Oliver Baumann/Dirk Martignoni*

Evaluating the New: The Contingent Value of a Pro-Innovation Bias**

Abstract

It is a central tenet in the literature on organizational change that firms need to explore novel courses of action in order to adapt and survive. Should firms thus exhibit a "pro-innovation bias" when evaluating novel decision alternatives? Or should firms rather assess new opportunities as objectively as possible? Our analysis of a simulation model suggests that a pro-innovation bias can have exploration-enhancing effects that increase long-run performance in complex and stable environments, but can also decrease performance substantially if the bias becomes too pronounced. However, under most other conditions, an unbiased, objective evaluation of novel opportunities is most effective. We also identify a set of contingency factors that strongly affect the value of a pro-innovation bias, which may explain why it is that we see so few firms with such a bias.

JEL-Classification: C63, D21, D83, O31.

Keywords: Innovation; Organizational Decision Making; Organizational Exploration and Adaptation; Organizational Search.

1 Introduction

It is a central tenet in the literature on organizational change that firms need to foster exploration to ensure their long-term adaptation and survival (Nelson and Winter (1982); Eisenhardt and Tabrizi (1995); Teece, Pisano, and Shuen (1997); Benner and Tushman (2003)). Otherwise, exploration may be driven out by exploitation, and firms run the risk of ending up in a success, or competence, trap (March (1991); Levinthal and March

* Oliver Baumann, Department of Marketing & Management, University of Southern Denmark, Campusvej 55, 5230 Odense M, Denmark, e-mail: oliv@sam.sdu.dk. Dirk Martignoni, Institute for Strategy and Business Economics, University of Zurich, Plattenstr. 14, 8032 Zurich, Switzerland, e-mail: dirk.martignoni@business.uzh.ch.

** For helpful comments, the authors would like to thank J. P. Eggers, Hans-Georg Gemünden, the editors Dodo zu Knyphausen-Aufseß and Arnold Picot, two anonymous reviewers, and participants at the European Meeting on Applied Evolutionary Economics, the Israel Strategy Conference, the Organization Science Winter Conference, and the Workshop of the Organization Section of the German Academic Association of Business Research.

(1993)). Parts of this literature even seem to propagate the idea that exploration is desirable per se, and that embracing novel opportunities is generally superior to sticking with the status quo (as observed, e.g., by Sheth (1981); Van de Ven (1986); Sturdy and Grey (2003); Rogers (2005)). That is, they suggest that firms should seek to adopt a "pro-innovation bias" to ensure sufficient exploration.

However, reality paints a different picture. Rather than proactively embracing new ideas and opportunities, the structures, processes, and tools that firms apply to evaluate novel ideas often suppress exploration systematically by letting firms perceive new ideas as less attractive than they actually are (Christensen, Kaufman, and Shih (2008)). Moreover, organizational decision makers often exhibit a systematic "status-quo bias" (Kahneman, Slovic, and Tversky (1982); Samuelson and Zeckhauser (1988); Gilovich, Griffin, and Kahneman (2002)), preferring the existing state of affairs over new alternatives, unless there are extremely compelling incentives to change. Consequently, history is replete with firms that missed attractive opportunities by sticking to their current solutions.

This observation proposes a puzzle. If we assume that managers can affect how their organizations evaluate "the new" (Christensen, Kaufman, and Shih (2008)), then we must ask why firms often seem to exhibit a status quo bias when presented with the choice between the current state of affairs and the exploration of new opportunities. Given the above arguments, should firms not rather develop a pro-innovation bias, adjusting their evaluation policies and processes in a way that systematically promotes rather than suppresses exploration? Or would it be preferable to always evaluate novel alternatives as objectively as possible?

To address these questions, we build on the behavioral accounts of organizational decision making that envision firms as evolving through adaptive search, i.e., by generating and evaluating novel decision alternatives (Simon (1955); March and Simon (1958); Cyert and March (1963)). We develop a simulation model in which firms search for solutions to a set of interdependent choices that lead to performance (Levinthal (1997)). In doing so, firms differ in how they evaluate new opportunities (Knudsen and Levinthal (2007)), ranging from firms that possess a pro-innovation bias that makes it possible for them to overestimate the value of the new opportunities, to those with a status-quo bias that results in fresh opportunities being underestimated (Kahneman, Knetsch, and Thaler (1991))[1]. By controlling this evaluation bias we can systematically explore its impact on adaptation and performance. Although we are not the first to study the impact of imperfect decision making on adaptation (see, e.g., Denrell and March (2001); Knudsen and Levinthal (2007); Christensen and Knudsen (2010)), our study is the first to disentangle the effects of a systematic evaluation bias.

The analysis we perform in our model shows that contrary to what we might think, a slight pro-innovation bias can have exploration-enhancing effects that increase long-run

1 "Bounds" to or deviations from rationality can take a variety of different forms. In our study, we particularly focus on one deviation from rationality – systematic biases in the evaluation of new opportunities. Thus, in the context of our study, a "rational" evaluator would correctly anticipate the true value of an opportunity.

performance in complex and stable environments, but may also decrease performance substantially if the bias becomes too pronounced. However, under most other conditions, an unbiased, objective evaluation of novel opportunities is optimal.

Furthermore, we identify four contingency factors that strongly affect the value of having a pro-innovation bias, which may help explain why it is that we see so few firms with such a bias: (1) the complexity and turbulence of the task environment, (2) the breadth of organizational search, (3) the degree of managerial control over the implementation and execution of a firm's evaluation process, and (4) the selection pressures that firms face.

Our study is structured as follows: In Section 2 we review prior research. In Section 3 we describe our model, and in Section 4 we present the results of our simulation experiments. Section 5 discusses our findings and concludes.

2 Prior Research

2.1 Organizational Exploration and Adaptation

Innovation is one of the primary ways in which organizations adapt to their environment (Nelson and Winter (1982); Eisenhardt and Tabrizi (1995); Teece, Pisano, and Shuen (1997); Benner and Tushman (2003)). Without exploring new ideas, organizations are likely to prematurely embrace suboptimal solutions (March (1991)). Hence, how to sustain a sufficient level of exploration has been considered a primary organizational challenge (Holland (1975); March (1991); Levinthal and March (1993)). But it is also a well-established finding that strategies that overemphasize exploration are also suboptimal, since firms that follow these strategies tend to be too erratic and thus suffer the costs of experimentation without gaining many of its benefits (Kanter (1988); March (1991)). Thus it has become a basic tenet in organization theory and strategic management that firms must balance the exploration of new ideas with the exploitation of old certainties (Holland (1975); March (1991)).

Achieving this balance requires two key processes that shape organizational exploration and adaptation: an effective process for generating decision alternatives, and an effective process for selecting among these alternatives (March (1991)). The latter process is the focus of our study. Prior research along these lines has investigated how organizational structures differ in terms of their ability to prevent evaluation errors (Sah and Stiglitz (1996); Christensen and Knudsen (2010)), or has focused on the costs and benefits of errors (Knudsen and Levinthal (2007)) or on the lack of precision when firms evaluate new alternatives (Denrell and March (2001)). This body of research has shown that the potential benefits of imperfect evaluation processes stem from the fact that imprecision and uncertainty can induce additional exploration, which in turn can lead to a better balance between exploration and exploitation. This argument is consistent with March's (2006) assertion that organizations have to introduce "some elements of foolishness." But in doing so, this line of research focuses primarily on nonsystematic errors in the evaluation process. At the same time, we know from experimental economics that many deviations from objective decision making "are too widespread to be ignored, too systematic to be dismissed as random errors" (Kahneman and Tversky (2000, 210)). In a similar vein, Cyert and March (1963)) argue that organizations often suffer from predictable biases rather than random errors. In our model, we extend this line of inquiry by disentangling the effects of firms that have a systematic evaluation bias either in favor of "the new" or against it.

2.2 Expected Impact of a Systematic Evaluation Bias

Research on organizational change has repeatedly highlighted the difficulties of (incumbent) firms to adapt to technological change (Henderson and Clark (1990); Christensen (1997)). The notion that individuals and organizations fail to embrace new opportunities has been discussed in different streams of the literature under a variety of labels such as "anti-innovation bias" (Christensen, Kaufman, and Shih (2008)), "organizational inertia" (Hannan and Freeman (1984)), "status quo bias" (Samuelson and Zeckhauser (1988)), "reluctance to change" (McGrath (1997)), "program persistence bias" (Williamson (1975)), or the "not invented here syndrome" (Katz and Allen (1982)). Although diverse in their theoretical foundations, these contributions all concur in the observation that for many organizations, the well-known and relatively risk-free current strategies and practices often appear to be more attractive than exploring the unknown and potentially risky decision alternatives. Furthermore, these studies unanimously agree that such a bias for the status quo comes at a cost. For example, organizations with a status-quo bias may develop "core rigidities" (Leonard-Barton (1992)) or get caught in "competency traps" (Levitt and March (1988)). Although their focus on current strategies, technologies, knowledge, and capabilities may enable these firms to reap immediate profits, it will eventually foster stagnation and leave the firms vulnerable to market and technological changes. Anecdotally, the press routinely documents the failure of formerly prominent firms such as Polaroid, DEC, PanAm, RCA, Sears, and Bethlehem Steel (Sull (2005)), whose decline is often attributed to the development of an exploration-suppressing culture that made them inert and unable to change. Further, a long-standing debate is concerned with the question of whether organizations can adapt at all. The most pessimistic perspective argues that organizations are largely inert and thus ultimately fail (Hannan and Freeman (1984); Hannan and Carroll (1992)). This body of research suggests:

Proposition 1: A status-quo bias has negative effects on the performance of organizational exploration and adaptation.

The idea of a bias in favor of the new, i.e., of organizations that prefer novel decision alternatives over their status quo, thereby overestimating the value of the new ("new means better, old is worse"), has rarely been discussed explicitly in the organizational literature. Still, there is some evidence that such decisions can be found in various organizational contexts. For example, there are numerous firms that cut their costs by outsourcing or offshoring parts of their operations. Many of these firms often grossly overestimate the savings that outsourcing can deliver, so the literature abounds with discussions of organizations that reintegrated functions that had been outsourced (Tadelis (2007)). In contrast, when it comes to the performance implications of a pro-innovation bias, prior work is less consistent. Most research points to a positive effect, arguing that firms in innovation-driven environments can only achieve high performance, and, ultimately, survive, if they manage to persistently adapt and change (D'Aveni (1994); Brown and Eisenhardt (1997)). Similarly, the literature on search and adaptation suggests that a pro-innovation bias may inflate an organization's expectations on the value of adopting a novel alternative, i.e., of deviating from the status quo. Inflated expectations can, in turn, encourage exploration and may lead to a better balance between exploration and exploitation (Sutton and Barto (1998)).

In an entrepreneurial context, being overly optimistic may be an important antecedent to entrepreneurial activity, i.e., to the exploratory efforts that give rise to new organizations, and ultimately, to new products and services (Lowe and Ziedonis (2006)). For example, when comparing their current status as employees with the prospects of becoming self-employed, (eventual) entrepreneurs often greatly overestimate the potential returns from their intended entrepreneurial activities (e.g., Hayward, Shepherd, and Griffin (2006)). Without this optimism, many potential entrepreneurs might refrain from their explorative activities, some of which do succeed.

However, several studies have taken a somewhat different position. These studies criticize research on innovation as containing a "pro-innovation bias" (Rogers (1962); Rogers and Shoemaker (1971); Van de Ven (1986)), i.e., to assume or conclude that innovation is necessarily positive. Instead, these studies have argued that because firms tend to overestimate new decision alternatives, they often pursue ideas or adopt practices that turn out to be not only inefficient but also often hamper performance (Abrahamson (1991)). If organizations overemphasize the pursuit of new alternatives and ignore their own capabilities, they can end up in an exploration (or failure) trap (e.g., Levinthal and March (1993)). Put differently, they run the risk of seeking new opportunities at the expense of today's operations (Birkinshaw and Gibson (2004)), which can turn them to frenzies of experimentation and change by a dynamics of failure, i.e., failure that leads to search and change, which leads to failure, which leads to change and so on.

Hence, while some contributions suggest that a pro-innovation bias has dysfunctional performance effects, others are less pessimistic and emphasize the functional effects of such a bias. Our review of the above literatures points toward two opposing propositions:

Proposition 2a: A pro-innovation bias has negative effects on the performance of organizational exploration and adaptation.

Proposition 2b: A pro-innovation bias has positive effects on the performance of organizational exploration and adaptation.

3 Model

To accomplish our goal of studying the impact of an evaluation bias on the dynamics of organizational search and adaptation, we develop an agent-based simulation model (Macy and Willer (2002)). Decision-making agents are confronted with controlled environments. They are equipped with heuristics to react to their environment, and the resulting performance is recorded over time. By varying the behavior of the agents and the structure of the environment, we systematically explore the impact and interdependence of the variables under consideration. Although this approach grants high degrees of freedom to the modeler, we follow an established tradition to develop simple yet insightful (and traceable) models (March and Simon (1958); Cohen, March, and Olsen (1972); Nelson and Winter (1982); Burton and Obel (1995)).

Computational models have gained broad popularity in studies of organizational adaptation and learning (March (1991); Levinthal (1997); Gavetti and Levinthal (2000); Davis, Eisenhardt, and Bingham (2009)), particularly because such models allow a more rigorous analysis than does verbal discussion, forcing the modeler to make all underlying assumptions explicit. In contrast to algebraic approaches, computational models make it possible to incorporate a more comprehensive set of features into the analysis. Although they cannot yield "exact solutions" like closed-form modeling, computational models allow the researcher to model conditions under which closed-form approaches would be intractable. But most importantly, we are concerned with the question of how the search of organizational decision makers that have only bounded rationality is affected by an evaluation bias. Exploring the underlying dynamics of search can be easily achieved with computational models, whereas analytic models tend to be concerned with equilibria and not with the question of how, or whether, those equilibria will be attained. Hence, "[s]imulation is particularly useful when the theoretical focus is longitudinal, nonlinear, or processual, or when empirical data are challenging to obtain" (Davis, Eisenhardt, and Bingham (2007, 481)) – all of which applies to the context of our question.

Our model contains three components: a task environment that firms explore and to which they adapt, a process by which firms generate new decision alternatives, and a (biased) process of evaluating new alternatives.

3.1 Task Environment

We conceptualize firms as facing a set of interdependent decisions that determine firm performance (Porter (1996); Levinthal (1997); Siggelkow (2002)). For a firm to explore and adapt to its environment, its managers must make many decisions. For example, the manager of a manufacturing firm might have to decide about making her firm's production system more flexible, or whether to expand the firm's product variety. Furthermore, many of these decisions interact with each other. The value of a flexible manufacturing system, for instance, increases as a firm increases its product variety.

In the model, each firm must resolve N decisions $a_1,\ a_2,\ \ldots,\ a_N$. Without loss of generality, we assume that each decision is binary. For instance, a_1 might denote the decision to increase product variety $(a_1 = 1)$ or not $(a_1 = 0)$. Thus, a firm faces 2^N possible configurations of choices, each of which can be represented by a binary vector $\boldsymbol{a} = (a_1,\ a_2,\ \ldots,\ a_N)$.

In computational studies of firms as complex adaptive systems, it is common to interpret the payoffs to configurations of interdependent choices as performance landscapes (Levinthal (1997); Rivkin (2000)). A performance landscape consists of N "horizontal" dimensions, which are the N decisions that the firm needs to make; and one "vertical" dimension, which denotes the corresponding performance of each configuration. Thus, a performance landscape represents a mapping of each configuration $\boldsymbol{a}$ (each "point" on the landscape) to a performance value $V(\boldsymbol{a})$ (the "height" of the particular point).

We create performance landscapes with a variant of the NK model (Kauffman (1993); Kauffman (1995)) – stochastically, yet in a well-controlled manner. Although the NK model was created in evolutionary biology, it has been applied to a number of organizational issues (e.g., Levinthal (1997); Rivkin (2000); Baumann (2010)). In the NK model, each decision a_i is assumed to make a contribution c_i to the performance $V(\boldsymbol{a})$ that a firm receives from a particular configuration of choices $\boldsymbol{a}$. The contribution c_i of each decision a_i not only depends on how a_i is resolved (0 or 1), but also on how K other decisions $(\boldsymbol{a_{-i}})$ that interact with a_i are resolved. Hence, K controls the degree of interdependence between the decisions. When $K = 0$, all decisions are independent, and the performance contribution of each decision depends only on how the decision itself is resolved. In this case, the performance landscape is smooth and contains only a single peak. In contrast, if $K = N - 1$, the value of each decision depends on how all other decisions are resolved. The landscape is now rugged, exhibiting numerous local peaks. The identity of the K decisions $\boldsymbol{a_{-i}}$ that influence the value of each decision a_i is determined randomly for each decision variable. Particular values for all possible c_i's are determined randomly by drawing from a uniform distribution over the unit interval, i.e., $c_i(a_{i;}\ \boldsymbol{a_{-i}}) \sim u[0;1]$. We calculate the value $V(\boldsymbol{a})$ of a given set of choices $\boldsymbol{a}$ as the average of its N performance contributions, i.e., $V(\boldsymbol{a}) = [c_1(a_{1;}\ \boldsymbol{a_{-1}}) + c_2(a_{2;}\ \boldsymbol{a_{-2}}) + \ldots + c_N(a_{N;}\ \boldsymbol{a_{-N}})] / N$.

Hence, the landscape metaphor allows an intuitive representation of organizational search and adaptation: subject to its configuration of choices $\boldsymbol{a}$, a firm inhabits a particular point on the performance landscape. The firm searches for improvements to its current situation by identifying and evaluating alternative configurations, i.e., it tries to move uphill and reach high points on the performance landscape, configurations of choices that create high performance.

Although the landscape is supposed to remain stable throughout this process, we also consider the notion of turbulent environments. Here, the landscape undergoes "correlated" shocks in periodic intervals. In particular, once we create a landscape, every ten periods we replace each contribution value c_i by $0.2^*c_i + 0.8^*u$, where u is a new draw from a uniform distribution over the unit interval.

3.2 Generation of Alternatives

In each period, each firm considers one alternative $\tilde{\boldsymbol{a}}$ that differs in one decision from its status-quo set of choices $\boldsymbol{a}$. Thus, if the firm is currently at 1000 (assuming $N = 4$), it would have four alternatives available: 1001, 1010, 1100, and 0000. For instance, a product manager might consider different ways of modifying her firm's product portfolio. Among the N possible alternatives, the manager picks one at random. This procedure for generating alternatives that are similar to the current configuration of choices represents a strategy of local search, which is a central feature in both theoretical (March and Simon (1958); Cyert and March (1963); Nelson and Winter (1982)) and empirical accounts (Stuart and Podolny (1996); Rosenkopf and Almeida (2003)) of organizational decision making and adaptation. According to this body of work, cognitive bounds prevent managers from coming up with radically innovative ideas, i.e., with alternative configurations that differ in many dimension from the status quo. (In Section 4.3, we relax this assumption and explore how our findings are affected if the bounds on managers' rationality are less severe, i.e., if they can tweak multiple dimensions of their current alternative to come up with a different configuration of choices.)

3.3 Evaluation of Alternatives

Subsequently, the firm needs to evaluate the newly identified alternative $\tilde{\boldsymbol{a}}$. Ideally, this process would proceed as follows: if the firm finds that the alternative denotes a performance improvement, i.e., if $V(\tilde{\boldsymbol{a}}) > V(\boldsymbol{a})$, then it will adopt the alternative and move from point $\boldsymbol{a}$ to the nearby point $\tilde{\boldsymbol{a}}$ on the landscape. But if the firm finds that the value of $\tilde{\boldsymbol{a}}$ is lower than or equal to the value of the firm's current alternative ($V(\tilde{\boldsymbol{a}}) \leq V(\boldsymbol{a})$), then the firm will discard the alternative and remain on its current "spot" on the landscape, generating another local alternative in the next period.

However, we assume that this evaluation process is affected by the firm's evaluation bias α. We assume that rather than perceiving the actual value $V(\tilde{\boldsymbol{a}})$ of a new alternative, a firm is biased for or against fresh ideas, which translates into the firm perceiving a biased value $V(\tilde{\boldsymbol{a}}) + \alpha$. Hence, if $\alpha < 0$, then the firm will systematically underestimate the true value of any new alternative. Doing so would imply that the firm's status quo is more attractive to it than any alternative that is objectively superior by an amount of up to α. If, in contrast, $\alpha > 0$, then the firm will systematically overestimate new alternatives and is therefore considered to posses a pro-innovation bias. Thus, even alternatives that are objectively inferior to the status-quo set of choices will, up to a difference of α, appear preferrable to the firm.

We assume that despite having a biased perception of potential decision alternatives, the firm learns about the true (i.e., objective) value of a new alternative after it adopts it[2]. This modeling assumption reflects the fact that cognitive, "offline" evaluations of new

2 In our model, since we are primarily concerned with the implications of a bias in evaluating new alternatives, this updating process occurs instantaneously.

alternatives are often more crude and less reliable than are experiential, "online" learning processes (Gavetti and Levinthal (2000)).

Once a firm has implemented a configuration and has learned that it cannot be further improved by any local alternative, the search ends. In this case, the firm has reached a local peak on the landscape. (A local peak is a configuration of choices $\boldsymbol{a}$ with $V(\boldsymbol{a}) > V(\tilde{\boldsymbol{a}})$ for all $\tilde{\boldsymbol{a}}$ that differ from $\boldsymbol{a}$ in one decision.) Because interdependencies between the different decisions result in rugged landscapes (Kauffman (1993); Levinthal (1997)), there are many internally consistent configurations, i.e., local peaks, that may act as "competency traps" (Levinthal and March (1981); Levitt and March (1988)) and terminate a firm's adaptive search.

4 Results

To study how the biased evaluation of novel decision alternatives affects exploration and adaptation, we place firms that differ in their evaluation bias α onto random points of our stochastically generated performance landscapes. We then let the firms search for 500 periods, by which time either they all reach a local peak or converge to an ultimate performance level. We measure the performance of each firm relative to the global peak in each landscape, i.e., firm performance is 1.0 if the firm reaches the global peak. To ensure that performance differences are characteristic of the model and that they do not result from any stochastic effects, we repeat each experiment for 5,000 different landscapes and calculate the average performance for each type of firm across all landscapes.

4.1 Biased Evaluation That Boosts or Restricts Exploration

In our first experiment, we study the general implications of an evaluation bias on a firm's performance in the short, medium, and long run (*Figure 1*). To do so, we assume a stable environment of medium complexity ($K = 5$). When we apply a short-term perspective, we find an inverted u-shaped relation that is roughly symmetric around $\alpha = 0$. Hence, because both biases yield similar performance penalties, firms that evaluate new alternatives in an objective manner outperform those that over- or underestimate new options. Precise evaluators make only correct decisions rather than running the risk of missing some (objectively) good alternatives (as they would if they had a status-quo bias) or of adopting some (objectively) inferior alternatives (as they would if they had a pro-innovation bias), which induces them to adapt efficiently and achieve performance improvements quickly.

Figure 1: Baseline result

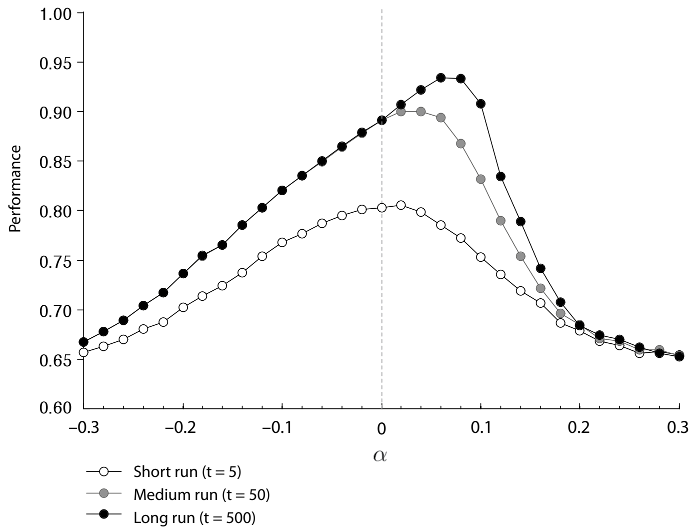

This figure reports the average firm performance in the short, medium, and long run over 5,000 stable landscapes with $N = 8$ and $K = 5$. Firms differ in their evaluation bias (α).

In contrast, in the medium to long run, the relation between the (biased) evaluation of new alternatives and firm performance becomes nontrivial. Although the relation remains an inverted u-shape, it is now asymmetric and skewed to the right. Hence, an objective evaluation no longer represents the optimum. Instead, a slight pro-innovation bias yields a long-term performance that is significantly higher than if a firm were a perfect evaluator. But if the bias becomes too great, then the resulting performance drops are severe.

To examine the effects that drive these results, we consider the individual performance development of two representative firms with different degrees of α. A firm that is evaluating alternatives in a fully objective manner ($\alpha = 0$) either makes a performance-improving decision or remains at its current performance level. Since this firm can perfectly discriminate between better and worse alternatives, it will only change if an alternative yields a real performance advantage. In contrast, a firm with a slight pro-innovation bias sometimes overestimates decision alternatives that in fact have a slightly lower performance advantage than does the status-quo set of choices. Although they set out to make a performance-improving change, these firms experience a slight performance drop after adopting the new alternative. This behavior implies that the firm will not always climb the nearest local peak on the performance landscape. Instead, by overestimating and implementing some slightly worse alternatives, it will go downhill temporarily, but by doing so it increases its chances of not getting stuck on lower local peaks and thus can move towards the foothills of higher peaks that it could not reach without these seeming mistakes. In contrast, a perfect evaluator will never make a downward movement but will instead climb the next-best local peak.

Thus, we ask why it is that *Figure 1* shows that only positive α values in a certain range increase long-run performance, but higher levels of α lead to steep performance drops. This result is driven by how a biased evaluation affects a firm's search behavior. In *Figure 2*, we report four measures that shed light on the underlying dynamics.

Figure 2: Effects of the evaluation bias on the dynamics of search

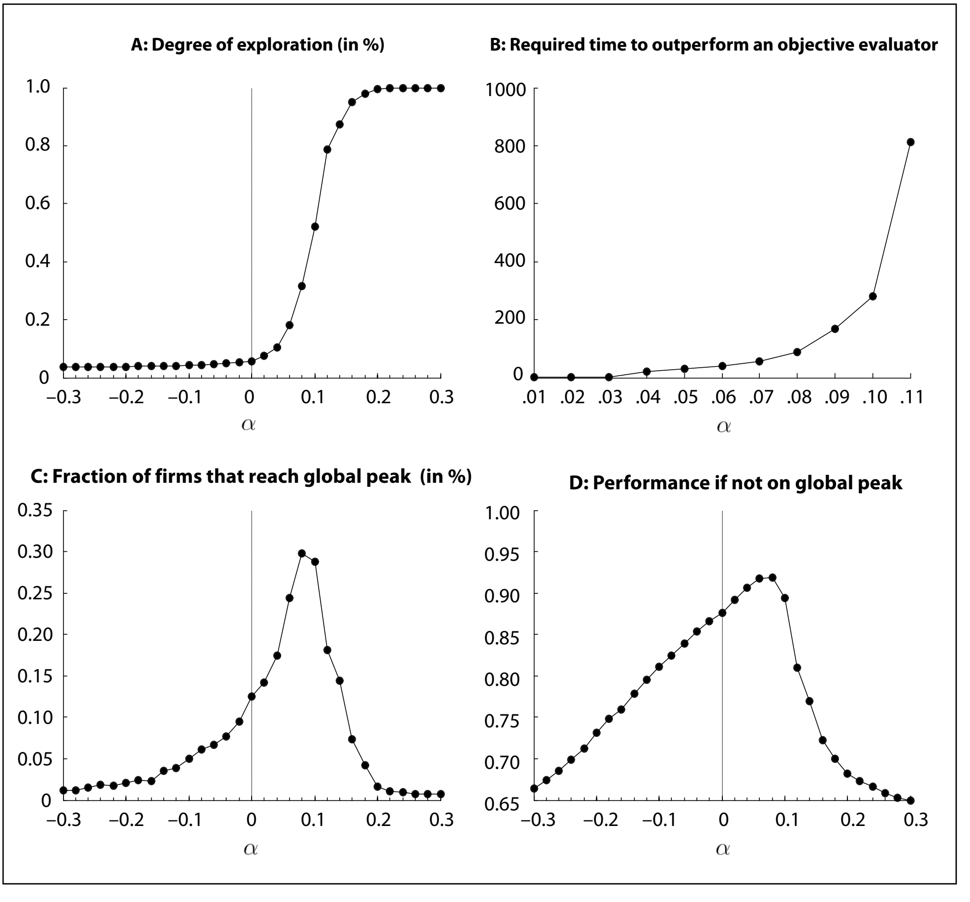

Panel A reports the fraction of all different performance contributions (c_i) that a firm evaluates during its search. Panel B reports the time required by a firm with a pro-innovation bias to outperform a firm that evaluates alternatives fully objectively. Panel C reports the percentage of firms that eventually end up on the global peak. Panel D reports the performance of those firms that do not reach the global peak. Firms differ in their evaluation bias (α). All values are averages in period 500 over 5,000 stable landscapes with $N = 8$ and $K = 5$.

As argued above, a slight pro-innovation bias may let a firm move up and down the landscape rather than only upward. As Panel A shows, this behavior results in broader search, i.e., in a higher number of alternatives that are screened by the firm. In consequence, Panel C shows that the chances increase that a firm identifies the global peak rather than only an average local peak. In addition, in Panel D, even if the firm does not reach the global peak, it will, on average, still exhibit a higher performance. However, if a firm's pro-innovation bias is too large, and thus overly exploration-promoting, e.g., if $\alpha > 0.1$, the firm will implement too many bad alternatives; it will have difficulty improving its performance, instead moving aimlessly across the landscape. A slight pro-innovation may help a firm reach an above-average peak by disrupting the firm's hill-climbing efforts and making it move downward at times, but if the exploration-promoting effect becomes too large, then it is extremely difficult for the firm to keep track of the "right" direction. Because this firm overestimates a large fraction of the potential alternatives it encounters, it explores very broadly (Panel A), but it cannot reap the benefits of this exploration (Panels C and D). In other words, although the firm implements good alternatives, it has difficulties sticking to them, because its biased evaluation makes various other alternatives so appealing that the firm is again drawn away from its better choices. Nevertheless, Panel B shows that the additional time to outperform a fully objective evaluator significantly increases as a firm's pro-innovation becomes larger. In sum, *Figure 2* illustrates how the power and pitfalls of a pro-innovation bias can either boost exploration to a healthy degree or result in an exploration trap.

4.2 Effects of Environmental Complexity and Turbulence

Our baseline case had assumed a complex, stable environment. But now we ask how environmental complexity and turbulence moderate the impact of the evaluation bias on firm performance. *Figure 3* compares the long-run performance in both simple and complex environments (white and black dots, respectively) that are either stable (solid lines) or turbulent (dotted lines).

This figure reports the average firm performance in the long run (period 500) over 5,000 landscapes with $N = 8$ and different levels of task interdependence (K). Environments are either stable or turbulent. In turbulent environments, a "correlated" shock occurs every 10 periods, and each contribution value c_i is replaced by $0.2^{*}c_i + 0.8^{*}u$, where u is a new draw from a uniform distribution over the unit interval. Firms differ in their evaluation bias (α).

Figure 3 shows an inverted u-shaped relation in all cases. In simple, stable environments $(K = 0)$, and in contrast to our baseline result in Section 4.1, objective evaluation yields the highest performance. The reason perfect evaluators now outperform firms with a slight pro-innovation bias is that simple environments result in smooth, rather than rugged, performance landscapes. Given such environments, even precise evaluators that engage in simple hill-climbing will reach the global peak. Firms with a pro-innovation bias also get close to the global peak, but because they overestimate alternatives with a similar, but inferior, performance, they do not always stick with the global peak once they have reached it, thus resulting in a slightly lower performance. Furthermore, in simple environments, turbulence does not affect the optimal approach, the unbiased evaluation. However, performance suffers because the environment changes before the firm can sufficiently explore it.

Figure 3: Effects of environmental complexity and turbulence

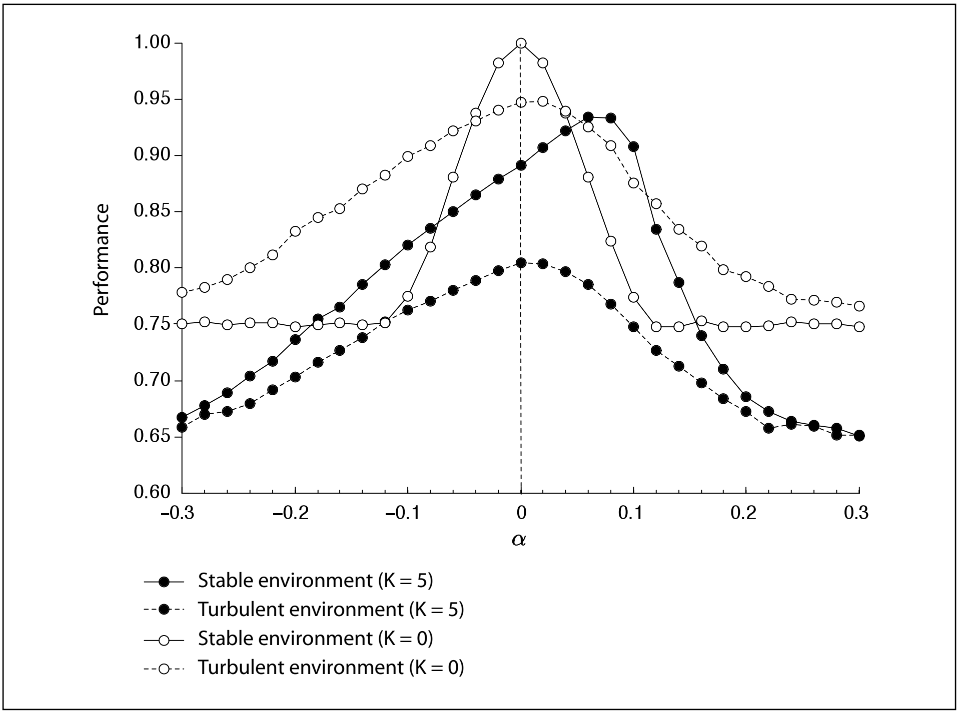

The case of a complex, stable environment repeats the experimental setting in Section 4.1. However, if we introduce turbulence into complex environments, we find that firms that evaluate new alternatives in a fully objective manner ($\alpha = 0$) achieve the highest performance. The reason for this difference is that similar to a short-run perspective in a stable environment, the ongoing shifts in the environment are best addressed by increasing performance in the most efficient manner. Thus, a pro-innovation bias leads to lower performance, because the environment shifts before the benefits of the higher levels of exploration that the bias entails can be reaped.

4.2 Effects of the Breadth of Search

In *Figure 4*, we explore the interdependence between a firm's processes of generating alternatives and of evaluating them. In the analyses above, firms generated alternatives by modifying only one dimension of the current alternative, what we called the "local search". Now, we assume that when generating new decision alternatives, organizations can search more broadly and modify more than one dimension simultaneously. (We note that the variable "search radius" denotes the number of the N dimensions that the organization might experiment with simultaneously.) In doing so, we again assume a complex and stable environment, the setting in which we should observe a beneficial effect of a pro-innovation bias.

Figure 4: Interdependence of alternative generation and alternative evaluation

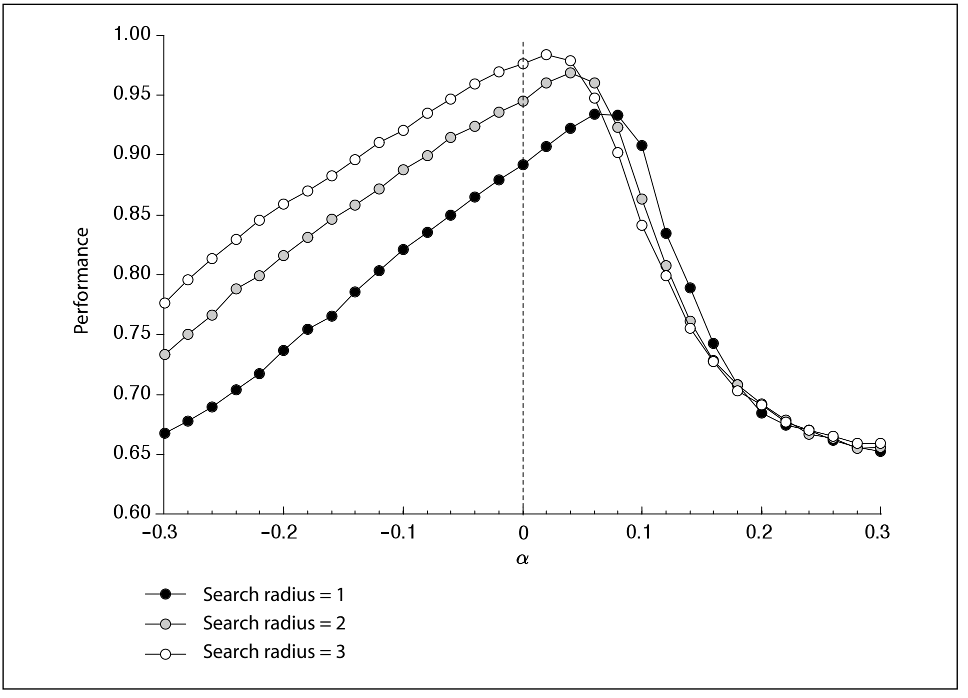

This figure reports the average firm performance in the long run (period 500) over 5,000 stable landscapes with $N = 8$, $K = 5$. Firms differ in their evaluation bias (α) and in their search radius.

Our general finding remains robust. A slight pro-innovation bias is optimal even if the firm can come up with more radically different alternatives. However, *Figure 4* also indicates a moderation effect: the optimal level of α is decreasing with the breadth of search. This result is due to the fact that although a slight pro-innovation bias can induce a healthy degree of additional exploration, the rate of exploration is already higher than in our benchmark case, and firms need to make more precise evaluations if they are to avoid becoming too explorative. Clearly, the more broadly a firm can search, the more objectively it should evaluate the alternatives that its search turns up.

4.3 Effects of Managerial Control over the Evaluation Process

In the analyses above, we assumed that a firm's management has no problems in implementing and maintaining the structures, processes, and tools that are associated with a particular evaluation process. We now relax this assumption. First, we assume that managers may have difficulties implementing the desired evaluation process. Specifically, we assume that there is a random, normally distributed deviation ε between the intended bias and the bias that is actually implemented. Hence, $\varepsilon = 0$ characterizes our baseline case. A deviation occurs for $\varepsilon \neq 0$, so firms that intend to implement α will instead implement $\alpha + \varepsilon$.

Figure 5: Implications of an imprecise implementation of the evaluation bias

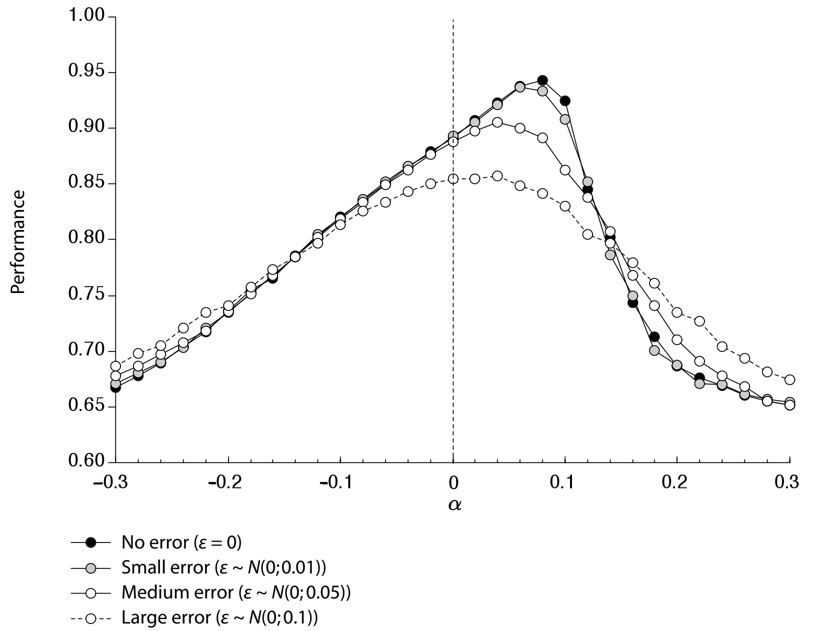

This figure reports the average firm performance in the long run (period 500) over 5,000 stable landscapes with $N = 8$, $K = 5$. Firms differ in their imprecision (ε) when implementing a desired evaluation bias *(α)* in the first period, i.e., they implement $\alpha + \varepsilon$ instead of α.

Figure 5 shows that not being able to perfectly implement the intended evaluation bias comes at a cost, especially if the intended bias is close to the optimum. For very poorly performing evaluation biases, deviations are either costless (for $\alpha < -0.1$) or may even create some value (for $\alpha > 0.1$). But most importantly, what denotes the optimal bias depends on management's ability to implement the intended bias. For highly capable managers or managers who exert strong control over the implementation, the optimal evaluation bias in environments of moderate complexity has a positive alpha value. In contrast, when $\varepsilon \neq 0$, i.e., when managers are less capable or exert less control over the implementation, the optimal evaluation bias gradually moves towards $\alpha = 0$ as ε gets larger.

However, management is probably responsible not only for controlling the implementation of the evaluation process, but also its execution. In *Figure 6*, we explore the effect of a lack of managerial control in the execution process, i.e., when firms have difficulties maintaining a chosen evaluation bias over time. We note that technically, this perspective means that ε is now redrawn from a normal distribution in each period. When there is imperfect implementation, ε is only drawn once in $t = 0$ for each replication.

Figure 6: Implications of an imprecise execution of the evaluation bias

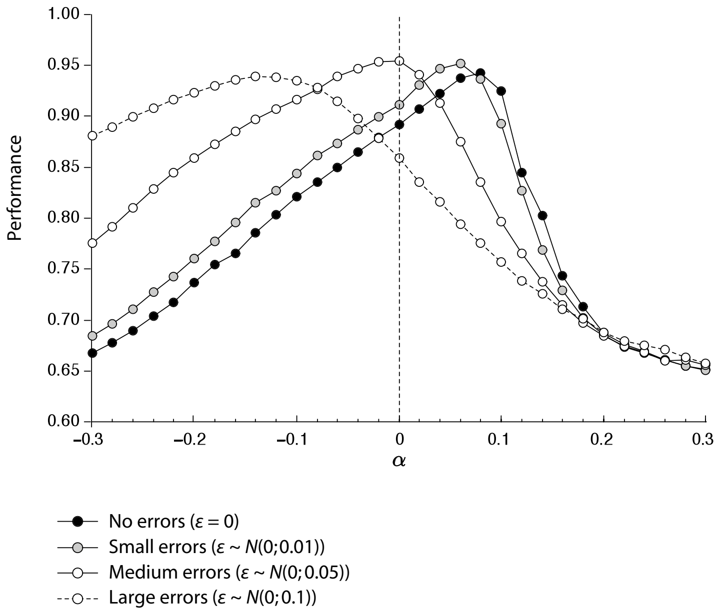

This figure reports the average firm performance in the long run (period 500) over 5,000 stable landscapes with $N = 8$, $K = 5$. Firms differ in their imprecision *(ε)* while executing a desired evaluation bias *(α)*, i.e., they implement $\alpha + \varepsilon$ instead of α, with ε being randomly drawn in each period.

On average, varying the lack of managerial control in the execution of the evaluation process has no major effect on the highest-possible performance. Instead, for all levels of control, the maximum performance ranges between 0.94 and 0.96, but the type and level of the evaluation bias (α) that result in this performance vary dramatically. Indeed, for low levels of control (high ε), the evaluation bias that yields the highest performance may even become negative. This result arises because, depending on the level of control over the execution process, different evaluation biases will induce the sporadic, slight overestimation of new alternatives. These are the choices that prove most beneficial by increasing exploration to a healthy degree, and which, over time, result in higher performance.

4.4 Effects of Environmental Selection Pressures

The analyses above establish that a slight pro-innovation bias results in the highest long-run performance in complex and stable environments, when managers search locally, and when they have full control over the implementation and execution of the evaluation process. Given these findings, we might conclude that if managers are interested in long-run survival and performance, they should adopt such a bias. However, if organizations face selection pressures, then firms with this particular evaluation bias may not survive to reap the advantages of a pro-innovation bias.

In *Figure 7,* we report the survival chances of organizations with values of α that in the long run are either optimal (a high-performing pro-innovation bias), fully objective, or suboptimal (a low-performing status-quo bias). In $t = 0$, we generate a population in which these three types of organizations are equally likely. In the left-hand panel of *Figure 8*, we report their share in the surviving firms if the lowest-performing 1%, 50%, 95%, or 99% of all firms are deselected in the short run $(t = 20)$. The right-hand panel reports the according result for the long run $(t = 200)$.

Figure 7: Impact of the selection environment

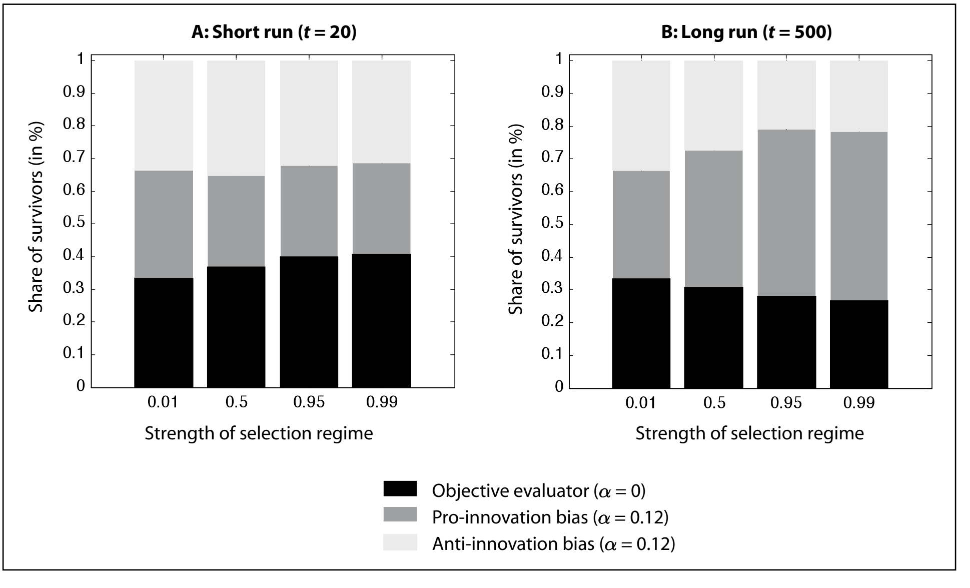

This figure reports, for selection regimes of different strength, the share of surviving firms given that selection occurs in the short run (period 20) or the long run (period 200). Firms differ in their evaluation bias (α). In the first period ($t = 0$), one third of the population of firms evaluates new alternatives fully objectively ($\alpha = 0$), one third has a pro-innovation bias ($\alpha = 0.12$), and one third has an anti-innovation bias ($\alpha = -0.12$), respectively. All numbers are averages over 5,000 stable landscapes with $N = 8$, $K = 7$.

If selection occurs in the short run $(t = 20)$, then organizations with an evaluation bias that is optimal in the long run are more likely to be deselected than are firms with a bias that is less suited to the long run. The reason is that short-term selection favors objective evaluators over firms with a pro-innovation bias, which may yield a long-term advantage but temporarily result in high-variance performance trajectories. We illustrate this feature of exploration-promoting evaluation processes in *Figure 8*. The figure reports the individual performance history of two firms, one of which is a fully objective evaluator and with the other of which has a slight pro-innovation bias.

Figure 8: Performance variation for a single run

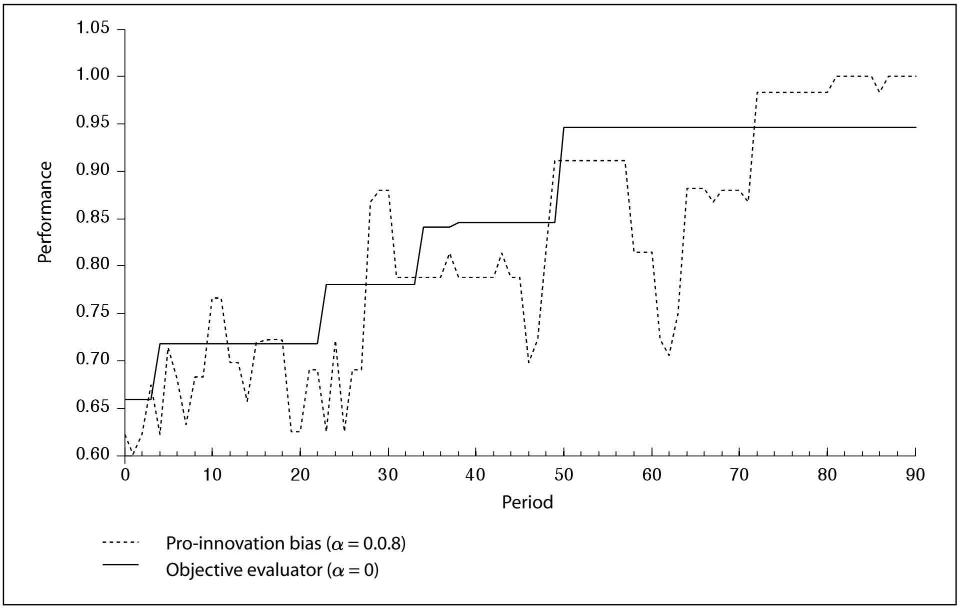

This figure reports the individual performance development of two firms on a single stable landscape with $N = 8$, $K = 5$. The two firms differ in their evaluation bias *(α)*.

If we were to refer back to *Figure 7*, we might expect that given their considerably better performance in the long run, firms with a pro-innovation bias come to dominate the population if selection occurs rather late in time, particularly in the presence of an extremely strong selection regime. However, although the surviving firms that possess a pro-innovation bias clearly outnumber both the objective evaluators and the firms with a status-quo bias, they never completely dominate the population; firms with a suboptimal evaluation bias still have a considerable chance to survive. The reason is that optimal organizational biases are associated with both a higher average performance and a higher variance in performance. In a population of organizations, organizations with extremely high and low performance are thus disproportionally characterized by an optimal bias. These low-performing organizations are very likely to be deselected.

5 Discussion and Conclusion

In this study, we have examined how systematic biases in the evaluation of new alternatives affect the ability of organizations to adapt to their environments. Our analysis was motivated by a puzzle: if, as certain studies on organizational change suggest, organizations should benefit from a pro-innovation bias, then why would organizations fail to adopt a more efficient evaluation strategy? A parsimonious simulation model allowed us to discuss three potential answers to these questions.

The first answer could be that the conjecture is simply incorrect, and that there is no value in having a pro-innovation bias (Rogers (2005)). However, our results indicate that just the opposite is true. We find that a slight pro-innovation bias can yield a considerable performance advantage in many circumstances, for instance, in medium to highly complex, stable environments, when long-term considerations matter, and when firms search locally. This result arises because a slight pro-innovation bias can trigger a firm to search more broadly and identify better solutions by implementing some alternatives that temporarily decrease performance but later help the firm identify better solutions. In contrast, a pro-innovation bias that is overly pronounced will render the search process inefficient and result in an exploration trap.

The second possible answer to our question might be that although the predicted relation is fully valid, it might be hard to detect empirically or be otherwise obscured (Siggelkow and Rivkin (2009)). Certainly, because of the asymmetric relation between the strength of a pro-innovation bias and the resulting organizational performance, an empirical investigation looking for an inverted *u*-shaped pattern might conclude that having no bias at all, i.e., evaluating new ideas correctly, denotes the optimal adaptation strategy. Even if based on empirical samples that are not subject to a survival bias, unless researchers are deliberately searching for an asymmetric relation, statistical analyses can thus lead to the wrong conclusions. Moreoever, our study shows that the positive performance effects of a pro-innovation bias are highly contingent on a number of internal and external factors, which could make them hard to detect statistically in a random sample of firms.

The third possible answer suggests that there are barriers that prevent organizations from achieving or securing a beneficial pro-innovation bias. In this context, our results are consistent with two arguments from prior research. One is that a status-quo bias can be a consequence of organizational learning, because the tendency to reproduce successful actions comes at the cost of underestimating unknown alternatives (Denrell and March (2001)). The other argument is that because innovation-seeking firms can exhibit high-variance performance trajectories, their survival rates decrease in the sense that firms may be deselected in the short run, even though their long-run performance might be superior (Levinthal and Posen (2007)). We extend these arguments by showing that even in the absence of learning and selection, trying to achieve a beneficial evaluation bias involves a high risk, partly because of its asymmetric performance implications and partly because its value is highly contingent on the complexity of the environment. Therefore, firms will in most cases be better off erring on the side of a conservative pro-innovation bias or even an anti-innovation bias.

Our study considers several other facets. First, it provides insights on the performance implications of deviations from rationality in the context of organizational search. We show that once an agent deviates from the assumptions of rationality (for example, in the sense that the agent is very limited in the number of alternatives that he can generate), or is trying to meet all other assumptions (such as a fully objective evaluation of the newly-generated alternatives) this divergence will no longer lead to the best outcome (see also Lipsey and Lancaster's (1956)) theory of the second best). Instead, if one assumption is not met, violating other assumptions may actually be performance-increasing. As a result, firms that have a biased evaluation process, but that commit systematic errors in evaluating new alternatives, may even outperform firms with unbiased evaluation processes. If an organization can generate high levels of variety, its evaluation skills should be similarly well developed. In the extreme case, i.e., when an organization is a global searcher, any bias in the evaluation of alternatives would have negative performance consequences.

This interdependence between processes of alternative generation and evaluation also has managerial implications. Consider the case in which an organization has identified the optimal level of a pro-innovation bias subject to its ability to generate alternatives. It may now seek to further strengthen its ability to adapt to its environment by increasing the number of alternatives it generates. Perhaps this step to becoming closer to being omniscient has positive performance implications. However, our results suggest that if the evaluation bias remains unchanged, then performance actually suffers. Similarly, if, to become more rational, the organization tries to reduce its pro-innovation bias, then performance will suffer if the firm does not simultaneously increase the number of alternatives it generates. Instead, if an organization does not wish to experience negative performance consequences, then both abilities must co-evolve.

Our study also points to the managerial challenges of becoming more rational. Ideally, a perfectly rational organization would consider every alternative in its analysis and would have no bias in its assessment. However, human decision makers are bounded in their ability to generate alternatives and to evaluate these alternatives objectively. Yet the attempt to become more rational is not accompanied by consistent increases in performance. Clearly, an organization with a smaller search radius can be considered less rational than an organization with a higher search radius. At the optimum level of evaluating alternatives, increasing the search radius can have severe negative performance consequences. Similarly, becoming less biased in the evaluation of new alternatives (while leaving the search radius constant) can hamper performance. However, with a pro-innovation bias that is at less than the optimal level, becoming more rational in terms of alternative generation can significantly increase performance. Thus, these interdependencies may also explain why most organizations "prefer" a pro-innovation bias that is below its optimal level. Decreasing the extent of deviation from perfect rationality does not always increase performance, and it may become more attractive for organizations to stick to their local optimum of imperfect rationality.

We also expand the work on search in complex systems (Rivkin and Siggelkow (2007)). When it comes to exploration, the pitfalls of local search in complex systems are widely accepted. Since trade-offs between the different elements of a complex system result in a highly rugged performance landscape, firms that are searching locally may become trapped on low local peaks. Our results indicate that a pro-innovation bias may denote a helpful antecedent to exploration. Overestimating some (objectively) inferior alternatives affects the perceived ruggedness of the landscape. Thus, a firm might perceive the landscape as "neutral" (Huynen, Stadler, and Fontana (1996)) or even turn valleys into mountains, creating networks that induce movement on the landscape, thus letting the firm muddle through towards better solutions.

References

Abrahamson, Eric (1991), Managerial fad and fashion: The diffusion and rejection of innovations, *Academy of Management Review* 16, 586-612.

Baumann, Oliver (2010), Search, failure, and the value of moderate patience, *sbr* 62, 233-258.

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

Birkinshaw, Julian and Cristina Gibson (2004), Building ambidexterity into an organization, *MIT Sloan Management Review* 45, 47-57.

Brown, Showna L. and Kathleen M. Eisenhardt (1997), The art of continuous change: Linking complexity theory and time-paced evolution in relentlessly shifting organizations, *Administrative Science Quarterly* 42, 1-34.

Burton, Richard and Borge Obel (1995), The validity of computational models in organization science: From model realism to purpose of the model, *Computational and Mathematical Organization Theory* 1, 57-72.

Christensen, Clayton M. (1997), *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*, Harvard, MA: Harvard Business School Press.

Christensen, Clayton M., Stephen P. Kaufman, and Willy C. Shih (2008), Innovation killers: How financial tools destroy your capacity to do new things, *Harvard Business Review* 86, 98-105.

Christensen, Michael and Thorbjørn Knudsen (2010), Design of decision-making organizations, *Management Science* 56, 71-89.

Cohen, Michael D., James G. March, and Johan P. Olsen (1972), A garbage can model of organizational choice, *Administrative Science Quarterly* 17, 1-25.

Cyert, Richard M. and James G. March (1963), *A Behavioral Theory of the Firm*, Englewood Cliffs, NJ: Prentice Hall.

D'Aveni, Richard A. (1994), *Hypercompetition: Managing the Dynamics of Strategic Maneuvering*, New York: Free Press.

Davis, Jason P., Kathleen M. Eisenhardt, and Christopher B. Bingham (2007), Developing theory with simulation methods, *Academy of Management Review* 32, 480-499.

Davis, Jason P., Kathleen M. Eisenhardt, and Christopher B. Bingham (2009), Optimal structure, market dynamism, and the strategy of simple rules, *Administrative Science Quarterly* 54, 413-452.

Denrell, Jerker and James G. March (2001), Adaptation as information restriction: The hot stove effect, *Organization Science* 12, 523-538.

Eisenhardt, Kathleen M. and Behnam N. Tabrizi (1995), Accelerating adaptive processes: Product innovation in the global computer industry, *Administrative Science Quarterly* 40, 84-110.

Gavetti, Giovanni and Daniel A. Levinthal (2000), Looking forward and looking backward: Cognitive and experiential search, *Administrative Science Quarterly* 45, 113-137.

Gilovich, Thomas, Dale W. Griffin, and Daniel Kahneman (2002), *Heuristics and Biases: The Psychology of Intuitive Judgement*, Cambridge: Cambridge University Press.

Hannan, Michael T. and John Freeman (1984), Structural inertia and organizational change, *American Sociological Review* 49, 149-164.

Hannan, Michael T. and Glenn R. Carroll (1992), *Dynamics of Organizational Populations: Density, Competition, and Legitimation*, New York: Oxford University Press.

Hayward, Mathew L. A., Dean A. Shepherd, and Dale Griffin (2006), A hubris theory of entrepreneurship, *Management Science* 52, 160-172.

Henderson, Rebecca and Kim B. Clark (1990), Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms, *Administrative Science Quarterly* 35, 9-30.

Holland, John H. (1975), *Adaptation in Natural and Artificial Systems*, Ann Arbor, MI: University of Michigan Press.

Huynen, Martijn A., Peter F. Stadler, and Walter Fontana (1996), Smoothness within ruggedness: The role of neutrality in adaptation, *Proceedings of the National Academy of Sciences USA* 93, 397-401.

Kahneman, Daniel, Paul Slovic, and Amos Tversky (1982), *Judgment under Uncertainty: Heuristics and Biases*, Cambridge: Cambridge University Press.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler (1991), Anomalies: The endowment effect, loss aversion, and status quo bias, *The Journal of Economic Perspectives* 5, 193-206.

Kahneman, Daniel and Amos Tversky (2000), *Choices, Values, and Frames*, Cambridge: Cambridge University Press.

Kanter, Rosabeth M. (1988), When a Thousand Flowers Bloom: Structural, Collective, and Social Conditions for Innovation in Organizations, in Barry M. Staw and L. L. Cummings (eds.), *Research in Organizational Behavior*, Vol. 10, Greenwich, CT: J.A.I. Press, 169-211.

Katz, Ralph and Thomas J. Allen (1982), Investigating the Not Invented Here (NIH) syndrome: A look at the performance, tenure and communication patterns of 50 R&D project groups, *R&D Management* 12, 7-19.

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

Kauffman, Stuart A. (1995), *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*, New York: Oxford University Press.

Knudsen, Thorbjørn and Daniel A. Levinthal (2007), Two faces of search: Alternative generation and alternative evaluation, *Organization Science* 18, 39-54.

Leonard-Barton, Dorothy (1992), Core capabilities and core rigidities: A paradox in managing new product development, *Strategic Management Journal* 13, 111-125.

Levinthal, Daniel A. and James G. March (1981), A model of adaptive organizational search, *Journal of Economic Behavior and Organization* 2, 307-333.

Levinthal, Daniel A. and James G. March (1993), The myopia of learning, *Strategic Management Journal* 14, 95-112.

Levinthal, Daniel A. (1997), Adaptation on rugged landscapes, *Management Science* 43, 934-950.

Levinthal, Daniel A. and Hart E. Posen (2007), Myopia of selection: Does organizational adaptation limit the efficacy of population selection?, *Administrative Science Quarterly* 52, 586-620.

Levitt, Barbara and James G. March (1988), Organizational learning, *Annual Review of Sociology* 14, 319-340.

Lipsey, Richard G. and Kelvin Lancaster (1956), The general theory of second best, *The Review of Economic Studies* 24, 11-32.

Lowe, Robert A. and Arvids A. Ziedonis (2006), Overoptimism and the performance of entrepreneurial firms, *Management Science* 52, 173-186.

Macy, Michael W. and Robert Willer (2002), From factors to actors: Computational sociology and agent-based modeling, *Annual Review of Sociology* 28, 143-166.

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

March, James G. (1991), Exploration and exploitation in organizational learning, *Organization Science* 2, 71-87.

March, James G. (2006), Rationality, foolishness, and adaptive intelligence, *Strategic Management Journal* 27, 201-214.

McGrath, Rita Gunther (1997), A real options logic for initiating technology positioning investments, *Academy of Management Review* 22, 974-996.

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

Porter, Michael E. (1996), What is strategy?, *Harvard Business Review* 74, 61-78.

Rivkin, Jan W. (2000), Imitation of complex strategies, *Management Science* 46, 824-844.

Rivkin, Jan W. and Nicolaj Siggelkow (2007), Patterned interaction in complex systems: Implications for exploration, *Management Science* 53, 1068-1085.

Rogers, Everett M. (1962), *Diffusion of Innovation*, New York: Free Press.

Rogers, Everett M. and Floyd F. Shoemaker (1971), *Communication of Innovations: A Cross-Cultural Approach*, New York: Free Press.

Rogers, Everett M. (2005), *Diffusion of Innovation*, New York: Free Press.

Rosenkopf, Lori and Paul Almeida (2003), Overcoming local search through alliances and mobility, *Management Science* 49, 751-766.

Sah, Raaj K. and Joseph E. Stiglitz (1996), The architecture of economic systems: Hierarchies and polyarchies, *American Economic Review* 76, 716-727.

Samuelson, William and Richard Zeckhauser (1988), Status quo bias in decision making, *Journal of Risk and Uncertainty* 1, 7-59.

Sheth, Jagdish N. (1981), Psychology of innovation resistance: The less developed concept (LDC) in diffusion research, *Research in Marketing* 4, 273-282.

Siggelkow, Nicolaj (2002), Evolution toward fit, *Administrative Science Quarterly* 47, 125-159.

Siggelkow, Nicolaj and Jan W. Rivkin (2009), Hiding the evidence of valid theories: How coupled search processes obscure performance differences among organizations, *Administrative Science Quarterly* 54, 602-634.

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

Stuart, Toby E. and Joel M. Podolny (1996), Local search and the evolution of technological capabilities, *Strategic Management Journal* 17, 21-38.

Sturdy, Andrew and Christopher Grey (2003), Beneath and beyond organizational change management: Exploring alternatives, *Organization* 10, 651-662.

Sull, Donald N. (2005), *Why Good Companies Go Bad and How Great Managers Remake Them*, Boston, MA: Harvard Business School Press.

Sutton, Richard S. and Andrew G. Barto (1998), *Reinforcement Learning: An Introduction*, Cambridge, MA: MIT Press.

Tadelis, Steven (2007), The innovative organization: Creating value through outsourcing, *California Management Review* 50, 261-277.

Teece, David J., Gary Pisano, and Amy Shuen (1997), Dynamic capabilities and strategic management, *Strategic Management Journal* 18, 509-533.

Van de Ven, Andrew H. (1986), Central problems in the management of innovation, *Management Science* 32, 590-676.

Williamson, Oliver E. (1975), *Markets and Hierarchies: Analysis and Antitrust Implications: A Study in the Economics of Internal Organization*, New York: Free Press.