Oliver Baumann*

Search, Failure, and the Value of Moderate Patience†

Abstract: Experimentation in complex and novel environments entails a high risk of failure. Both intuition and a substantial body of research thus suggest that successful innovation requires patience – persistence despite failures – to ensure broad search for good solutions. But is patience generally beneficial? And what trade-offs are associated with different levels of patience? Using an agent-based simulation model, I show that the link between patience and innovation is intricate: Moderate levels of patience do promote broad and effective search. High levels of patience, in contrast, can have unintended effects and even decrease performance despite further increasing the overall degree of search. This result arises as firms fail to confine the exploration that patience brings about, instead searching too erratically and "drifting" away from potentially good solutions. Furthermore, as translating the gains of patience into performance improvements requires time, low levels of patience are generally optimal for shorter time frames. The findings show that paying attention to how patience affects explorative and exploitative aspects of search alike, rather than exploration alone, can shed light on when patience may effectively boost innovation.

Keywords: Search, experimentation, innovation, complexity, agent-based simulation

JEL classification: C63, D83, O31

---

* Oliver Baumann, Munich School of Management, Ludwig-Maximilians-Universität München, Ludwigstr. 28, 80539 München, Germany, Tel.: +49-89-2180-2982, Fax: +49-89-218-99-2982, E-Mail: baumann@lmu.de

† For helpful comments on a prior version of this paper, I thank Shmuel Ellis, Donald Hatfield, Mu-Yen Hsu, Markus Lang, Dan Levinthal, Dirk Martignoni, Nils Stieglitz, two anonymous reviewers, and participants at the DRUID Summer Conference, the AOM Annual Meeting, the Israel Strategy Conference, and the Workshop of the Organization Section of the German Academic Association of Business Research. All remaining errors are mine.

Search, Failure, and the Value of Moderate Patience

1 Introduction

When firms need to solve complex and novel problems, they face a managerial dilemma: As complexity results in high-dimensional solution spaces with numerous local optima, firms must search broadly for good solutions (Simon (1962); Levinthal (1997)). Under conditions of genuine novelty, on the other hand, action-performance links are largely undefined, and firms often cannot but evaluate potential solutions through "online" experiments, i.e., by implementing an alternative and observing its performance (Gavetti and Levinthal (2000)). Yet while such experiential learning creates new knowledge, it comes at a high risk of failures. Both intuition and various accounts in the academic (e.g., Fleming (2001); Thomke (2003); Cannon and Edmondson (2005)) and practitioner (Farson and Keyes (2002); Petroski (2006)) literatures thus suggest that successful innovation requires patience, i.e., persistence despite failures.[1]

Consider, for illustration, the well-known anecdote of how Thomas A. Edison designed the first commercially useful incandescent light bulb (Israel (1998)). While searching for a high-performing configuration of the various interdependent design elements (which eventually included using a U-shaped carbon filament in an oxygen-free environment), Edison experimented with thousands of alternative design variants, appreciating the numerous failures he experienced ("I have not failed. I've just found 10,000 ways that won't work") and committing himself to being patient ("I am not discouraged, because every wrong attempt discarded is another step forward").

Hence, this line of reasoning suggests that patience is positively linked to innovation – firms that continue to search despite experiencing failures will experiment with a higher number of alternatives and are thus more likely to identify innovative solutions. But is patience generally beneficial? And what performance trade-offs are associated with different levels of patience?

The issue of choosing a degree of patience relates to a broader class of problems that go beyond the challenge faced by firms that are experimenting with different design variants. Similar to complex technologies, firms as a whole have been conceptualized as systems of interdependent activities (Milgrom and Roberts (1995); Porter (1996); Sig-

[1] The notion of "patience" comes with a number of (slightly) different connotations. In the course of this paper, it is conceptualized as the perseverant active search despite failures, rather than the passively waiting for an unpleasant situation to pass by.

gelkow (2002); Gavetti and Rivkin (2007)). Managers within firms are seen as trying to find coherent and high-performing sets of activities. Yet when an environmental shock deteriorates their existing configuration, firms may have to experiment with completely new courses of action. How patient should firms be when searching for new activity configurations? A similar problem arises for new ventures. Start-up firms have to engage in a range of activities before they can launch (Garud and Van De Ven (1992); Ravasi and Turati (2005)). Yet in emerging industries, business models and knowledge about successful practices are not yet established. Again, what are the implications of exhibiting different levels of patience when searching for a set of good activity choices?

To shed light on these issues, this paper draws from two related but somewhat separate literatures. One stream of work has generated insight into the role of failure and patience in innovation (e.g., Vincenti (1990); Petroski (1992); Thomke (1998); Lee et al. (2004)), but does not speak to how patience affects the search processes that lead to performance. The other stream has modeled problem-solving search in complex and novel environments (e.g., Levinthal (1997); Rivkin (2000); Winter et al. (2007)), but has largely assumed that firms can evaluate new alternatives through "offline" analysis (i.e., without having to implement them), whereas the notion of "online" trials has received little attention (for an exception, see, e.g., Gavetti and Levinthal (2000)). In consequence, little is known about how search is affected by different levels of patience, and whether and how the value of patience is contingent upon the complexity and novelty of a problem.

This paper provides some initial integration of these two literatures by embedding stylized features of failure and patience in problem-solving search into an agent-based simulation model. The model contains firms that face complex and novel problems and that search for better alternatives to their current solutions. By controlling how many failures firms are willing to tolerate, I can systematically analyze how patience affects the dynamics of search and, ultimately, firm performance.

Even in this simple model set-up, counterintuitive effects can arise that show how patience and innovation interact in non-trivial ways. I find that moderate levels of patience – tolerating some degree of failure before abandoning a search path and starting over – do promote broad and effective search. Also, the more complex and novel a problem is, the more distinct this benefit of moderate patience becomes. High levels of patience, in contrast, can have unintended effects and even decrease performance despite further increasing the overall degree of exploration. This result arises as firms fail to confine the exploration that patience brings about, instead searching rather erratically and "drifting" away from potentially good solutions. Furthermore, as translating the gains of patience into performance improvements requires time, low levels of patience that result in more local search are generally optimal for shorter time frames.

The results of this paper relate to the oft-cited trade-off between exploration and exploitation (Holland (1975); March (1991)). They suggest that in order to gain insight into the link between patience and innovation, paying attention to how patience translates into exploration is necessary, but not sufficient. Rather, attention must be paid to how patience affects the process by which the exploration of a novel context will eventually be replaced by the exploitation of the newly identified opportunities. The findings also suggest that, if managers seek to boost innovation, trying to act upon their organization's level of patience with respect to affecting this transition is more valuable than merely trying to increase the sheer amount of search.

2 Literature and propositions

2.1 Prior research on failure and patience in innovation

Given the costs and wasted efforts that failed experiments entail, problem solvers will set a limit to how many failures they are willing to tolerate. But what determines this level of patience? Clearly, one factor is personality. For example, the high degree of persistence that can often be observed among entrepreneurs or engineers (Garud and Van De Ven (1992); Forbes (2005); Lowe and Ziedonis (2006)) is typically attributed to behavioral traits such as overconfidence, optimism, passion, or foolishness (Kahneman et al. (1982); Baron (1988); March (2006)). Most other individuals, in contrast, tend to avoid courses of action in which failures are likely (Thomke (1998); Lee et al. (2004)).

Another determinant is the experimentation environment that affects both the costs of failing and the costs and ease of reversing an experiment. The latter issue relates to the fact that the distribution of alternatives is not known in advance, and that alternatives are usually encountered sequentially. Hence, if a decision maker hopes to find an even better alternative, he may discard his current solution despite the alternative being among the best in the population. If he is able to return to an option that has been permitted to "pass by", he is said to "recall" the alternative, thus reversing the previous experiment(s) (Gigerenzer et al. (1999)). Consider, for example, a software designer. If she finds that her recent experiments have deteriorated the performance of the project, she may easily revoke an old version that she knows has a satisfactory performance. In this case, the cost to recall will only pertain to the time invested into making and testing the (eventually fruitless) modifications. A firm that is unsatisfied with its performance after an acquisition, in contrast, can likewise recall a previous state. However, the degree of sunk costs and the time required to "un-do" the investment will be of a different order of magnitude. In consequence, patience will typically be higher in a lab environment, where experiments are frequently undertaken with the knowledge that they may fail, whereas in most organizations, there is high pressure to avoid failure (Denrell and March (2001); Thomke (2003); Lee et al. (2004)).

A third determinant relates to the time horizon. In basic R&D, for example, it is common wisdom that – in order to explore freely and broadly – problem solvers need sufficient time, during which they are willing to tolerate longer periods in which their search efforts may not bear fruit. On the other hand, consider product development in a fast-paced environment (Bourgeois and Eisenhardt (1988); Eisenhardt (1989); Fine (1998)). Here, decisions need to be made quickly, and product designers or managers are usually given little time to demonstrate the feasibility of a concept and improve the performance of the initial prototypes. Rather, by techniques such as "concurrent engineering" or "front-loading" (Loch and Terwiesch (1998); Thomke and Fujimoto (2000); Loch et al. (2001)), firms are trying to speed up the search process and evaluate the potential of a product early on. In this context, product designers may not be able to exhibit much patience for tinkering around with a concept. Instead, they may have to make the best out of a given idea and limited time, i.e., exploit a current concept as efficiently as possible.

But why, in the first place, should firms be patient? One major advantage of online experiments, despite the failures they may entail, is that they generate knowledge when the paths of cause and effect are uncertain (Allen (1977); Thomke (1998); Fleming (2001); Thomke (2003)). Hence, failures are appreciated as they often denote an interim stage on a problem solver's journey of knowledge creation, and tolerating them may be necessary to accumulate a "sufficiently" large and focused body of knowledge (Popper (1959)). Vincenti (1990)), for instance, gives detailed insight into the processes of knowledge creation in the early days of the aviation industry: As engineers faced the challenge of designing airfoils that exhibited the desired lift and drag needed for a particular aircraft, they tested hundreds of different airfoils in an online manner. Only after experience had generated sufficient insight into what worked (and what didn't) under which conditions, aerodynamic theory could be derived which could then substitute for many online trials. In other words, further alternatives could be evaluated in an offline manner, i.e., without the risk of failures. A different appreciation of failure in engineering design is given by Petroski (1992); 1994); 2006)), who points out for various cases from engineering history how multiple failures often proceeded the evolution of a successful design. In his reading, failures can act as a stepping stone towards more breakthrough solutions – by triggering further search that may lead to the conception of high-performing designs which might not have been cognitively conceivable without taking the intermediate steps that proved to be failures. In sum, these arguments suggest a positive relationship between patience and performance:

> *Proposition 1a. Higher levels of patience lead to broader search that will result in higher performance in the long run.*

However, some research also points out that the value of failure not only stems from creating knowledge or acting as a stepping-stone, but from the fact that it can trigger change. For instance, failures may convince problem solvers to abandon their current search path and try other routes instead (Petroski (1992)). In the design process, for example, "the final version is [sometimes] closer to the first than any of the intervening versions" (Petroski (1992, p.77)). Hence, if a chosen search path is a dead end, high patience – often driven by an escalation of commitment – may translate into extensive search, but still result in the eventual failure of the particular project (Biyalogorsky et al. (2006); Välikangas (2007)). Hence, understanding failures becomes important to improve performance (Thomke (2003)): As "products are the result of as many failed experiments as successful ones [,] an innovation process […] is at least partially based on 'accumulated failure' that has been carefully understood" (Thomke (2003: p. 27)). In other words, failures need not only be endured, but the resulting knowledge must be utilized to correct one's actions (Mach (1905)). Firms, however, often possess no systematic culture and process for learning from failure, or fail to address this issue altogether (Tucker and Edmondson (2003); Baumard and Starbuck (2005); Cannon and Edmondson (2005)), "sweeping" failures "under the carpet". These considerations raise an alternative view:

> *Proposition 1b. Higher levels of patience may not necessarily lead to higher organizational performance if the firm fails to learn from the increased level of search.*

Finally, the above discussion also suggests that when the level of patience is low, a firm will search less broadly but explore only the neighborhood of its current solution, as it will not be willing to tolerate failures for an extended period of time. At the same time, however, the local knowledge that is created enables the firm to take actions in this local context more quickly. To summarize:

Proposition 2. Low levels of patience lead to local search and support the short-term improvement of a firm's current status.

2.2 Prior research on search in complex and novel environments

Due to bounded rationality (Simon (1955); Simon (1956)), firms need to search for new decision alternatives rather than optimize over a collectively known set of options as assumed in neoclassical theory (March and Simon (1958); Cyert and March (1963)). How, then, do the complexity and novelty of a problem affect the dynamics of search and the value of different degrees of patience? Following Simon (1962)), I conceive of complex problems as systems that consist of a large number of elements that have many interactions. A well-known property of such systems is that with a rising number of interactions between its elements, local peaks – internally consistent configurations of the system elements that cannot be improved through incremental changes – proliferate (Kauffman (1995); Levinthal (1997); Rivkin and Siggelkow (2007)). This increases the risk that constrained exploration will trap a firm on a low local peak, thus making broader exploration more beneficial. Furthermore, a higher number of system elements, ceteris paribus, makes problem-solving search more demanding as it significantly increases the number of potential combinations of these elements. In other words, the space in which problem solver conduct their search for good solutions expands as systems become larger (Simon (1996)). Hence, in order to identify superior configurations, higher levels of search become crucial. With respect to the complexity of a problem, this suggests:

Proposition 3. High interdependence between the system elements makes high levels of patience more valuable.

Proposition 4. A large number of system elements makes high levels of patience more valuable.

But how does a firm know whether a new alternative is satisfactory? While prior modeling efforts have mainly focused on the generation of alternatives, i.e., the discovery aspects of organizational search, the mechanisms by which new alternatives are evaluated have received little attention. Here, the dichotomy of "offline" and "online" evaluation (Lippman and McCall (1976); Levitt and March (1988); Gavetti and Levinthal (2000)) offers helpful assistance: Offline (or cognitive) evaluation refers to situations in which actors can evaluate the usefulness of an alternative through thought experiments, calculations, computer simulations, or laboratory experiments, i.e., various mechanism that are close to Freud's (1912)) notion of thinking as "internalized experimental action" ("internalisiertes Probehandeln"). Hence, a firm will only adopt an alternative if offline reasoning has proven the superiority of the idea. Online (or experiential) evaluation, in contrast, is characterized by a strong "try-it-and see" aspect, requiring a firm to implement an alternative in order to learn about its value. Online experiments are thus obviously more risky than offline assessments, as they require to leave the status quo and because the outcome of the trial is uncertain and may result in a failure. Hence, when a firm is faced with problems, but already possesses a certain degree of domain-specific knowledge, it can make a higher number of evaluations in an offline manner, decreasing the need to explore experientially. Because the firm can avoid adopting (eventually) fruitless ideas, it can efficiently build upon its existing knowledge and quickly improve its performance – a behavior much in line with March's (1991)) notion of exploiting old certainties. If, in contrast, the firm has only little initial knowledge, it cannot but explore experientially initially, with higher levels of exploration proving more beneficial to generate broad knowledge. Regarding the novelty of a problem, this suggests:

> *Proposition 5. High levels of initial knowledge in the domain of the search support the exploitation of a firm's current situation and correspond to low levels of patience.*

> *Proposition 6. Low levels of initial knowledge in the domain of the search support the exploration of new alternatives and make high levels of patience more valuable.*

3 Model

To study the impact of different degrees of patience on the dynamics of search, I develop a simple agent-based simulation model. Computational models have gained broad popularity in studies of organizational search and learning (March (1991); Levinthal (1997); Gavetti and Levinthal (2000); Mihm et al. (2003); Winter et al. (2007)) for a variety of reasons (Davis et al. (2007); Harrison et al. (2007)). One is that they allow a more rigorous analysis than verbal analysis, forcing the modeler to make all underlying assumptions explicit. In contrast to algebraic approaches, on the other hand, computational models allow to incorporate a richer set of features into the analysis. Although they cannot yield "exact solutions" like closed-form techniques, they allow to model conditions of complex interactions under which algebraic approaches such as, e.g., the supermodularity framework for studying complementarities (Milgrom and Roberts (1990); Milgrom and Roberts (1995)) would be intractable. Most importantly, however, the paper is concerned with the question of how patience affects problem-solving search by decision makers that possess only bounded rationality. While exploring the underlying dynamics of search can be easily achieved with computational models, analytic models tend to be concerned with equilibria and not with the question of how, or whether, they will be attained.

The basic principle of agent-based simulation is straightforward (Macy and Willer (2002)): Decision-making agents (e.g., firms, managers, or designers) are confronted with controlled environments, they are equipped with heuristics to react to their environment, and the resulting behavior is recorded over time. By varying the behavior of the agents and the structure of the environment, I can systematically explore the impact and interdependence of the variables under consideration. In the following, I thus describe the environment that my modeled firms face as well as their behavior with regard to search and patience. Although computational approaches grant high degrees of freedom, the model does not represent any real-world context. Instead, it contains stylized elements that are essential to shed light on the abstract problem under investigation, thus following an established tradition in computational research to develop simple yet insightful models (Cohen et al. (1972); Nelson and Winter (1982); Burton and Obel (1995)). As such, it is the aim of many computational models (including the one developed in this paper) to generate thought-provoking insight about problems for which common intuition may be deceptive or inconsistent.

3.1 Complex problems

I conceptualize firms as facing a set of interdependent decisions (Porter (1996); Levinthal (1997); Siggelkow (2002)). In order to explore and learn about this environment, a firm's managers or designers need to make a multitude of decisions. For instance, the designers of a jet engine might have to decide on the power of the engine or the materials that are employed to manufacture it. A manager, for instance, might have to decide about the firm's product variety or about the features of its production system. Furthermore, many of these decisions interact with each other. For instance, the value of a powerful jet engine will depend on whether the structural properties of the material can endure the corresponding forces. Likewise, the value of flexible manufacturing capabilities will increase as a firm increases its product variety.

In the model, each firm must resolve N decisions a_1, a_2, …, a_N. Without loss of generality, I assume that each decision is binary. For instance, a_1 might denote the decision to increase product variety ($a_1 = 1$) or not ($a_1 = 0$). In consequence, a firm faces 2^N

possible configurations of choices, each of which can be represented by a binary vector $\mathbf{a} = (a_1, a_2, \ldots, a_N)$.

In computational studies of firms as complex adaptive systems, it has become 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 (the N decisions that the firm needs to make), and one "vertical" dimension that denotes the corresponding performance of each configuration. A performance landscape thus 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).

I create performance landscapes with a variant of the *NK* model (Kauffman (1993); Kauffman (1995)) – stochastically, yet in a well-controlled manner. The *NK* model has been developed in evolutionary biology and has recently been applied to a number of organizational issues (e.g., Levinthal (1997); Rivkin (2000); Ethiraj and Levinthal (2004); Lenox et al. (2006)). In the 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 ($\boldsymbol{a}_{-i}$) are resolved that interact with a_i. 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. Now, the landscape is highly 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 performance landscape. Particular values for all possible c_i's are drawn from a uniform distribution over the unit interval, i.e., $c_i(a_{i;}\ \boldsymbol{a}_{-i}) \sim u[0;1]$. Finally, the value $V(\boldsymbol{a})$ of a given set of choices $\boldsymbol{a}$ is calculated as an average of its N performance contributions, i.e., $V(\boldsymbol{a}) = [c_1(a_{1;}\ \boldsymbol{a}_{-1}) + c_2(a_{2;}\ \boldsymbol{a}_{-2}) + \ldots + c_N(a_{N;}\ \boldsymbol{a}_{-N})] / N$. In sum, the parameters N and K thus allow to tune the complexity of a firm's environment in terms of size (N) and interdependence (K).

The landscape metaphor allows an intuitive representation of organizational search: A firm inhabits – subject to its configuration of choices $\mathbf{a}$ – 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 reach high points on the performance landscape – configurations of choices that together create a high performance. Below, I describe how firms generate and evaluate new alternatives and how they react to failure that unknown alternatives may bring about.

3.2 Problem-solving search

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 (given $N = 4$), it would have four alternatives available: 1001, 1010, 1100, and 0000. For instance, a jet engine designer might come up with the idea to modify the composite material that was used so far. In a similar manner, a manager might think about modifying the firm's current production system by introducing a new process control software. Among the N possible local alternatives, each manager picks one at random. Hence, this procedure for generating alternatives that are very similar to the existing configuration of choices represents a behavior of local search – a central feature in both theoretical (March and Simon (1958); Cyert and March (1963); Nelson and Winter (1982)) and empirical (Stuart and Podolny (1996); Rosenkopf and Almeida (2003)) accounts of organizational learning and adaptation. In other words, cognitive bounds prevent managers from coming up with highly innovative ideas, i.e., with alternative configurations that differ in multiple dimensions from the status quo. (In the robustness section, I relax this assumption and show that the main results also hold if the bounds on managers' rationality are less severe.)

Subsequently, the evaluation of a newly identified alternative $\tilde{\boldsymbol{a}}$ can either proceed in a cognitive (offline) or experiential (online) manner. If the firm needs to explore experientially, it must first adopt the new alternative, i.e., move from point $\boldsymbol{a}$ to the nearby point $\tilde{\boldsymbol{a}}$ on the landscape. Subsequent to the adoption of $\tilde{\boldsymbol{a}}$, the firm learns about its value $V(\tilde{\boldsymbol{a}})$, i.e., about whether the new alternative denotes an improvement over the previous configuration ($V(\tilde{\boldsymbol{a}}) > V(\boldsymbol{a})$) or not ($V(\tilde{\boldsymbol{a}}) < V(\boldsymbol{a})$). In the latter case, the firm has experienced a failure. Independent from the exact performance of the new alternative, however, the firm has created knowledge about the performance implications of a configuration that was (as yet) unknown. Should the firm encounter the same configuration again during the course of its further search, it could evaluate its relative attractiveness in an offline manner.[2] Finally, in the next period, the firm will generate another local alternative based on its new status quo set of choices.

The offline evaluation of a new alternative $\tilde{\boldsymbol{a}}$ proceeds in the reverse order. Here, the firm can determine the value $V(\tilde{\boldsymbol{a}})$ of the alternative without implementing the idea. If the firm finds that the alternative denotes an improvement, i.e., if $V(\tilde{\boldsymbol{a}}) > V(\boldsymbol{a})$, the firm will adopt it and move from point $\boldsymbol{a}$ to the nearby point $\tilde{\boldsymbol{a}}$ on the landscape.[3] If, in contrast, the value of $\tilde{\boldsymbol{a}}$ is lower (or equal) than the value of the firm's current alternative ($V(\tilde{\boldsymbol{a}}) \leq V(\boldsymbol{a})$), the firm will discard the alternative and remain on its current "spot" on the landscape, generating another local alternative in the next period.

Which alternatives a firm can assess offline is determined by the parameter *KNOW* (with $0 \leq KNOW \leq 1$) that represents the firm's initial degree of knowledge in the domain of the new problem. Thus, *KNOW* can be regarded as an inverse measure for the fundamental novelty of the problem that the firm needs to explore. The value of *KNOW* may be influenced by the fact that the firm has encountered related problems before. It may also be actively shaped by investing into human capital or new technologies that increase the scope of offline evaluation capabilities. Hence, when $KNOW = 0$, the firm starts its exploratory search without any offline knowledge, i.e., it initially needs to evaluate any alternative experientially. If, in contrast, $KNOW = 1$, then the firm has the expertise to assess the value of any potential alternative in an offline manner. For all intermediate values of *KNOW*, the firm possesses offline evaluation capabilities for the corresponding (randomly determined) fraction of all possible configurations.

Once a firm has implemented a configuration and knows that this alternative cannot be further improved by any local alternative, the search ends. In this case, the firm has reached a local peak on the landscape, which acts as a "competency trap" (Levinthal and March (1981); Levitt and March (1988)) and terminates a firm's exploratory search.[4] On the other hand, a firm can only be sure to have reached a local peak when it has generated enough knowledge to determine by means of offline reasoning that all local alternatives would yield a worse performance than its status quo set of choices. If

[2] This assumes that a single evaluation act is sufficient to understand the (real) value of a new alternative. Issues such as search depth (Katila and Ahuja (2002)) or reinforcement learning (Sutton and Barto (1999)) are thus beyond the scope of this model. Also, it is assumed that no "organizational forgetting" occurs during the search process. All of these aspects might denote fruitful potential for further work.

[3] To focus exclusively on the dynamics of patience and search, the model assumes that superior (offline) alternatives are always adopted. It thus abstracts from many real-world intricacies involved in the adoption of innovations such as, for instance, the role of promoters (Hauschildt and Gemünden (1999)).

[4] 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.

there remains only one local alternative that the firm cannot assess offline, the firm can also not be sure that it has reached a local peak, as the unknown alternative might yield an even higher performance. (As I will show below, governing this tension may help a firm to avoid low local peaks.)

3.3 *Different levels of patience*

The uncertainty involved in an online evaluation poses the risk that the newly-adopted alternative yields a lower performance than the previous one.[5] Yet in order to explore uncharted areas on the landscape, firms may have to tolerate some degree of failure, i.e., they must be willing to accept a certain number of downhill moves on their quest for high points on the performance landscape. To represent different approaches of how firms might deal with this issue, I introduce a patience parameter, *PAT*, that denotes the maximum number of periods that a firm is willing to tolerate – while exploring experientially – a performance that is below the performance of the best alternative encountered so far.[6] Hence, the firm always "remembers" the present best alternative, and during its online evaluations, counts the number of periods in which it does not come up with an alternative that exceeds this benchmark.[7] If a firm finds a better alternative, it sets its "failure counter" back to zero and uses the newly-found alternative as the new benchmark. If, in contrast, the firm does not come up with an alternative after *PAT* online trials, it will discontinue its current path of exploration and re-implement the benchmark alternative. For instance, if *PAT* = 1, the firm will be rather exploration-averse and make no more than one online evaluation before returning to its previous alternative, should the trial be unsuccessful. If *PAT* = 10, in contrast, a firm is much more exploration-seeking and will tolerate up to 10 periods of underperformance.

[5] In the model, I assume that firms have no assumptions about unknown alternatives. I also make no assumptions about a firm's risk preferences. The modeled firms do not seek nor try to avoid risky online trials subject to their current performance. They simply generate and assess new alternatives, independent of which "type" of evaluation (online, offline) is possible.

[6] Offline evaluations that do not yield a performance improvement are not counted as failures, because they do not require a firm to leave its status quo. Incorporating them into the failure count, however, does not qualitatively affect the results.

[7] I do not discriminate between different degrees of failure, which may be a fruitful avenue for further research, but count each period of below-benchmark performance as a failure, independent of whether performance is only slightly or rather significantly below the benchmark.

4 Results

To study the effects of patience on the dynamics of search, I placed firms that differed in their level of patience (*PAT*) onto randomly chosen points of stochastically generated performance landscapes. Setting the firm's degree of initial knowledge (*KNOW*) and the size (*N*) and interdependence (*K*) of the performance landscapes allowed me to additionally tune the novelty and complexity of the environment. I then observed the firms' behavior for 1,000 periods, by which time all firms had reached a peak. The performance of each firm was measured relative to the global peak in each landscape, i.e., firm performance was 1.0 if the firm reached the global peak. In order to ensure that performance differences were inherent to the model and did not result from any stochastic effects, I repeated each experiment for 1,000 different landscapes and calculated the average performance for each type of firm across all landscapes. All reported performance differences were significant at the 0.001 level.

The results are ordered as follows: First, I report the core result of this study – that moderate levels of patience are beneficial, whereas too much patience will be futile and will even reduce performance – and describe the mechanisms that are driving it. I then show how this result is contingent upon the complexity and novelty of the firms' environment. The robustness of the findings is discussed at the end of this section.

4.1 Core result: The value of moderate patience

The firms that I consider first need to explore a completely novel (*KNOW* = 0) and moderately complex (*N* = 8, *K* = 4) environment, while differing in their patience (*PAT*) (see Figure 1).

< *Insert Figure 1 about here* >

Consider, at first, a firm that has only little patience (*PAT* = 1). Its performance improves very quickly and then stabilizes. A firm with a medium level of patience (*PAT* = 5 or *PAT* = 10) requires more time to improve, but eventually reaches a higher performance level. Finally, when the firm's patience is even larger (*PAT* = 100), performance improves even more slowly and eventually reaches a level that is *worse* than what resulted from a lower level of patience. While these findings clearly support the discussion about low levels of patience as summarized by Proposition 2, they seem to be ambivalent as far as higher levels of patience are concerned: Whereas Proposition 1a is supported for a medium level of patience, the findings given high patience favor Proposition 1b. To probe into the mechanisms that yield these performance implications, I repeated the same experiment for a broad range of the patience parameter. Figure 2 reports various measures that give insight into how different levels of patience affect the dynamics of search.

< Insert Figure 2 about here >

A firm with little patience engages in online evaluations, but returns to the best alternative that has been encountered so far after a single or a few trials that prove to be performance-decreasing. Hence, the firm never dares to move “far away” from its respective benchmark configuration. It will thus explore only a small, local fraction of the overall performance landscape (Figure 2, left chart) and often “move back” to its respective benchmark (Figure 2, middle chart). As the firm thus quickly creates knowledge about a local part of the landscape, offline evaluation sets in that allows the firm to make only performance-improving changes anymore, drawing the firm uphill on a nearby local peak. In other words, because the overall degree of exploration is lower for a firm that has little patience, its initial position strongly determines where it will finally end up. Nonetheless, however, the firm can efficiently exploit its local exploration of the landscape and quickly improve its performance.

For levels of moderate patience, the firm is willing to accept longer periods of underperformance, i.e., it will automatically “move” more intensively on the landscape, letting it explore more broadly (Figure 2, left chart). Given the higher degree of exploration, chances increase that the firm identifies superior alternatives, while the firm’s limited patience still frequently forces it to return to its respective benchmark configuration after a certain number of periods with below-benchmark performance has elapsed. At the same time, however, higher degrees of exploration also require more time to improve performance than in the case of very low levels of patience (Figure 2, right chart).

Yet when the level of patience is even higher, the results are intriguing. Although they yield even higher overall exploration of the landscape (Figure 2, left chart), they also make it very unlikely that the firm will return to a benchmark configuration (Figure 2, middle chart). Because high levels of patience make the firm highly exploration-seeking, it “drifts” around the landscape erratically, thereby getting exposed to a high variety of different alternatives. Thus, rather than focusing its exploration efforts, it chases after every idea that comes its way. When the firm eventually builds up a sufficient amount of local offline knowledge to start hill-climbing, it is at a rather "average" position and will consequently settle on a rather average local peak.[8]

Hence, the less patient a firm is, the more quickly it can improve its performance by creating local offline knowledge, whereas the more patient it is, the more broadly it will explore and the more slowly its performance will increase. However, to exploit the higher degree of exploration in the latter case, a moderate value of patience appears to be most efficient. What makes moderate patience valuable is that it helps strike a healthy balance between loosening and confining search efforts: On the one hand, tolerating underperformance for a moderate number of periods may allow a firm to move "downhill", entering a valley and reaching the "foothills" of a better peak, i.e., the firm may "escape" from the neighborhood of its current peak; on the other hand, the fact that patience is strictly limited forces the firm to still confine its exploration activities and create a critical amount of local (broadly defined) knowledge that can eventually be exploited, allowing the firm to move uphill. Very high levels of patience, in contrast, only further increase exploration in a time-consuming process, but yield no more gains. Rather, the firm fails to exploit the exploration that higher levels of patience entail and gets drawn away from potentially good configurations, only to end up on a rather average local peak.

4.2 The impact of complexity on the appropriate level of patience

Propositions 3 and 4, which are based on complex systems theory, state that a higher number of system elements as well as interdependencies between them will make higher levels of patience – by boosting exploration – more beneficial. Consider, first, the implications of interdependencies. Varying the degree of interdependence among the elements that the firm needs to explore, Figure 3 (right chart) finds partial support for Proposition 3: In the long run, higher degrees of interdependence do make patience more beneficial, yet medium rather than high levels of patience again yield the highest performance.[9] In the short run, in contrast, low levels of patience always prove benefi-

[8] This behavior reflects – in accordance with behavioral accounts of organizational decision making – that only repeated failure will force a firm to give up its current path and make the major "move back" to some configuration that is distant both in space and time. However, once a firm can determine by offline evaluation that its current position is a local peak, the high opportunity costs of giving up a current situation will rather lead to organizational inertia (Greve (2003)).

[9] The general performance decline for higher levels of interdependence results from the rising number of local peaks that tend to, on average, trap firms on configurations with lower performance (Kauffman (1995)).

cial, independent from the degree of interdependence, thus yielding additional support for Proposition 2 (Figure 3, left chart).[10]

< *Insert Figure 3 about here* >

If complexity is high and patience is rather low, a firm will explore only a confined neighborhood around its starting position: As complex environments contain many local optima (Kauffman, 1995), the firm will be likely to search in the neighborhood of an average local peak to which it will get drawn once offline search sets in. While this is detrimental in the long run (Figure 3, right chart), it helps the firm to quickly improve its performance in the short run (Figure 3, left chart). If, in contrast, patience is very high, chances are equally high that the firm will have – during its rather random drift across the landscape – explored the neighborhood of an average peak that intensively (creating that much offline knowledge) that it will get drawn to this peak. Instead, a moderate level of patience allows maximizing overall exploration by giving up less promising paths before getting drawn to a peak. For lower levels of complexity, the level of patience matters less, as the number of local peaks is likewise lower, and chances increase that a firm can reach an above-average peak by either a restrained neighborhood search or by broad exploration (which, however, requires more time than having only a low level of patience). Thus, the more complex the environment, the more the firm should try to find a balance between impatience and patience, whereas being impatient (or even overly patient) is less detrimental in a less complex environment.

Turning to the impact of system size, Figure 4 yields three main findings regarding the performance advantages of higher levels of patience over the exploitation-focused case of highly limited patience ($PAT = 1$): First, the "larger" a problem gets, the higher the optimal level of patience becomes; second, as systems get larger, the performance advantage of higher levels of patience grows; and third, medium levels of patience again yield the highest performance, whereas overly high levels of patience prove detrimental. While findings one and two thus support Proposition 4, it is not supported by the third finding, which again supports the robustness of my core result.

< *Insert Figure 4 about here* >

What drives this impact of system size? As a higher number of system elements increase the search space that needs to be explored, firms need to allow for a higher

[10] In Figure 3, I define "short run" as period 40, yet the results are insensitive to this particular choice.

level of patience to ensure that the exploration process is sustained for a sufficiently long time. Also, being patient proves to be more beneficial as larger systems contain more local peaks and prolonged exploration (higher patience) increases the chances to move towards one of the better ones. Nonetheless, independent from the size of the system, the detrimental "drift" effect of being overly patient eventually sets in. To summarize, then, larger systems increase the value of being patient as well as the notion of what denotes a "moderate" level of patience.

4.3 The impact of novelty on the appropriate level of patience

The experientially-searching firms above started their search without any offline knowledge. In many situations, however, a firm may have at least partial knowledge of a problem domain, e.g. when a technological shock has altered but not destroyed a firm's existing competencies in a particular domain. In this case, the firm may be able to assess certain alternatives in an offline manner once it has generated them. To gain further insight into the role of patience under these conditions, Figures 5 varies the firms' initial knowledge.

< Insert Figure 5 about here >

First, consider the left chart in Figure 5 that reports the resulting performance values given a short-term planning horizon. It clearly points to the value of (1) having offline knowledge or (2) gaining it quickly. In the short term, fast exploitation of existing competencies and the identification of "sufficient" solutions are more valuable than a longsome exploration for superior alternatives. This finding denotes a clear support for Proposition 5: A firm that possesses a high degree of offline knowledge can make many offline evaluations and improve its performance directly and, hence, quickly. A firm with less initial knowledge, in contrast, will have to search experientially first and will exhibit a lower performance initially because it will experience some failures. Yet although the performance discount for lower degrees of initial knowledge is severe, it can be mitigated significantly by very low levels of patience. Again, this supports Proposition 1: If a firm with little initial knowledge explores experientially but abandons its current search path after each (or only a low number of) period(s) of below-benchmark performance, it will quickly generate enough local knowledge to make hill-climbing improvements like a firm that started with high degree of domain-specific knowledge. Hence, the degree of initial knowledge need not be an obstacle for short-term exploitation if the firm chooses the right sampling strategy that supports a generation of offline knowledge that is geared toward the short term.

However, if the firm's planning horizon is geared towards the long run, the above wisdom is reversed (Figure 5, right chart). Now, broad exploration becomes valuable. Yet the higher the initial degree of knowledge, the less broadly a firm will explore as offline considerations will prevent it from implementing performance-decreasing alternatives.[11] However, as shown above, tolerating failures is necessary to leave one's local neighborhood and explore also more distant and possibly superior regions. A firm with less offline knowledge will (have to) make these mistakes and potentially be able to reap their benefits, whereas a firm with a higher degree of offline knowledge will more often be reluctant, and chances are higher that it remains in its initial region of the landscape. The lower a firm's degree of knowledge is, in contrast, the better it can make use of it to increase overall exploration, and the more the level of patience matters. For the reasons discussed above, a moderate level of patience, as well as starting over when exploration during this interval does not pan out, is especially valuable as it will both broaden a firm's degree of exploration and at the same time help the firm not to get stuck at lower local peaks when offline considerations finally take command. This finding partially supports Proposition 6 and yields additional support for my core result.

4.4 Robustness

As shown above, the core result is highly robust with regard to changes in the parameters problem complexity, i.e., size (*N*) and interdependence (*K*), as well as the degree of initial knowledge (*KNOW*). Here, I consider robustness with respect to a different parameter, a firm's search radius, which indicates how broadly the firm can search for new alternatives. Intuition suggests that if firms have a larger search radius, i.e., if managers face less severe cognitive bounds, they can come up with more innovative alternatives and will not benefit that much from patience to increase exploration. To test this idea, I considered firms that have a search radius of two (as compared to a search radius of one in Figure 1), i.e., that can identify new alternatives that differ in up to two decisions from their status quo.

[11] I assume that only experimenting with an unknown alternative allows firms to experience a performance drop. Firms do not move down-hill on purpose, i.e., they do not implement alternatives which they know would yield a worse performance than their status quo. As firms have no cognitive understanding of the landscape, firms thus behave (intentionally) rational.

< Insert Figure 6 about here >

As Figure 6 indicates, intuition is partially correct: A higher search radius decreases the performance differences between firms that have low levels of patience and those that have a moderate level of patience. Nonetheless, my core result – that medium levels of patience outperform both low levels and high levels of patience – proves robust again.[12]

5 Discussion

The issue of how alternatives are evaluated denotes a fundamental aspect of problem-solving search: Sometimes, a decision maker may evaluate an idea cognitively and free of risk; often, however, he may have to accept that he is "blind", having to implement an alternative in order to learn about its value. Hence, when firms need to solve complex and novel problems, for instance when exploring a novel technology or a strategy in a new business environment, they often cannot but take the risk of failures and search for a solution through trial-and-error experimentation. In consequence, a number of literatures – from engineering history (Vincenti (1990); Petroski (1992); Petroski (1994); Flyvbjerg et al. (2003)) to organizational research (Thornhill and Amit (2003); Starbuck and Farjoun (2005)) or psychological accounts of problem solving (Dörner (1997)) – yield plenty of cases that document failure that resulted from such online experiments.

Of course, as Levinthal (2005)) stresses, the dichotomy between offline and online modes of evaluation is not a strict one. Rather, there is a large gray area between the two poles in which most evaluation processes occur. Online trials need not require giving up the current alternative entirely. For instance, individual markets, pilot plants, or subsidiary firms may serve as a "test area" and provide the experiential basis on which the firm evaluates a new proposal. In aerospace or automotive design, for instance, wind tunnels are used to test the characteristics of new design alternatives, substituting for either the risky and costly (online) test of a full-fledged prototype or, on the other hand, the fully cognitive route of analytical analysis and computer simulations. As Levinthal (2005: 16)) further remarks: "In some sense, the issue of on- or off-line search becomes less a categorical distinction than a set of factors that influence the costs, risk, and possible accuracy of the evaluation process. Online search often entails a particular

[12] Also, the other measures introduced in Figure 2 remain qualitatively similar when the search radius is higher and the degree of patience is varied.

sort of cost, the opportunity cost of not making use of established options. [. . .] the degree to which current operations need [to] be disrupted by the need to evaluate a proposed alternative influences how painful that trade-off is." Hence, albeit only a stylized dichotomy, discriminating between cognitive and experiential evaluation remains helpful to gain insight into the dynamics of search. Whenever designers or managers need to solve complex and novel problems, the value of a particular search strategy depends crucially on which form of evaluation mechanism is possible.

Furthermore, failure and patience are closely linked to a fundamental strategic challenge faced by any firm: the trade-off between the (short-term and less risky) exploitation of old certainties, and the (long-term and more risky) exploration of new possibilities (Holland (1975); March (1991)). Here, patience has been shown as a mechanism that induces exploration by letting firms accept some performance-decreasing (down-hill) moves, thus contributing to overcoming the pitfalls of search that would otherwise be overly local. However, this does not guarantee effective search, as high levels of patience lead to excessive search, but yields only little gains, as "adaptive systems that engage in exploration to the exclusion of exploitation are likely to find that they suffer the costs of experimentation without gaining much of its benefit" (March (1991: p. 71)). However, firms may be able to balance exploration with exploitation by being determined to eventually abandon their current search path and to return to a previously encountered benchmark. Prior research has shown how organizational design may help firms to explore broadly but likewise stabilize around good decisions once they are discovered (Rivkin and Siggelkow (2003)). Here, I suggest that moderate levels of patience may have a similar function. On the one hand, they induce broad search, while on the other, they bound the search space such that "local" problem-specific knowledge can accumulate and eventually be exploited.

In many cases, however, firms will have too little rather than too much patience, which may confer a long-term disadvantage due to the risk of getting stuck prematurely on a low local peak. Consider the current case of the big car makers that have to restructure fundamentally in order to adapt to a changing business landscape. In searching for a new strategy that can restore their long-term viability, however, these firms may have to endure some significant short-term pain, requiring a large amount of patience in order not to get discouraged (Denrell & March, 2001). But because of the immediate performance feedback and public expectations to show good results quickly, these firms will most likely have less patience than might be necessary, a fact that can be most likely have less patience than might be necessary, a fact that can be considered to drive out exploration quickly and prematurely force the firms back into exploitation.

On the other hand, low levels of patience can also be beneficial. Even if a firm starts with little understanding of what works and what doesn't, and has only little time available, an adequate (i.e., low) level of patience can help it create local offline knowledge quickly and exploit its current position. For instance, a mechanism of this kind may be underlying concepts such as prototyping in software development (Thomke (1998); Loch et al. (2001)): Building (and potentially discarding) a high number of prototypes helps to quickly build up knowledge about the value of different alternatives, and hence contributes to rapid exploitation rather than longsome exploration.

Despite the abstract nature of the findings reported in this paper, they still have implications for managers that seek to affect their organization's level of patience. For examples, measures such as appropriate recruitment policies, the installation of test plants, or funding pet projects might help to increase the level of patience that a firm's decision makers are willing to exhibit, at least in those domains of the organization that are key to innovation. On the other hand, embedding constraints into critical search processes (e.g., milestones and review meetings in new product development projects) may serve to constrain the patience of organizational agents (e.g., of engineers) that might otherwise be tempted to explore overly broadly.

While verbal discussions already abstract from many real-world intricacies, formal modeling efforts go even further into this direction. This paper has been no exception, and various aspects deserve further attention. Here, I point to four potential avenues. One is cognition. While early writings in the tradition of the Carnegie School have offered local search by boundedly rational decision makers as an alternative draft to the fully rational optimizer of neoclassical theory, current research is starting to treat a middle ground (Gavetti (2005)). Hence, even when faced with novel and complex problems, initially "blind" problem solvers may apply cognitive devices to interpret the knowledge that is generated about the landscape (Farjoun (2008)). Gradually, they might form cognitive maps that provide a "big picture" of a problem domain (Gavetti and Levinthal (2000)) or help to identify a "preferred" direction (Winter et al. (2007); Nelson (2008)). With their growing understanding of the performance landscape, firms might then adjust their patience and search behavior accordingly. On the other hand, as Rosenberg (1995)) remarks, it is a central characteristic of innovation that even pioneers may lack vision. Further investigation of the role of intentionally rational behavior in problem solving thus appears to be promising.

Another direction for further research concerns the issue of delayed feedback that has become known as the "credit-assignment problem" (Holland (1998); Denrell et al. (2004)). Often, the performance implications of a particular alternative may not be observed instantaneously but only after some time has elapsed, during which the firm may have "moved on". Performance feedback can also be noisy. How should a firm then link current performance feedback to past actions? Also, performance feedback may be ambiguous, requiring multiple trials to establish a reliable link to performance. Furthermore, new alternatives may have both short-term and long-term effects. In evaluating a new drug, for instance, some effects may be assessable instantaneously, while knowledge about other effects can sometimes only be established through long-term studies (Nelson (2008)). Incorporating such considerations into models of search would make them more intricate, yet also more realistic.

A third limitation of this study pertains to the fact that it has been rather non-organizational. Most processes of problem-solving search and evaluation, however, occur in an organizational context that is characterized by a division of labor, by hierarchical relationships, and by various other formal and informal aspects of organizational reality (March and Simon (1958); Cyert and March (1963)). Despite a few exceptions (Rivkin and Siggelkow (2003); Siggelkow and Rivkin (2005)), models of adaptive search have largely shied away from applying an explicit organizational perspective. However, in order to bring our models closer to how organizations are actually evaluating alternatives, more work along these lines will be necessary (Gavetti et al. (2007)).

Finally, the paper has explicated why overly high levels of patience, despite inducing a significant amount of search, may have dysfunctional effects. Even under the "laboratory conditions" of the computational model developed above, i.e., in the absence of selection pressures or resource constraints, overly high levels of patience yield no further gains but rather reduce performance. Clearly, introducing a selection mechanism or equipping firms with exhaustible resources and cost considerations, all of which could be considered to dynamically affect the level of patience, might help shed light on further relevant aspects of patience in organizational search and denotes potential for further work.

6 Conclusion

Despite the above limitations, this paper has introduced an explicit notion of online experimentation into a model of search, shedding light on how patience is linked to innovation. I find that contrary to what intuition might suggest, high levels of patience are little desirable despite promoting a high level of exploration. Choosing a level of patience rather reflects a decision about how firms weigh exploration *and* exploitation. If managers want their organizations to innovate, they must embrace exploration of new possibilities – and be willing to tolerate failures that will inevitably occur along the way. At the same time, they also have to contain exploration, as it competes with the exploitation of the newly generated opportunities. If one seeks to boost innovation, achieving a healthy balance between the two becomes necessary. Under a robust set of assumptions, this can be a matter of moderate patience.

References

Allen, Thomas J. (1977), *Managing the Flow of Technology*, Cambridge, MA: MIT Press.

Baron, Jonathan (1988), *Thinking and Deciding*, Cambridge: Cambridge University Press.

Baumard, Philippe and William H. Starbuck (2005), Learning from failures: Why it may not happen, *Long Range Planning* 38, 281-298.

Biyalogorsky, Eyal, William Boulding and Richard Staelin (2006), Stuck in the past: Why managers persist with new product failures, *Journal of Marketing* 70, 108-121.

Bourgeois, L. J. and Kathleen M. Eisenhardt (1988), Strategic decision processes in high-velocity environments: Four cases in the microcomputer industry, *Management Science* 34, 816-835.

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.

Cannon, Mark D. and Amy C. Edmondson (2005), Failing to learn and learning to fail (intelligently): How great organizations put failure to work to innovate and improve, *Long Range Planning* 38, 299-319.

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.

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

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

Denrell, Jerker, Christina Fang and Daniel A. Levinthal (2004), From T-mazes to labyrinths: learning from model-based feedback, *Management Science* 50, 1366-1378.

Dörner, Dietrich (1997), *Die Logik des Mißlingens: Strategisches Denken in komplexen Situationen*, Reinbek: Rowohlt.

Eisenhardt, Kathleen M. (1989), Making fast strategic decisions in high-velocity environments, *Academy of Management Journal* 32, 543-576.

Ethiraj, Sendil K. and Daniel A. Levinthal (2004), Modularity and innovation in complex systems, *Management Science* 50, 159-173.

Farjoun, Moshe (2008), Strategy making, novelty, and analogical reasoning - commentary on Gavetti, Levinthal, and Rivkin (2005), *Strategic Management Journal* 29, 1001-1016.

Farson, Richard and Ralph Keyes (2002), *Whoever Makes the Most Mistakes Wins: The Paradox of Innovation*, New York: Free Press.

Fine, Charles H. (1998), *Clockspeed: Winning Industry Control in the Age of Temporary Advantage*, Reading, MA: Perseus Book.

Fleming, Lee (2001), Recombinant uncertainty in technological search, *Management Science* 47, 117-132.

Flyvbjerg, Bent, Nils Bruzelius and Werner Rothengatter (2003), *Megaprojects and Risk: An Anatomy of Ambition*, Cambridge: Cambridge University Press.

Forbes, Daniel P. (2005), Are some entrepreneurs more overconfident than others?, *Journal of Business Venturing* 20, 623-640.

Freud, Sigmund (1912), Formulierungen über zwei Prinzipien des psychischen Geschehens, in Sigmund Freud (ed.), *Gesammelte Werke*, London: Imago, 910-924.

Garud, Raghu and Andrew H. Van De Ven (1992), An empirical evaluation of the internal corporate venturing process, *Strategic Management Journal* 13, 93-109.
Gavetti, Giovanni and Daniel A. Levinthal (2000), Looking forward and looking backward: Cognitive and experiential search, *Administrative Science Quarterly* 45, 113-137.
Gavetti, Giovanni (2005), Cognition and hierarchy: Rethinking the microfoundations of capabilities' development, *Organization Science* 16, 599-617.
Gavetti, Giovanni, Daniel A. Levinthal and William Ocasio (2007), Neo-Carnegie: The Carnegie School's past, present, and reconstructing for the future, *Organization Science* 18, 523-536.
Gavetti, Giovanni and Jan W. Rivkin (2007), On the origin of strategy: Action and cognition over time, *Organization Science* 18, 420-439.
Gigerenzer, Gerd, Peter M. Todd and the ABC Research Group (1999), *Simple Heuristics That Make Us Smart*, New York: Oxford University Press.
Greve, Henrich R. (2003), *Organizational Learning from Performance Feedback: A Behavioral Perspective on Innovation and Change*, Cambridge, MA: Cambridge University Press.
Harrison, J. Richard, Zhiang Lin, Glenn R. Carroll and Kathleen M. Carley (2007), Simulation modeling in organizational and management research, *Academy of Management Review* 32, 1229-1245.
Hauschildt, Jürgen and Hans-Georg Gemünden (1999), *Promotoren - Champions der Innovation*, Wiesbaden: Gabler.
Holland, John H (1975), *Adaptation in Natural and Artificial Systems*, Ann Arbor, MI: University of Michigan Press.
Holland, John H. (1998), *Emergence: From Chaos to Order*, Reading, MA: Helix.
Israel, Paul (1998), *Edison: A Life of Invention*, New York: Wiley.
Kahneman, Daniel, Paul Slovic and Amos Tversky (1982), *Judgment under Uncertainty: Heuristics and Biases*, Cambridge: Cambridge University Press.
Katila, Riita and Gautam Ahuja (2002), Something old, something new: A longitudinal study of search behavior and new product development, *Academy of Management Journal* 45, 1183-1194.
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.
Lee, Fiona, Amy C. Edmondson, Stefan H. Thomke and Monica Worline (2004), The mixed effects of inconsistency on experimentation in organizations, *Organization Science* 15, 310-326.
Lenox, Michael J., Scott F. Rockart and Arie Y. Lewin (2006), Interdependency, competition, and the distribution of firm and industry profits, *Management Science* 52, 757-772.
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. (1997), Adaptation on rugged landscapes, *Management Science* 43, 934-950.
Levinthal, Daniel A. 2005. Footnotes to the behavioral theory of the firm: Wharton School, University of Pennsylvania.
Levitt, Barbara and James G. March (1988), Organizational learning, *Annual Review of Sociology* 14, 319-340.
Lippman, Stephen A. and John J. McCall (1976), The economics of job search: A survey, *Economic Inquiry* 14, 155-187.

Loch, Christoph H. and Christian Terwiesch (1998), Communication and uncertainty in concurrent engineering, *Management Science* 44, 1032-1048.
Loch, Christoph H., Christian Terwiesch and Stefan Thomke (2001), Parallel and sequential testing of design alternatives, *Management Science* 47, 663-678.
Lowe, Robert A. and Arvids A. Ziedonis (2006), Overoptimism and the performance of entrepreneurial firms, *Management Science* 52, 173-186.
Mach, Ernst (1905), *Erkenntnis und Irrtum*, Leipzig.
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.
Mihm, Jürgen, Christoph H. Loch and Arnd Huchzermeier (2003), Problem-solving oscillations in complex engineering projects, *Management Science* 49, 733-750.
Milgrom, Paul R. and John Roberts (1990), The economics of modern manufacturing: Technology, strategy, and organization, *American Economic Review* 80, 511-528.
Milgrom, Paul R. and John Roberts (1995), Complementarities and fit: Strategy, structure, and organizational change in manufacturing, *Journal of Accounting and Economics* 19, 179-208.
Nelson, Richard R. and Sidney G. Winter (1982), *An Evolutionary Theory of Economic Change*, Cambridge, MA: Harvard University Press.
Nelson, Richard R. (2008), Bounded rationality, cognitive maps, and trial and error learning, *Journal of Economic Behavior & Organization* 67, 78-89.
Petroski, Henry (1992), *To Engineer is Human: The Role of Failure in Successful Design*, New York: Vintage Books.
Petroski, Henry (1994), *Design Paradigms: Case Histories of Error and Judgment in Engineering*, Cambridge: Cambridge University Press.
Petroski, Henry (2006), *Success Through Failure: The Paradox of Design*, Princeton, NJ: Princeton University Press.
Popper, Karl (1959), *The Logic of Scientific Discovery*, New York: Harper.
Porter, Michael E. (1996), What is strategy?, *Harvard Business Review* 74, 61-78.
Ravasi, Davide and Carlo Turati (2005), Exploring entrepreneurial learning: A comparative study of technology development projects, *Journal of Business Venturing* 20, 137-164.
Rivkin, Jan W. (2000), Imitation of complex strategies, *Management Science* 46, 824-844.
Rivkin, Jan W. and Nicolaj Siggelkow (2003), Balancing search and stability: Interdependencies among elements of organizational design, *Management Science* 49, 290-311.
Rivkin, Jan W. and Nicolaj Siggelkow (2007), Patterned interaction in complex systems: Implications for exploration, *Management Science* 53, 1068-1085.
Rosenberg, Nathan (1995), Innovation's uncertain terrain, *McKinsey Quarterly* 3, 170-185.
Rosenkopf, Lori and Paul Almeida (2003), Overcoming local search through alliances and mobility, *Management Science* 49, 751-766.
Siggelkow, Nicolaj (2002), Evolution toward fit, *Administrative Science Quarterly* 47, 125-159.
Siggelkow, Nicolaj and Jan W. Rivkin (2005), Speed and search: Designing organizations for turbulence and complexity, *Organization Science* 16, 101-122.

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

Simon, Herbert A. (1956), Rational choice and the structure of the environment, *Psychological Review* 63, 129-138.

Simon, Herbert A. (1962), The architecture of complexity, *Proceedings of the American Philosophical Society* 106, 467-482.

Simon, Herbert A. (1996), *The Sciences of the Artificial*, Cambridge, MA: MIT Press.

Starbuck, William H. and Moshe Farjoun (Eds.) 2005. *Organization at the Limit: Lessons from the Columbia Disaster*. Malden, MA: Blackwell Publishing.

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

Sutton, Richard and Andrew Barto (1999), *Reinforcement Learning: An Introduction*, Cambridge, MA: MIT Press.

Thomke, Stefan and Takahiro Fujimoto (2000), The effect of 'front-loading' problem-solving on product development performance, *Journal of Product Innovation Management* 17, 128-142.

Thomke, Stefan H. (1998), Managing experimentation in the design of new products, *Management Science* 44, 743-762.

Thomke, Stefan H. (2003), *Experimentation Matters: Unlocking the Potential of New Technologies for Innovation*, Boston: Harvard Business School Press.

Thornhill, Stewart and Raphael Amit (2003), Learning about failure: Bankruptcy, firm age, and the resource-based view, *Organization Science* 14, 497-509.

Tucker, Anita L. and Amy C. Edmondson (2003), Why hospitals don't learn from failures: Organizational and psychological dynamics that inhibit system change, *California Management Review* 45, 55-72.

Välikangas, Liisa (2007), Rigidity, exploratory patience, and the ecological resilience of organizations, *Scandinavian Journal of Management* 23, 206-213.

Vincenti, Walter (1990), *What Engineers Know and How They Know It: Analytical Studies from Aeronautical History*, Baltimore: The Johns Hopkins University Press.

Winter, Sidney G., Gino Cattani and Alex Dorsch (2007), The value of moderate obsession: Insight from a new model of organizational search, *Organization Science* 18, 403-419.

Figure 1: Performance implications of different levels of patience

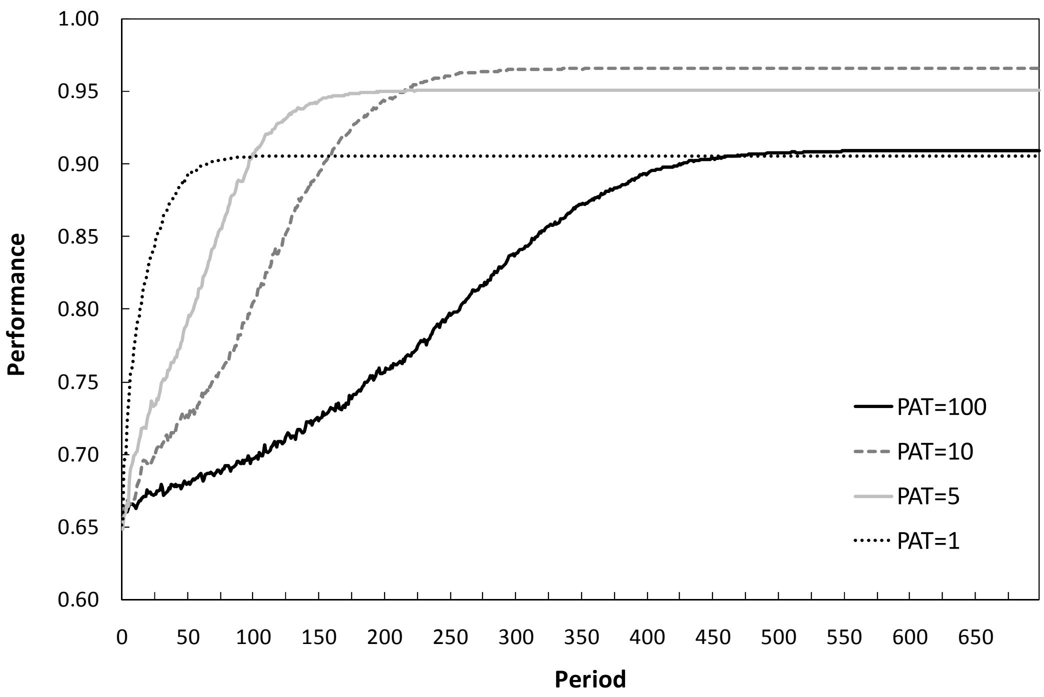

This figure reports the average firm performance over 1,000 landscapes with N = 8, K = 4. Firms differ in their level of patience (PAT). All firms start their search without any offline knowledge ($KNOW$ = 0).

Figure 2: Impact of different levels of patience on the dynamics of search

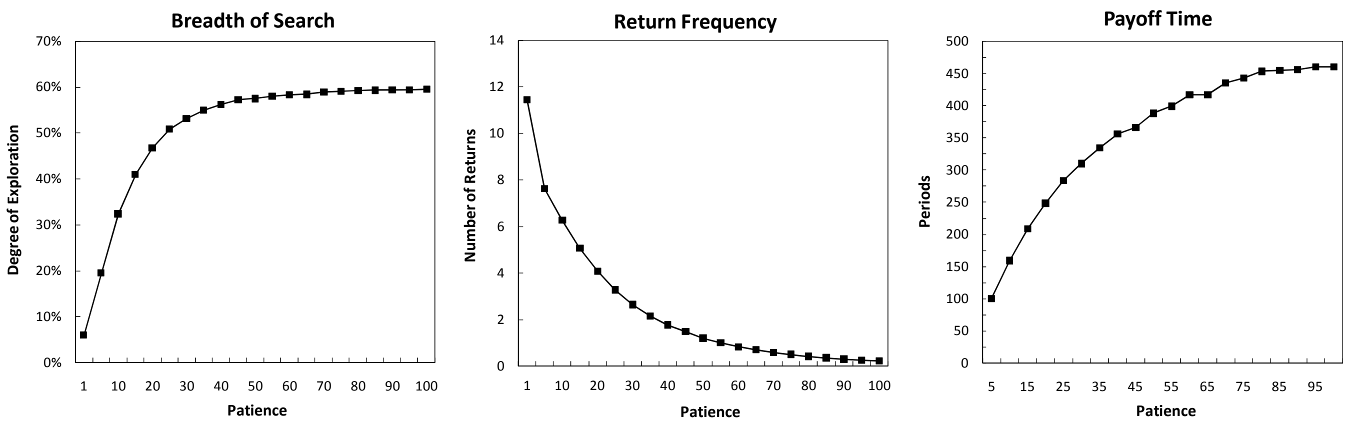

The chart on the left reports the fraction of all different performance contributions (c_i) that a firm assesses during its search. The chart in the middle reports the number of times that firms give up their search path and "start over" after persistent underperformance. The chart on the right illustrates the number of periods that firms with a higher level of patience (PAT) require to outperform a firm with a low level of patience (PAT = 1). All results are averages over 1,000 landscapes with N = 8, K = 4. Firms search for 1,000 periods. They differ in their level of patience (PAT) and start their search without any offline knowledge ($KNOW$ = 0).

Figure 3: Short-term and long-term effects of different levels of interdependence and patience

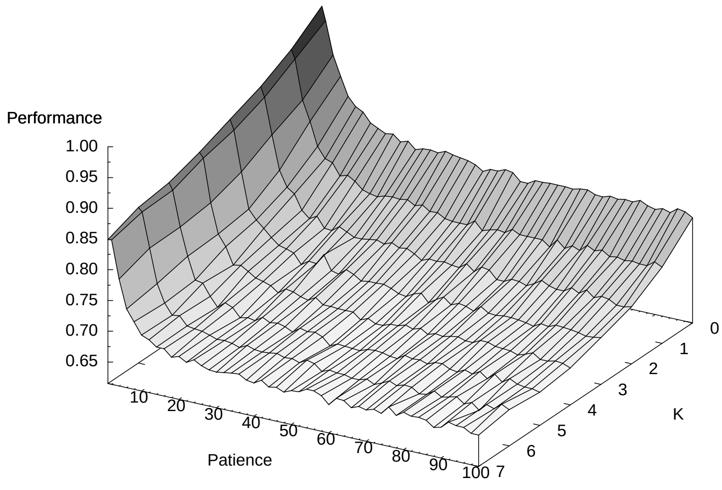

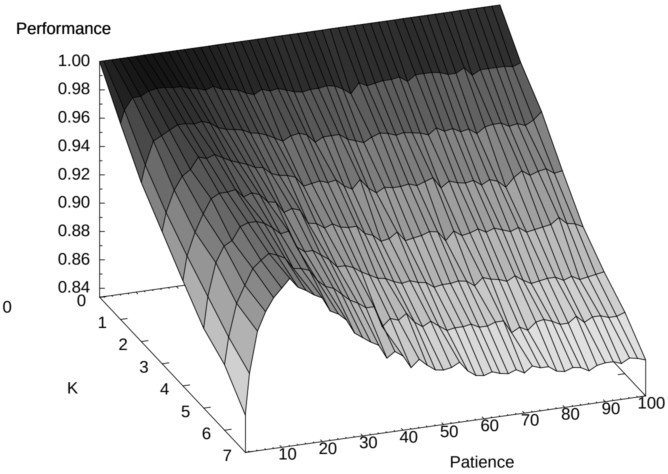

This figure reports the average firm performance in period 1,000 over 1,000 landscapes with $N = 8$ and different degrees of complexity ($0 < K < 7$). Firms differ in their degree of patience (*PAT*) and start their search without any offline knowledge ($KNOW = 0$). The left chart reports the performance in the short run (period 40). The right chart reports long-run performance (period 1,000).

Figure 4: Effects of different problem dimensions

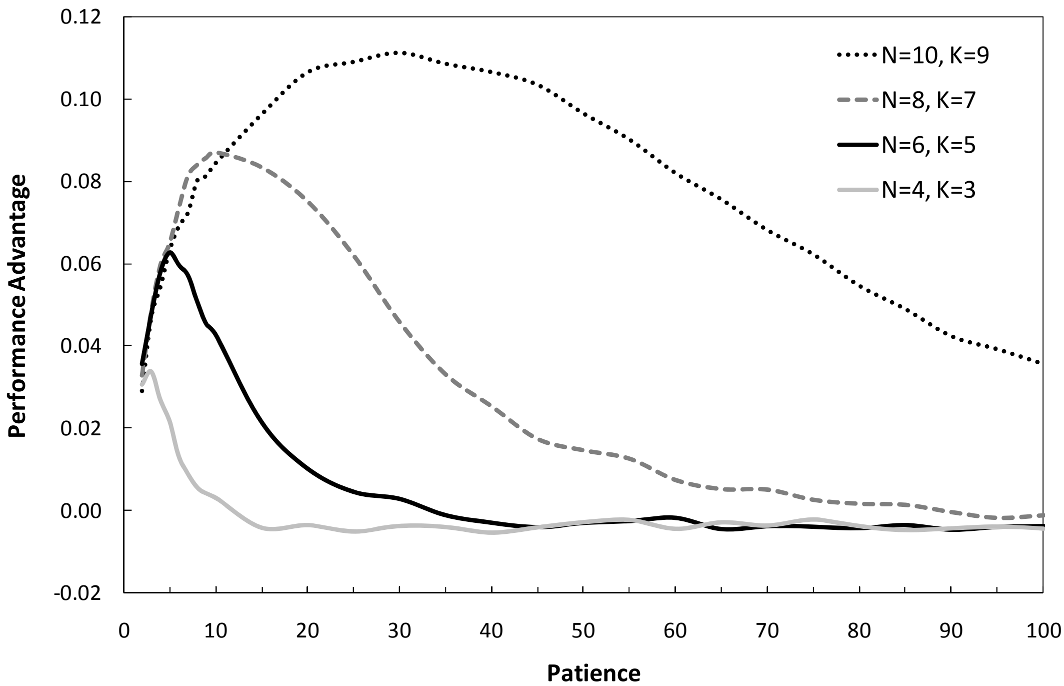

This figure reports the average performance difference between firms with higher levels of patience ($PAT > 1$) and firms with a very low level of patience ($PAT = 1$) over 1,000 landscapes. Landscapes differ in their size (*N*), while the degree of interdependence is set to $K = N$-1. All firms start their search without any offline knowledge ($KNOW = 0$).

Figure 5: Short-term and long-term effects of different levels of initial knowledge and patience

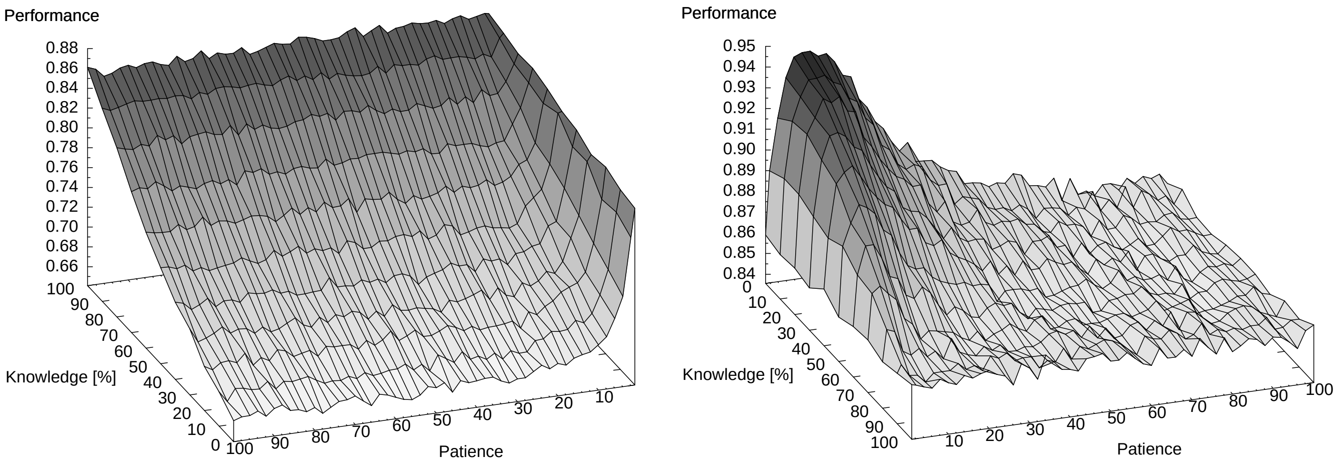

This figure reports the average firm performance over 1,000 landscapes with N = 8 and *K* = 7. Firms differ in their level of patience (*PAT*) and their degree of initial knowledge in the problem domain (0 ≤ KNOW ≤ 1). The left chart reports the performance in the short run (period 40). The right chart reports long-run performance (period 1,000).

Figure 6: Performance implications of different levels of patience given a higher search radius

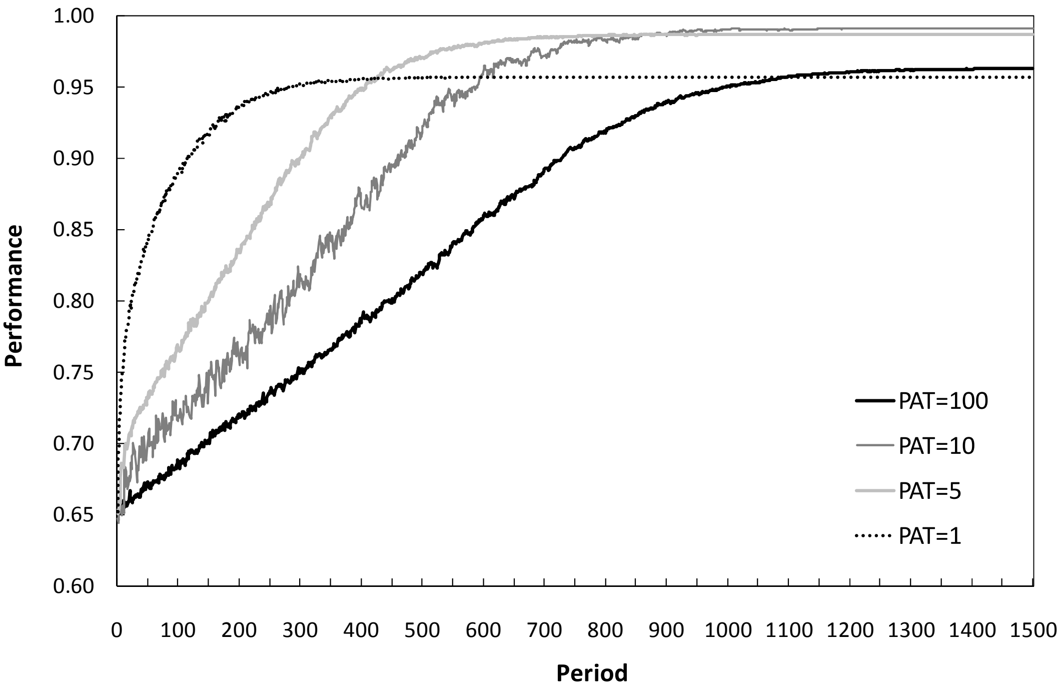

This figure reports the average firm performance over 1,000 landscapes with *N* = 8, *K* = 4. Firms differ in their level of patience (*PAT*). All firms start their search without any offline knowledge (*KNOW* = 0). Firms are less bounded in their rationality and have a search radius of 2, i.e., they can generate alternatives that differ in up to two decisions from their status quo set of choices.