data where possible.) Figure 6 presents the results of analysis.

With recall error, search is extended, since errors open up new trajectories of search. Recall error also

Figure 6 Forgetting: Memory and Neutrality Jointly Enable Evolvability When Information Is Forgotten and Recall Is Imperfect

(a) Fitness ratio

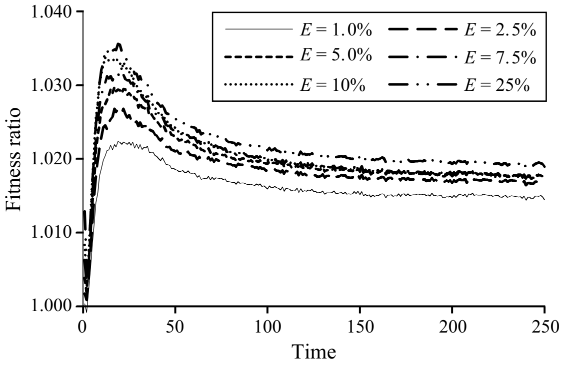

(b) NPV ratio

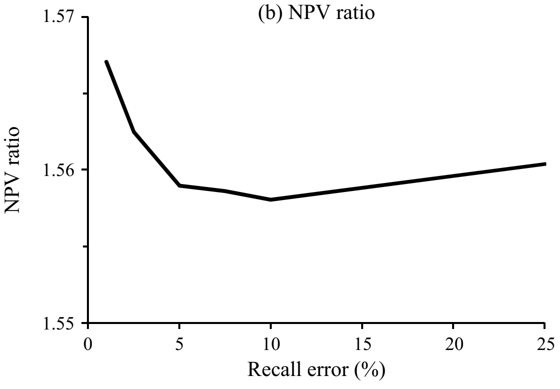

Notes. Forgetting is modeled as the imperfect recall of information retained in memory. Each attribute of every position retained in memory is recalled with an error rate between 1% and 25% (using neutrality = 0.10, memory = 25, $k = 4$, search cost = 0.5 per move, and discount rate = 0.43% per month). We compute the fitness and NPV ratios as the value for the population with imperfect recall (forgetting) to that of the population with no memory. With forgetting, search is extended as erroneously recalled positions open up new trajectories of search, and this results in greater fitness (panel (a)) and NPV irrespective of the rate of forgetting (panel (b)).

introduces an inefficiency: firms do not know with certainty if they have already visited a new position. Consequently, some relearning occurs as a result of recall error. As can be seen in Figure 6(a), as the rate of recall error increases, so does the fitness ratio: recalling prior positions with error opens up new trajectories of search that lead to greater fitness over time. Even though such extended trial and error is costly, the improved performance is of positive benefit in net. As a result, the NPV ratio plotted in Figure 6(b) is greater than 1 for all rates of recall error. All in all, our results indicate that even though there is substantial value to precise retention, forgetting can be beneficial to organizations. Anything that gets organizations out of their local optima traps and back onto the exploratory path has its rewards in gross.

Finite Repository. One of the common observations about memory is the constraints on the breadth of its registry. If the repository is full, then some other information must be "forgotten" for new information to be retained. Expanding the capacity of memory to retain more information, though, comes at a cost.

Figure 7 Forgetting (Finite Memory): Memory and Neutrality Jointly Enable Evolvability When the Size of the Memory Repository Is Small

(a) Fitness ratio

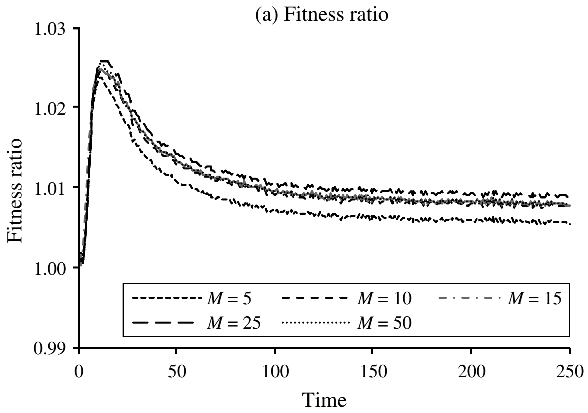

(b) NPV ratio

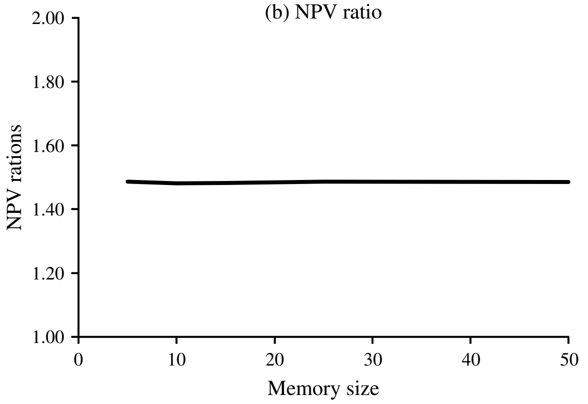

Notes. Forgetting is modeled taking into consideration finite memory sizes of 5, 10, 15, 25, and 50 positions (using neutrality $= 0.10$ and $k = 4$). We compute the evolvability statistics as a ratio of the variable values for a population with memory to another population without memory. As the size of memory increases, the fitness ratio (panel (a)) increases, but there are decreasing returns to memory size. Even a small memory size (five units) leads to improvements in performance over the memoryless case. Similarly, the NPV ratio is almost invariant with memory size. Thus, memory creates value, even when the capability is small.

We simulated forgetting and finite memory size as a robustness check using the evolvability metrics. These metrics are computed as the ratio of fitness (NPV) of a population of firms searching neutrally with memory of a given size to that of another population without memory. In this analysis, we limited the size of memory to 5, 10, 15, 25, and 50 positions only. We kept our other parameter values (neutrality, complementarities, costs, and discount rates) at their baseline rates.

Figure 7 reports the results of our investigations of the size of memory on evolvability. Panel (a) shows that the fitness ratio is greater than 1 and that it increases as memory size increases. However, increasing memory beyond 25 units provides very little additional benefit. Furthermore, the results indicate that even for low memory levels of five units there is a marked improvement in fitness relative to the case with no memory. To conclude, the NPV ratio (panel (b)) is considerably greater than 1 and almost invariant with memory size. These results indicate that although the effect of memory on fitness increases with memory size, retaining all information accumulated in organizations has little value, and even a small and partial memory provides substantive benefits. In other words, these results give reassurance that a lot of memory is wasteful and that a bit goes a long way.

Discussion: Robustness to Time and Costs

If evolution is guided by nature with infinite patience, the discount rate is essentially 0, and biological simulations have accepted this assumption. The tradition in NK modeling has stayed close to this practice. However, organizations care about time and the intertemporal trade-off between exploitation today and exploratory benefits tomorrow. We investigated further the mechanics of memory by investigating whether its effect on evolvability is robust to time (discount rates), costly change (search costs), and costs of retracing ones steps (costs of using memory).[14]

Our investigation of robustness to time and discount rates reveals a paradox: even though memory pertains to the acquisition, retention, and retrieval of *prior* experience, it creates more value in the *future*. We checked for robustness by varying discount rates from 0.43% to 25% per month (time period) and computed the NPV ratio for neutral search with memory (Q4) to that without memory (Q3). As discount rates increase, the benefits accrued from memory decrease. Since the present value of future payoffs decreases as the discount rate increases, this indicates that a substantial portion of value created by memory lies in the future.

The efficiency-inducing role of memory provides an additional implication: memory is particularly useful when costs of inefficiency are high. We varied search costs from 0 to 0.6 per move in search and computed the NPV ratio for neutral search with memory to that without it (see the online supplement). Our simulations reveal not only that our prior results remain robust over a broad range of search costs but also that the NPV ratio increases as search costs increase. That is, as a source of efficiency, memory is particularly useful when the costs of search are high.

Finally, memory results in more performance even when the costs to using memory are high. Reverting to a prior position retained in memory can be costly, but our prior analysis did not take this into consideration. For instance, if a firm introduces piece rates and assembly lines into its production processes and later realizes that they are incompatible, then reverting to prior practices would be beneficial. However, implementing prior practice is not is not easy and comes at some cost—a cost not taken into consideration in analysis to this point. To this end, we introduced a cost of using memory (see the online supplement) to check whether the contribution of memory to evolvability remains robust. Our results indicate that even though the NPV ratio decreases as costs of using memory increase, it remains greater than 1, and memory still results in greater performance.

Conclusions

The objective of this paper was to resolve the ambiguous assessment of an organization's ability to search and innovate by showing that neutrality and memory improve organizational evolvability. Their contribution is that they harness the evolutionary mechanism of plasticity without suffering oversetting costs arising out of additional search. By offering an escape from competency traps, neutrality and memory dynamically promote category change and improved fitness.

There has been, as noted above, long-standing pessimism regarding the abilities of organizations to change, and a fundamental ambiguity of the contribution of memory. Memory has an interesting property: without neutrality, it does not help very much. Memory leads to faster time to discovery in neutral search where the evaluation of alternatives is costly. With or without memory, adaptive search leads to local optima. These conditions lay the groundwork for a negative assessment of adaptive search and the role of memory on the evolvability of organizations. Neutrality improves innovation, but by itself, it incurs large tuition fees. Again, pessimism seems warranted.

Yet the history of organizational capabilities shows a massive and rapid improvement in productivity. Clearly, something is amiss in these adaptionist arguments. We have shown by simulations that learning stored in memory and random and neutral change within a category of activities uniquely generate evolvability in the presence of switching and temporal costs. The simulations used to evaluate this argument incorporated realistic considerations of costs and forgetting rates that are empirically grounded. Both neutrality and memory are *necessary* to expand the evolvability of organizations; that is, they are complements.

The resolution of the ambiguity of memory is easy to summarize once the net effects are parsed out: memory is good but takes the blame because of the bad adaptive company its keeps. Here, we can understand the contribution of neutrality and memory. It is in the evolutionary advantage of organizations to evolve capabilities that, though limited, can avoid dead end traps. Two such capabilities are neutral search and memory. Neutrality permits an organization to find ridges from its current prototype to improve prototypical ways of doing things, such as moving from craft to mass to Toyota production models. Memory enables firms to draw from prior experience and start new search trajectories when faced with dead ends in search. If organizations permit small and local experiments enlightened by history, they increase probabilistically their access to new and better prototypes and ways of doing things. Coupled with moderate chance caught on the wing, organizations with a bit of ability to remember their past are endowed with a capacity of evolvability and a higher potential for innovation.

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/orsc.2013.0841.

Acknowledgments

The authors thank INSEAD, the National University of Singapore [Grants R-535-000-005-133 and R-313-000-090-133], and the Sanford C. Bernstein & Co. Center for Leadership and Ethics at Columbia University for the funding for this research. They are grateful to senior editor Linda Argote and two anonymous reviewers for thoughtful and constructive feedback.

Endnotes

[1] The NK model was introduced into molecular biology by Kauffman (1993) and Fontana et al. (1993), among others. It has been a workhorse in analyzing adaptation by organizations since Levinthal (1997), who shows that adaptive search leads to competency traps.

[2] A more traditional treatment of intertemporal trade-offs is to subject organizations to selection and then look at the population frequencies at the end of the simulation. This metric for comparing adaptive and neutral evolution is less meaningful for neutrality, because many genetic sequences map into the same fitness value.

[3] The NK simulation has been used to study organizational design (Levinthal and Warglien 1999), cognition (Gavetti and Levinthal 2000), modularity (Ethiraj and Levinthal 2004b), underlying topology or social network (Rivkin and Siggelkow 2003), and how communication networks affect performance (Lazer and Friedman 2007).

[4] The following history draws on the study by Hounshell (1985). See Kogut et al. (2004) for an analysis using fuzzy set reduction to find the organizational prototypes in the auto industry.

[5] See, for example, Athey and Stern (2004) for an econometric example. Ragin's qualitative comparative analysis relies similarly on coding organizations (e.g., countries) as binary strings (Ragin 1987); March's (1991) article on exploitation and exploration codes organizational culture and individual knowledge as binary strings.

[6] For a recent review of the concept, see Masel and Trotter (2010). Ethiraj and Levinthal (2004a) studied "evolvability" in the context of how the structure of complexity affects the usefulness of bounded rational design efforts.

[7] Fontana (2003) explained this point by differentiating between dynamics and kinematics, with the former emphasizing why things move and the latter the structural constraints.

[8] See Kogut and Zander (1992) for a treatment of knowledge as information and know-how.

[9] Bonacich and Lu (2012, p. 18) provided an accessible description.

[10] We also checked for robustness by increasing the dimensionality to 20, or 1.05 million possible sequences, with no substantive change in our results.

[11] For formal results on this point, see Weinberger (1991).

[12] Our simulation treats each time period as a month. A discount rate of 0.43% per time period (month) corresponds to an annual discount rate of 5%. We consistently use this discount rate in all of our analyses.

[13] Using an annual discount rate of 5% (0.43% per simulation time period) and search costs of 0.5 per move, the NPV for each of the four experimental conditions can be computed as follows:

$$\begin{aligned}
\text{Q1} &= (0.47-0.50)+(0.59/0.0043)\times 0.9957=136.6,\\
\text{Q2} &= (0.47-0.50)+(0.59-0.50)\times 0.9957+(0.48/0.0043)\\
&\quad \times(0.9957)^2=110.7,\\
\text{Q3} &= (0.47-0.50)+(0.59/0.0043)\times 0.9957=136.6, \quad \text{and}\\
\text{Q4} &= (0.47-0.50)+(0.59-0.50)\times 0.9957+(0.48-0.50)\\
&\quad \times(0.9957)^2+(0.55-0.50)\times(0.9957)^3+(0.42-0.50)\\
&\quad \times(0.9957)^4+(0.79/0.0043)\times(0.9957)^5=179.8.
\end{aligned}$$

[14] The online supplement contains a detailed treatment of robustness to discount rates, search costs, and costs of using memory.

References

Argote L (2013) *Organizational Learning: Creating, Retaining and Transferring Knowledge*, 2nd ed. (Kluwer Academic Publishers, Norwell, MA).

Argote L, Beckman SL, Epple D (1990) The persistence and transfer of learning in industrial settings. *Management Sci.* 36(2):140–154.

Athey S, Stern S (2004) An empirical framework for testing theories about complementarity in organizational design. NBER Working Paper 6600, National Bureau of Economic Research, Cambridge, MA.

Benkard CL (2000) Learning and forgetting: The dynamics of aircraft production. *Amer. Econom. Rev.* 90(4):1034–1054.

Bonacich P, Lu P (2012) *Introduction to Mathematical Sociology* (Princeton University Press, Princeton, NJ).

Darr ED, Argote L, Epple D (1995) The acquisition, transfer, and depreciation of knowledge in service organizations: Productivity in franchises. *Management Sci.* 41(11):1750–1762.

Dutton JM, Thomas A (1984) Treating progress functions as a managerial opportunity. *Acad. Management Rev.* 9(2):235–247.

Ethiraj SK, Levinthal D (2004a) Bounded rationality and the search for organizational architecture: An evolutionary perspective on the design of organizations and their evolvability. *Admin. Sci. Quart.* 49(3):404–437.

Ethiraj SK, Levinthal D (2004b) Modularity and innovation in complex systems. *Management Sci.* 50(2):159–173.

Fontana W (2003) The topology of the possible. Working paper, Santa Fe Institute, Santa Fe, NM.

Fontana W, Stadler PF, Bornberg-Bauer EG, Griesmacher T, Hofacker IL, Tacker M, Tarazona P, Weinberger ED, Schuster P (1993) RNA folding and combinatory landscapes. *Physical Rev. E* 47:2088–2099.

Gavetti G, Levinthal D (2000) Looking forward and looking backward: Cognitive and experiential search. *Admin. Sci. Quart.* 45(1):113–137.

Hargadon A, Sutton RI (1997) Technology brokering and innovation in a product development firm. *Admin. Sci. Quart.* 42(4):716–749.

Hedberg BLT, Nystrom PC, Starbuck WH (1976) Camping on seesaws: Prescriptions for a self-designing organization. *Admin. Sci. Quart.* 21(1):41–65.

Hounshell D (1985) *From the American System to Mass Production, 1800–1932: The Development of Manufacturing Technology in the United States* (John Hopkins University Press, Baltimore).

Ichniowski C, Kathryn S, Giovanna P (1997) The effects of human resource management practices on productivity: A study of steel finishing lines. *Amer. Econom. Rev.* 87(3):291–313.

Jain A (2012) Forgetting in R&D activities: The depreciation of knowledge stocks in biotechnology research. *Management of Innovation Tech. (ICMIT), 2012 IEEE Internat. Conf., Bali, Indonesia*, 68–73.

Jain A (2013) Learning by doing and the locus of innovative capability in biotechnology research. *Organ. Sci.* 24(6):1683–1700.

Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47(2):263–291.

Kauffman SA (1993) *The Origins of Order: Self-Organization and Selection in Evolution* (Oxford University Press, New York).

Kimura M (1968) Evolutionary rate at the molecular level. *Nature* 217(5129):624–626.

Kimura M (1983) *The Neutral Theory of Molecular Evolution* (Cambridge University Press, Cambridge, UK).

King JL, Jukes TH (1969) Non-Darwinian evolution. *Science* 164(3881):788–798.

Kogut B, Zander U (1992) Knowledge of the firm, combinative capabilities, and the replication of technology. *Organ. Sci.* 3(3): 383–397.

Kogut B, MacDuffie JP, Ragin C (2004) Prototypes and strategy: Assigning causal credit using fuzzy sets. *Eur. Management Rev.* 1(2):114–131.

Lazer D, Friedman A (2007) The network structure of exploration and exploitation. *Admin. Sci. Quart.* 52(4):667–694.

Levine SS, Prietula MJ (2012) How knowledge transfer impacts performance: A multilevel model of benefits and liabilities. *Organ. Sci.* 23(6):1748–1766.

Levinthal DA (1997) Adaptation on rugged landscapes. *Management Sci.* 43(7):934–950.

Levinthal DA, Warglien M (1999) Landscape design: Designing for local action in complex worlds. *Organ. Sci.* 10(3):342–357.

Liang DW, Moreland R, Argote L (1995) Group versus individual training and group performance: The mediating factor of transactive memory. *Personality Soc. Psych. Bull.* 21(4):384–393.

Lobo J, Miller JH, Fontana W (2004) Neutrality in technological landscapes. Working paper, Santa Fe Institute, Santa Fe, NM.

MacDuffie JP (1995) Human resource bundles and manufacturing performance: Organizational logic and flexible production systems in the world auto industry. *Indust. Labor Relations Rev.* 48(2):197–221.

March JG (1972) Model bias in social action. *Rev. Educational Res.* 42:413–429.

March JG (1991) Exploration and exploitation in organizational learning. *Organ. Sci.* 2(1):71–87.

Masel J, Trotter MV (2010) Robustness and evolvability. *Trends Genet.* 26(9):406–414.

Monod J (1971) *Chance and Necessity* (Alfred A. Knopf, New York).

Moreland RL, Argote L, Krishnan R (1996) Socially shared cognition at work: Transactive memory and group performance. Nye JL, Brower AM, eds. *What's Social About Social Cognition? Research on Socially Shared Cognition in Small Groups* (Sage, Thousand Oaks, CA), 57–84.

Nelson RR, Winter SG (1982) *An Evolutionary Theory of Economic Change* (Belknap Press, Cambridge, MA).

Neustadt RE, May ER (1986) *Thinking in Time: The Uses of History for Decision Makers* (Free Press, New York).

Nystrom PC, Starbuck WH (1984) To avoid organizational crises, unlearn. *Organ. Dynam.* 12(4):53–65.

Ragin C (1987) *The Comparative Method: Moving Beyond Qualitative and Quantitative Strategies* (University of California Press, Berkeley).

Reagans R, Argote L, Brooks D (2005) Individual experience and experience working together: Predicting learning rates from knowing who knows what and knowing how to work together. *Management Sci.* 51(6):869–881.

Rivkin JW, Siggelkow N (2003) Balancing search and stability: Interdependencies among elements of organizational design. *Management Sci.* 49(3):290–311.

Thompson P (2007) How much did the Liberty shipbuilders forget? *Management Sci.* 53(6):908–918.

Walsh JP, Fahey L (1986) The role of negotiated belief structures in strategy making. *J. Management* 12(3):325–338.

Walsh JP, Ungson GR (1991) Organizational memory. *Acad. Management Rev.* 16(1):57–91.

Wegner DM (1986) Transactive memory: A contemporary analysis of the group mind. Mullen B, Goethals GR, eds. *Theories of Group Behavior* (Springer-Verlag, New York), 185–205.

Weinberger ED (1991) Local properties of Kauffman's N-K model: A tunably rugged energy landscape. *Physical Rev. A* 44(10): 6399–6413.

Amit Jain is an assistant professor of strategy and technology management at the Faculty of Engineering and NUS Business School at the National University of Singapore. He received his Ph.D. from INSEAD at Fontainebleau and a doctorate from the University of Paris, Dauphine. His interests include organizational learning and change, with a particular emphasis on the role of learning by doing and forgetting.

Bruce Kogut is the Sanford C. Bernstein & Co. Professor at Columbia University and is on leave at the Wissenschaftskolleg in Berlin in 2012–2013. His 2012 edited book *The Small Worlds of Corporate Governance* provides data (available at the Bernstein Center website) and algorithms estimating and simulating governance networks. His current work focuses on the sources of inequality and on the creation and transmission of productivity, from the Weimar Republic to microfinance in India.