University of Zurich[UZH]

Department of Business Administration

UZH Business Working Paper Series

(ISSN 2296-0422)

Working Paper No. 312

Learning Through Inaccurate Replication

Helmut Dietl, Markus Lang, Eric Lucas, Dirk Martignoni

7 December 2012

University of Zurich, Plattenstrasse 14, CH-8053 Zurich,
http://www.business.uzh.ch/forschung/wps.html

UZH Business Working Paper Series

Contact Details

Helmut M. Dietl

University of Zurich

helmut.dietl@business.uzh.ch

Markus Lang

University of Zurich

helmut.dietl@business.uzh.ch

Eric Lucas

University of Zurich

eric.lucas@business.uzh.ch

Dirk Martignoni

University of Zurich

dirk.martignoni@business.uzh.ch

Learning Through Inaccurate Replication*

Helmut Dietl, Markus Lang, Eric Lucas, Dirk Martignoni **

University of Zurich

7 December 2012

Abstract

Replicating a successful "template" or best practice across time ("temporal replication") or across a number of different economic settings ("spatial replication") is an important strategy for organizational growth and performance improvement. In this paper, we use an NK landscape model to examine how organizations may innovate and adapt to their environment through "inaccurate replication". We identify conditions under which inaccurate replication can result in higher long-run performance than accurate replication. We also uncover the specific mechanisms through which replication errors may affect organizational performance. In that, our study also sheds a new light on how organizations may learn from (replication) errors.

Keywords: Replication, Learning, Errors, NK Simulation Model

Very Preliminary

- Please do not cite or distribute without permission -

* Earlier versions of this paper were presented at the OLKC 2012 and DRUID 2012. We would like to thank conference participants for helpful comments and suggestions. We also gratefully acknowledge the comments by Jan Rivkin, Hart Posen, Felipe Csaszar, Torben Pedersen, Nils Stieglitz, and Wolfgang Guettel. Financial support was provided by the Ecoscientia Foundation and the Foundation for the Advancement of Young Scientists (FAN) of the Zürcher Universitätsverein (ZUNIV).

** All the authors from the Department of Business Administration, University of Zurich, Plattenstrasse 14, 8032, Zurich, Switzerland. Emails: helmut.dietl@business.uzh.ch, markus.lang@business.uzh.ch, dirk.martignoni@business.uzh.ch, eric.lucas@uzh.ch

1. Introduction

There is much consensus that the replication[1] of organizational routines, business models, and best practices is an important means of both organizational growth and performance improvement (Nelson and Winter 1982; Teece et al. 1997; Argote and Ingram 2000; Rivkin 2000; Winter and Szulanski 2001; Winter et al. 2011). While there is also much consensus that accurate replication is hard to impossible to achieve (Szulanski and Winter 2002; Rerup 2004), there is much less consensus on the performance implications of inaccuracies in the replication process.[2] In our study, we seek to identify conditions under which these inaccuracies are harmful or may even prove beneficial.

There is a large yet inconsistent body of research on the performance implications of replication errors. On the one hand, there is considerable empirical evidence that firms should seek to replicate a successful practice as accurate as possible. For example, Winter et al. (2011) found that inaccurate replication has, on average, negative performance consequences. Obviously, if the template to be replicated is literally a best practice (i.e. there is no further room for improvement), any replication errors will be dysfunctional (Rivkin 2001). Moreover, firms such as Intel or McDonald's often put a very strong emphasis on replicating a successful practice at a number of different sites as completely and accurately as possible (Reinhardt 1997; Iansiti 1998; McDonald 1998; Szulanski and Winter 2002). On the other hand, primarily theoretical arguments suggest that there should be also benefits to inaccurate replication. For example, Becker et al. (2006) portray replication errors as a source of novelty and variation. Szulanski and Winter (2002) and Winter (2005) argue that perfect or accurate replication impedes any adaptation and improvement and thus might be costly, in particular, if adaptation to the local context is required (Teece 1997) or the current practice is suboptimal (Winter et al. 2007). Indeed, there seems to be a trade-off between accuracy and learning associated with the replication process, a trade-off that is often labeled as "replication dilemma" (Szulanski and Amin 2001; Winter and Szulanski 2001; Szulanski and Jensen 2004).

In our study, we will rely on the apparatus of the NK performance landscape model (Levinthal 1997; Rivkin 2001) in our attempt to reconcile these conflicting views on the implications of inaccuracies in the replication processes. Specifically, we are interested in how the performance consequences of replication errors might be moderated by the size of the replication errors, the time horizon, and the complexity of the organization's environment. Although firms may also deviate from an existing

---

[1] Replication may be best understood as an attempt to reproduce an existing (successful) practice at a number of different locations (Nelson and Winter 1982).

[2] Although the terms "imperfect replication", "inaccurate replication", "imprecise replication", "incomplete replication", or "partial replication" may carry different connotations, technically, however, they are very similar. Incomplete or partial replication is a subset of inaccurate or imperfect replication: with incomplete replication, some aspects of a practice may or may not be replicated correctly. Those aspects that are not replicated are recreated correctly or incorrectly. For more details, see also Section 3.3.

template to achieve a better adaptation to the local environment (Winter 2005; Winter et al. 2011) and to secure commitment and cooperation at the recipient site of the replication template (Szulanski 2000), these aspects are beyond the scope of our study. We are only interested in whether and how deviations from an existing (and at least locally optimal) template can induce innovation and, in turn, may improve firm performance.

The central argument of our study is that in the short-run, all replication errors have negative performance consequences if the replication template is a local or global optimal; in the long-run, in contrast, small replication errors may even enhance performance, in particular in environments of moderate complexity. Surprisingly, even large replication errors have no effect on long-run performance in highly complex environments, while they have negative consequences in low to moderately complex environments. A necessary condition for any performance improvement is that a replication error makes a firm *abandon* its current replication template and *discover* a different and hopefully better practice. Interestingly, however, an inaccurate replica is a worse "starting position" than any random position: the practices discovered starting from an inaccurate replica are, on average, worse performing than those discovered starting from any random position ("de novo reinvention", see Winter 1995). As we will demonstrate in our study, the positive performance effect of small replication errors is driven primarily by the fact that they make firms abandon particularly low-performing replication templates.

Given the importance of replication processes, it is not surprising that we are not the first to address the question of the implications of incomplete or inaccurate replication. Some studies have also identified a positive effect of small replication errors on long-run performance. However, our paper differs from these contributions in several ways.

First, those studies that report a positive effect of errors in replication and imitation often identify these positive effects on the *population level* (Holland 1975; Aldrich 1979) or assume within population-learning (Csaszar and Siggelkow 2010).[3] Imitation or replication errors introduce variation, which in turn, is the raw material for successful population adaptation (Holland 1975). In our study, in contrast, we are interested in the effect on the firm level. This is an important difference because errors or mutations may have a favorable effect for the population of firms (if they can learn from each other) but an adverse effect on the firm that suffers from these errors (March 1991; Csaszar and Siggelkow 2010).

Second, those studies (e.g., Rivkin 2001) that identify a consistent negative performance effect of all types of replication errors (even small replication errors) often assume that the template to be

[3] Besides imitation, Csaszar and Siggelkow (2010) also analyze the performance implications of random long jumps. They find that random long jumps nearly always lead to a decrease in short- and long-run performance.

replicated is literally a "best practice", i.e. it cannot be improved anymore; it is already a "global peak" in a performance landscape. By definition, there is no upside to replication errors; one cannot improve beyond the global peak. In our study, in contrast, we relax this assumption and also examine cases in which the practice to be replicated is a good practice (local peak) but not always a best practice (global peak).

Third, those studies that highlight the positive performance implications of deviations from a replication template often attribute these positive effects to adaptations to the local context (Williams 2007). In our study, in contrast, we assume that the source of the replication template and the recipient are operating in the same context.

Fourth, we adopt a dynamic perspective of organizational search and adaptation in which adaptive processes lead to improvements over time but also often fail to achieve global maxima on a complex[4] performance landscape (Nelson and Winter 1982; Becker et al. 2006). Both the source and recipient organization seek to improve the template through these search and adaptation processes. The inaccurately replicated practice only serves as a starting point for subsequent adaptation processes. Studies with a static perspective often focus only on the immediate consequences of replication errors. Yet, as we demonstrate in our study, these immediate performance consequences might be quite different from the long-run performance consequences.

Fifth, while our study focuses on the effects of replication errors on the search and adaptation processes, contributions with perspectives like institutional theory (Meyer and Rowan 1977; DiMaggio and Powell 1983) would highlight the (negative) effects of inaccurate replication on the acquisition of acceptance and legitimacy. These effects, however, are beyond the scope of our study.

Finally, Knudsen and Levinthal (2007) demonstrate in the context of the NK landscape model that inaccurately evaluating the performance associated with a practice can increase performance. In other words, their study (and related work in the context of experiential learning models, such as Denrell and March 2001) focuses primarily on the impact of errors in performance information while our study is concerned with the errors in choice information.[5]

These different conceptual foci and scopes may also explain why existing research on the performance implication of replication errors has generated mixed results. The primary contribution of our paper is that we seek to consolidate our knowledge about the performance implications of replication errors by means of series of systematic experiments with a well-established simulation

[4] Simon (1962) defines complexity as having two aspects: An item is complex if it consists of many elements and those elements interact richly.

[5] Transferred to our context, these studies examine the impact of errors in assessing the performance of the current position on the landscape incorrectly.

model. On an abstract level, the purpose of our study is twofold. First, we will seek to identify conditions under which replication errors may prove functional and dysfunctional. Prior research has generated inconsistent findings on this question. By taking into account moderators that have been neglected so far (e.g., size of the replication error, complexity, and time horizon), we can reconcile these inconsistencies. Second, we will try to uncover the specific mechanisms through which replication errors may affect organizational performance, thereby shedding new light on the interplay between learning and errors. Specifically, we show that the positive long-run performance effects of small, unintended replication errors only materialize if they are hard to detect and correct. "Intended" errors (e.g., global search or long jumps) or deviations that are easier to correct, in contrast, have lower positive long-run performance effects.

The remainder of the paper is structured as follows. In Section 2, we briefly review the existing literature with respect to the theoretical building blocks of our model. Section 3 introduces our extensions of the standard NK landscape model (Levinthal 1997). In Section 4, we systematically analyze the performance implications of replication errors and identify the boundary conditions of our findings. Finally, in Section 5 we conclude by discussing our results and implications.

2. Literature Review

In this section, we briefly outline the theoretical buildings blocks of our model and review the corresponding literature. We organize this literature into five broad themes: first, we discuss how replication can create any value. Second, we focus on the particular value of accurate replication. In the third subsection, we provide some evidence for the prevalence of replication errors. In the final two subsections, we discuss how errors may become a source of innovation and learning and how our study relates to the broader literature on errors and learning.

2.1 Benefits of Replication

Replication is an important strategy to both growth and performance improvement (Winter and Szulanski 2001). Organizations may leverage their knowledge assets by replicating successful practices at different sites, thereby avoiding costly and risky de novo reinvention of those practices (Nelson and Winter 1982). Ideally, replication efforts at new sites generate economic benefits that are comparable to those achieved at the source or template site (Winter 2010). Some organizations such as Intel with its "Copy Exactly" strategy for ramping up semiconductor production are able to reap considerable economic benefits from replicating the manufacturing process developed at a lead factor as accurately as possible at other production sites (Reinhardt 1997; Iansiti 1998; McDonald 1998; Szulanski and Winter 2002). From the start on, production yields almost immediately matched those of the lead factory that was the source of the replication template. With traditional approaches of repeating the learning process at each site ("re-invention"), it could have taken several months to match the efficiency of the lead factory.

Re-invention also almost always involves higher costs; investments in creating the replication template at the lead factory are basically sunk costs and the costs of copying it subsequently at different sites are decreasing (Winter 1995). In addition, as argued by Levinthal (1997) and Winter (1995), as a result of the path and context-dependent character of learning, there is no guarantee that re-invention will lead to the same outcome; there is always some uncertainty regarding the outcome and costs of this learning process (Winter 1995): "These latter uncertainties imply that, quite apart from the costs of reinvention de novo, the economic considerations that would typically motivate replication may simply not apply to reinvention. The profitability and quality performance of the existing routine are not likely to be predictive of the reinvented one; and the latter is likely to require accommodating changes in complementary routines that would not be needed under replication" (p. 156).

In addition to these direct positive economic benefits of replicating superior knowledge, adopting the same practice across the entire organization also increases consistency. As Dave Van Lear, CEO of BancOne Services Corporation, explains: "Having the banks on one set of systems, with standard operations, will tremendously increase our affiliates' ability to provide a consistent level of high quality service. It also means we can make changes quickly, introduce new products, and operate at lower cost" (Szulanski 2000: p. 90). These indirect effects, however, are beyond the focus of our study; we are only interested in the direct effects of replication for the recipient organization.

In sum, by replicating a successful practice, firms attempt to avoid costly and uncertain re-invention. Through replication, firms also can scale up their operations faster than by repeating the learning process.

2.2 Close and Accurate Replication

The direct economic benefits of replication seem to be particular pronounced if replication is as accurate as possible. Winter (2005), Winter and Szulanski (2001), and Szulanski and Jensen (2006) caution against any deviations from a template: any deviation might be costly and may have unintended consequences and the presence of interaction effects may even amplify these negative consequences. Thus, Winter and Szulanski (2001: p.248) recommend that organizations should keep "the list (of deviations) as short as possible." The reason is that "modifications introduced to adapt the established template may create new problems; problems that will have to be solved in-situ through a costly process of trial and error" (Winter and Szuslanksi 2001: p.737). Organizations such as Intel with highly complex manufacturing processes often follow a "Copy Exactly" strategy. They are investing considerable resources and effort to copy every detail of the manufacturing processes established in one lead factory to other manufacturing sites (Iansiti 1998: Chapter 8), often even replicating aspects that seem to be not contributing to the success of a practice. For example, at Intel, they even replicated the color of the paint of the lead factory (Reinhardt 1997). Managers at Intel even joked that “the height of process technicians must be identical in all fabs” (Iansiti 1998: p.168). There seems to be some truth in this joke: (Szulanski 2012) describes Intel’s replication process in the following way: “They copied everything, including a door in the new factory, even if it leads to nowhere (in the original design the door had led to another building)”.

In sum, accurate replication seems particularly important in complex environments; any deviation from the replication template will come at a cost. Specifically, unintended deviations in the core elements that contribute to the success of a template may cause severe problems. Yet, isolating these core elements is often difficult and therefore firms rather replicate everything, even elements that might not be important.

2.3 Inaccurate Replication is the Rule

While there seems to be strong evidence for the benefits of replication, perfect or accurate replication is rarely observed in reality. Indeed, as observed by Nelson and Winter (1982: p.121), “perfect replication is not more of an ultimate objective than perfect control.” One reason why we don’t observe accurate replication in reality is that it is hard to impossible to accomplish: “even if exact replication is aimed at, it will never be fully achieved in practice” (Winter 1995). Accurate replication is often hard to achieve because (i) knowledge embedded in the template is often tacit and cannot be easily transferred (Nelson and Winter 1982; Kogut and Zander 1992; Szulanski 2003), and (ii) it is not even clear what constitutes the elements of success of a template (“causal ambiguity”, see Lippman and Rumelt 1982).

Another reason why we don’t observe accurate replication in reality is that even if accurate replication were possible, the objective is often only partial rather than accurate replication (Winter 1995). One reason is that accurate replication is often too costly, in particular if it is not clear what constitutes the elements and boundaries of a successful practice. Organizations may also only seek partial replication because by deviating from the original template (i.e., replicating it inaccurately), organizations may achieve better local adaptation (Csaszar and Siggelkow 2010, Winter et al. 2010). Different contexts may have different optimal solutions and, by implication, using the solution that is optimal in one context doesn’t guarantee high performance in another context, in particular, if the contexts are quite different (Kaufmann and Eroglu 1999). Sticking closely to a template inhibits local adaptation (Bartlett and Ghoshal 1999). Another reason why firms may seek only inaccurate replication is that it helps to secure commitment and cooperation (Szulanski 2000). The recipient might be reluctant to adopt an externally developed practice (“not-invented-here syndrome”, see Katz and Allen 1982).

Szulanski (2000) describes that replicating a common practice across BancOne was difficult because it required some organizational units to discard their own superior products or systems. These aspects of local adaptation and commitment (El Akremi et al. 2011), however, are beyond the scope of our study.

In sum, in reality, inaccurate replication is the rule rather than the exception. Accurate replication may not even be the goal in all instances.

2.4 Innovation and Learning Through Inaccurate Replication

In our study, we are particularly interested in how inaccurate replication may affect an organization's ability to innovate and learn. If a practice is always replicated accurately (both temporarily and spatially), reliability in performance is high but, at the same time, the practice does not change and improve. At Intel, the copy-exactly strategy ensures that the manufacturing processes have not to be re-invented again at different production sites. Yet, as Intel managers also observed (Iansiti 1998), while it may guarantee high performance from the start on, it inhibits further improvements later on. When replicating a successful manufacturing process, Intel is often replicating old and outdated technology (Terwiesch and Xu 2004). As Bernie Wood, director of marketing at an equipment vendor to Intel, explains "They [Intel] use some of the oldest equipment in the business, some of the most outdated equipment, because of their copy-exactly requirements" (Dorsch 2000).

Inaccurate replication, in contrast, may become a source of innovation. Love and Miller (1995) describes how, in the inter-organizational context, such imitation errors lead to innovation in the fast food industry. In the intra-organizational context, Szulanksi (2001) and Winter and Szulanksi (2001) discuss the paradox of accuracy in the replication process ("replication dilemma"). One the one hand, by accurately replicating a good (or even a best) practice, firms can realize the economic benefits in the recipient organization that are comparable to those already achieved in the organization that created the replication template (Winter 2010). On the other hand, accurate replication also impedes any learning and improvement of a firm's practice (Winter 1975; Winter 2005). POrganizations may innovate by imperfect attempts to replicate. Similarly, but invoking a different label ("productivity dilemma"), Abernathy (1978) and Benner and Tushman (2003) argue that if firms focus too much on adhering and controlling the executing of their current practice, this inhibits their flexibility and ability to innovate.

While the replication dilemma is akin to the trade-off between exploration and exploitation (March 1991), there are some fundamental differences. Most importantly, exploration is often associated with "distant search" or "global search" (Kauffman 1993; Levinthal 1997), i.e. deliberately considering alternatives that are "not local in a temporal sense or with respect to the organization's existing capabilities" (Levinthal 2008: p.95). The focus of our study is different: we are interested in *unintentional* deviations from an existing practice, i.e. when seeking to replicate an existing template, a firm may unwittingly get some aspects of the template wrong. Firms may also intentionally deviate from the replication template to achieve better fit to its local context (Winter et al. 2011). Yet, these cases are beyond the scope of our study.

We are interested in only those cases in which the recipient organization cannot easily detect in which aspects it may or may not deviate from the replication template. A showcase of the kind of unintentional deviation we have in mind is recounted by Polanyi (1959): "I have myself watched in Hungary a new, imported machine for blowing electric lamp bulbs, the exact counterpart of which was operating successfully in Germany, failing for a whole year to produce a single flawless bulb." (p. 52). Clearly, if the reason had been obvious to the Hungarian operators, it would have been much easier to correct. Without this knowledge, however, correcting unintentional deviations (to the way the machine was operated in Germany) was very difficult.[6]

In sum, in order to further improve a practice, a firm must deviate from it. Without deviations, there is no change. These deviations can be intentional or unintentional. We are particularly interested in the effect of unintentional deviations.

2.5 Errors and Learning

What do we mean by "inaccurate replication as a source of innovation and learning"? This concept may sound similar to concepts like "learning from errors" (Argyris 1982; Edmondson 1996), but there are important differences: the basic idea behind "learning from errors" is error minimization or prevention. By enhancing the knowledge about what went wrong, this error may be avoided in the future. In our case, avoiding inaccuracies in replication cannot lead to any change or innovation. Indeed, the opposite is true: in the context of our study, this kind of learning would suppress any deviation from the replication template and, as a result, decrease the probability of innovation. Although it resembles trial-and-error learning, there is an important difference: with trial-and-error learning, experimentation, or deliberate perturbations, one deviates intentionally from the current template. By definition, the aspects in which these deviations occur are known; if they don't proof performance enhancing, they can be reverted back to their original values. The concept of inaccurate replication, however, implies that this information is not available.

Inaccurate replication might be best understood as mutation of the replication template. As described by Nelson and Winter (1982) "the copy is, at best, likely to constitute a substantial mutation of the original" (p. 123). To the organization, it is not clear which elements of the replication template are

[6] For another illustrative example of unintended and hard to correct errors, see HBS Case Corning Glass Works: The Z-Glass Project

mutated. Depending on the complexity of the organization's environment, it can be quite tedious to discover the mutations as illustrated by Polanyi's light bulb example. After some time, the Hungarian operators learned how to operate the light bulb machine; in more technical terms, they were able to revert mutations of the replication template. Learning from inaccurate replication, however, would have implied something like discovering a better way to run the light bulb machine. Thus, in this aspect, learning from inaccurate replication, is closely related to concepts like "accidental discovery" (Campbell 1965).

3. Model

Like the standard NK performance landscape model (Levinthal 1997), our version of the NK landscape model has three basic features: (1) a complex performance landscape, (2) firms that are represented by a position on this performance landscape, and (3) a process of local search through which firms improve their position on the performance landscape.[7] The (complex) performance landscape is a mapping of firm decisions to performance. A firm is associated with a specific decision vector in a given period. Firms seek to continuously improve their positions on the landscape through a process of local search. In our study, we relax the implicit assumption of many NK models that accurate temporal and spatial replication is always possible. In the following subsections, we provide detailed descriptions of the elements and processes of our model.

3.1 Complex Performance Landscapes and their Properties

The starting point of our model is an N-dimensional vector $\boldsymbol{a}=(a_1, a_2 ,.., a_N)$ of binary decisions $a_i \in \{0,1\}$ with $i \in I=\{1,..,N\}$, yielding a total of 2^N possible combinations of choices. In our model, we interpret the vector $\boldsymbol{a}$ as representing a firm or, more precisely, a firm's configuration of decisions. Some of these decisions may be interdependent and others may not.

The degree of interdependence among a firm's decisions is determined by the parameter $K \in \{0,..,N-1\}$, which describes the number of decisions a_j that (co-)determine the performance effect of decision a_i.[8] This effect is characterized by the contribution function $c_i = c_i (a_i, a_{i_1}, a_{i_2},..,a_{i_K})$, where $i_1,i_2...,i_K$ are K distinct decisions other than i. The realizations of the contribution function are drawn from a uniform distribution over the unit interval, i.e., $c_i \sim U[0;1]$. The performance of a given decision vector $\boldsymbol{a}$ is calculated as the arithmetic mean of the N contributions c_i according to the performance function

---

[7] The NK model has been applied to a broad range of topics such as organizational development and change (Ruef 1997), innovation (Frenken 2000; Fleming and Sorenson 2001; Almirall and Casadesus-Masanell 2010), organizational design (Rivkin and Siggelkow 2003; Siggelkow and Levinthal 2003; Gavetti 2005), and strategy (Siggelkow and Rivkin 2005; Levinthal and Posen 2007; Csaszar and Siggelkow 2010). Porter and Siggelkow (2008) provide a comprehensive overview on NK models.

[8] The interactions can be represented in a so-called interaction matrix (see Section 4.6.3 for more detail),

$$\phi(a)=\frac{1}{N}\sum_{i=1}^{N}c_i(a).$$

A "landscape" represents a mapping from all 2^N possible outcomes of the decision vector onto performance values. We normalize each landscape to the unit interval such that the mean value of the normalized landscape equals 0.5 and the global maximum equals 1.0. The "local peaks" on the performance landscape represent decision vectors for which a firm cannot improve its performance through local search (Ganco and Hoetker 2009: p.8). The "global peak" is the highest of all local peaks in the landscape.

The parameter K is commonly interpreted as a measure for complexity (Rivkin and Siggelkow 2003; Siggelkow and Rivkin 2005). The lowest value K=0 means that the decisions do not depend on each other yielding a smooth performance landscape with a single global peak; the highest value K=N-1 characterizes a situation, in which each decision depends on all other decisions. In NK models, the number of local peaks is increasing with the complexity of the performance landscape, i.e., the parameter K (Kauffman 1993). In our model, the number of local peaks increases more or less linearly from 1 (K=0) up to 95 (K=9) for N=10. At the same time as the parameter K increases, the "peaks on the landscape spread apart" (Rivkin 2000: p.838) and the performance of the peaks "dwindles" to the average performance in the landscape. Kauffman (1993) refers to this phenomenon as the "complexity catastrophe."

Moreover, as prior research on NK models documented, there is a positive correlation between the performance of a local peak and the size of its basin of attraction (Kauffmann 1993; Levinthal 1997; Rivkin 2000; Gavetti and Levinthal 2008; Siggelkow and Levinthal 2003). According to Kauffman (1993) the basin of attraction of a certain local peak can be considered as the set of points in the landscape for which local search leads to this local peak. Low-performance local peaks have smaller basins of attractions than high-performance local peaks. With increasing complexity K the "tendency for the highest optima to have the largest drainage basins dwindles [...] but even for modest values of K, some very large basins persist" (Kauffman 1993: p.63). Moreover, for small values of K, Kauffman (1993) shows that there is a clustering of peaks, i.e., the highest peaks are near one another (Massif Central).

3.2 Local Search

Like in previous modeling efforts (Levinthal 1997), in our model the firm starts in period $t=0$ from a randomly determined decision vector $\boldsymbol{a}=(a_1, a_2,.., a_N)$, whose average performance amounts to 0.5 on the normalized performance landscape. In the following periods, the firm seeks to improve its performance through a process of local search: in each period, one decision a_i of the vector $\boldsymbol{a}$ is inverted. If the modified decision vector yields a higher performance, it is adopted and the search continues from this new vector in period $t+1$. Otherwise, this modification is discarded again and the next search step starts from the unchanged vector defined in period t.[9] This process may be interpreted as a search for better positions on a high-dimensional performance landscape ("hill climbing").[10] This kind of local search process implies that if the firm actually modifies its decision vector, performance will always increase; a firm will never adopt a modification that decreases performance. Sooner or later, the firm will get stuck at a decision vector whose performance cannot be improved by changing one of its N decisions. In this case, the firm is either at a local peak or the global peak.

3.3 Replication Process and Replication Errors

The local search process described above implicitly assumes that "temporal replication" is always accurate. In our model, we relax this assumption for one period. In period $t=R+1$, the replication process may not be accurate. Formally, $\varepsilon \in \{0,..,N\}$, decisions out of N decisions of the decision vector in $t=R$ to which we refer as the "replication template" or simply "template" may not be replicated accurately[11]. Technically, a replication error is implemented by an inversion of ε bits (from 0 to 1 or from 1 to 0) of the replication template. The ε decisions affected by replication errors are randomly chosen among the whole set of N decisions. Although we model a "temporal" replication process, our findings also generalize to "spatial" replication processes: the first R periods can be thought of as the performance of the organizational unit that creates the replication template ("source"), while the subsequent R periods can be thought of as the performance of the receiving unit.

Starting from a random position in t=0, to which we refer to as the "starting practice", the source organization gradually improves its practice through local search until period $t=R$. For our choice of the replication period (i.e., R=200), all firms have converged to either a local or global peak at the end of period 200.[12] Hence, the template represents a locally or globally optimal decision vector, which cannot be improved anymore through local search. In period $t=R+1$, this template is transferred (more or less correctly) to the recipient organization. We refer to the practice in $t=R+1$ as the (more or less accurate) "replica". In the following R periods, the receiving organization seeks to improve the replica again through local search. By period $t=2R$, the recipient organization has converged again to either a

[9] Thus, we (implicitly) assume that new decision vectors can be assessed in an "off-line" manner. For studies that analyze "online" search, see Levitt and March (1988), Gavetti and Levinthal (2000) and Denrell et al. (2004).

[10] In this local search process, the organization is not assumed to always choose the mutation with the steepest gradient. It simply chooses the first superior alternative it discovers. Yet, our results also generalize to a setting where firms test all local alternatives and choose the alternative with the highest performance.

[11] For analytical reasons, we suspend the local search process for one period, i.e., the firm does not search locally in period $t=R+1$.

[12] In more complicated or compound hierarchical organizations, Rivkin and Siggelkow (2002) show that a firm may get stuck at a decision vector (sticking point) that is not a local peak on the fitness landscape of the overall organization. In our model, all sticking points are either local or global peaks.

local or global peak. We refer to the decision vector the firm achieves by period t=2R as the "final practice". In our model, we consider the case in which a template is transferred to only one receiving unit. Nevertheless, our results also generalize to multi-unit transfers if we assume that there is no inter-unit learning.[13] Figure 1 gives a high-level overview of our model.

<Insert Figure 1 about here>

On a more abstract level, replication errors are mutations that introduce novelty but also move a firm away from an attractive position on the performance landscape because the replication template is assumed to be either a local or global peak. Although both replication errors and local search processes randomly flip some bits of the decision vectors, there is an important difference. The local search process never yields a decrease in performance because the "mutated" decision vector is only adopted if it increases performance. Otherwise, it is discarded and the search process continues from the original vector. In contrast, replication errors can, and most likely will, decrease performance in the subsequent period, in particular, if the replication template is locally or globally optimal. With replication errors, the status quo cannot be recovered easily because it is the very nature of replication errors that they are not easy to detect. Hence, it can be extremely tedious to correct replication errors (Polanyi 1959). Thus, a replication error is different from local search in that – even in the case of a negative short-run performance effect – the status quo cannot be recovered.

4. Results

4.1 Implementation details

In the following sections, unless mentioned otherwise, we report results for the case of N=10 and K=2. To make sure that the reported differences are inherent to our model (rather than a result of the stochastic elements of our model), we repeated each experiment 100'000 times with different starting seeds for both the random interaction matrices and the initial position of the firm on the performance landscape.[14] This procedure ensures that the reported simulation results are statistically significant at the 1% level.We observe the firms for 400 periods. In period t=R=200, the firm may or may not experience errors in replicating its practice. We measure the short-run and long-run performance of the firms as the average performances in periods t=R+25 and t=R+200, respectively.

[13] Inter-unit learning is the focus of, for example, Csaszar and Siggelkow (2010).

[14] For studies that analyze non-random interaction patterns, see e.g., Rivkin and Siggelkow (2007) and Baumann and Siggelkow (2012).

4.2 Short- and Long-Run Performance Effects of Replication Errors

In our first experiment, firms are searching in performance landscapes of low complexity (K=2). Figure 2 provides a systematic analysis of the short-run (gray line) and long-run (black line) performance implications (y-axis) of the full range of replication errors (x-axis), i.e., from accurate replication (zero errors) to completely inaccurate replication (ten errors). The gray line represents the average performance difference in the short run between the replication template in t=R and the more or less replicated replica in t=R+25 (left y-axis). In other words, this measure reflects the more immediate performance impact of inaccuracies in replication. The black line reflects the average performance difference in the long run between the final practice to which the firm converges in period t=R+200 and the replication template in t=R (right y-axis).

<Insert Figure 2 about here>

Not surprisingly and consistent with Winter and Szuslanksi (2001), we find that replication errors always have strong negative short-run performance implications even tough we allow the firms 25 periods to correct replication errors. Recall that the replication template is at least a local peak if not even a global peak, i.e., the template cannot be further improved through local search.[15] Deviations from such a peak are, on average, always performance decreasing. Moreover, the larger is the replication error the lower is the performance of the replica. Yet, in the long run, firms that experience small replication errors (i.e., errors of size 1 to 5) can outperform those firms that accurately replicate the template in period R.[16] Despite the negative short-run performance consequences, inaccuracies in replication can improve average long-run performance. This is particularly surprising given that the replication template is already either a local or global optimum. Obviously, those firms in global optima cannot improve any more but instead, replication errors can at best be corrected and in the worst case, move a firm permanently away from the global optimum. Particularly, we observe an inversely U-shaped relationship between the size of replication errors and the average long-run performance difference between the final practice and the replica that reaches a maximum for intermediate levels of replication errors.

One might be tempted to conclude that replication errors have the same long-run performance effects as "long jumps" (Kauffman 1993) or a process of "global search" (Levinthal 1997; Rivkin 2000). Interestingly, we find that replication errors in the long run always outperform a comparable one-time expansion of the search radius (long jump or global search). This result holds for the whole range of

[15] We find that by period 200, already 29% of all firms have converged to the global peak.

[16] This positive effect on long-run performance of small replication errors can be reinforced through repeated replication errors. Yet, the time span between the periods where the replication errors occur must be sufficiently large.

possible expansions of the search radius.[17] Like long jumps or global search, replication errors can move an organization to a more distant position on the performance landscape. Yet, with long jumps and global search, this move is contingent upon an improvement of performance: if there is no improvement compared to the current position, the firm will not move. With inaccurate replication, in contrast, the performance of the replica does not matter because the replication error is enforced no matter what its performance consequences are. Additionally, the firm does not know where the error has occurred. Indeed, inaccurate replication has always a negative effect on average short-run performance as seen above.[18] Yet, as we will see in the subsequent analysis, it is the unintended nature of errors and the willingness to adopt an (on average) inferior replica that is necessary for the positive long-run performance effects of replication errors to materialize.

While we are not the first to argue that learning processes may benefit from small errors or some randomness (e.g., March 1991; Denrell and March 2001; Knudsen and Levinthal 2007; Csaszar and Siggelkow 2010), our study is different from existing research in two important ways. First, prior studies often focus on errors in the performance feedback (Denrell and March 2001, Knudsen and Levinthal 2009). In the context of our study, such errors would imply that the firm may experience inaccuracies in perceiving the performance associated with a given decision vector: the firm may over- or under-estimate the performance of a given position on the performance landscape (Denrell and March 2001; Knudsen and Levinthal 2007). In our study, in contrast, errors occur at the level of the decision vector; the performance associated with the inaccurately replicated decision vector is always perceived correctly. Second, in our study, we focus on the average performance implication for a single, isolated firm while prior research often focuses on the implications of errors in a population of firms (Csaszar and Siggelkow 2010). In population of firms, an implicit assumption is that performance-decreasing errors at one firm can be corrected again by imitating other, better performing firms. In our study, this information is not available to firms.

In our analysis, we proceed as follows. First, we will seek to uncover the mechanisms that generate the inversely U-shaped relationship between the size of replication errors and the long-run performance difference by disentangling the effects of replication errors on discovering new practices and abandoning the current template (Sections 4.3-4.5). Second, we identify boundary conditions

[17] Csaszar and Siggelkow (2010) also examined the implications of replication errors (yet, under the label of random search, compare their Figure 5). Their results deviate from our results because long jumps or “replication errors” always occur once the firm got stuck at a local or global peak. As a result, their measure of long-run performance in t=100 may also include firms that experienced a replication error very recently, e.g., in t=99. In other words, their long-run performance is not a steady state performance. In our study, in contrast, we report steady state performances (i.e. the performance 200 periods after the replication error had occurred).

[18] Even though the average short-run performance decreases through replication errors, some firms actually benefit from replication errors in the short run because there is a small probability that the replica yields a higher performance than the template.

under which small replication errors have particularly positive long-run performance consequences (Section 4.6).

4.3 Discovering Better Peaks Through Replication Errors?

A necessary condition for long-run performance to improve through replication errors is that a firm permanently abandons its replication template and converges to a different practice. If the firm sticks or returns to its replication template, performance will not change. If it abandons the replication template, the quality of the new practice then determines whether performance improves or not through replication errors. Thus, one potential reason why small replication errors improve long-run performance could be that they trigger the discovery of better practices. Or, in more technical terms, a replication error permanently dislodges the firm from its current peak and, subsequently, the firm converges to a different and better peak. However, why should the "new peak" have a higher performance than the abandoned peak?

From Kauffman (1993), we know that for low to moderately complex landscapes (i.e., small values of *K*), peaks are not distributed randomly across the landscape, but instead are clustered around the global peak (compare Section 3.1). We also know that for small values of *K*, there is a positive correlation between the performance of a local peak and the size of its basin of attraction. One might expect that small replication errors may dislodge a firm from one of the local peaks and, given this positive correlation and the clustering of peaks, let it converge to a nearby better peak.

However, we find no evidence for such a dynamic. In contrast, we find that replication errors lead to the discovery of practices whose average performances are below the average performance of the replication template. The gray line in Figure 3 reports the average long-run performance of the new practice for those firms that had been permanently dislodged from their replication template for different sizes of the replication error (x-axis). The black line illustrates the average performance of the replication template.

<Insert Figure 3 about here>

The figure shows that the average performances of the peaks that are discovered by those firms that permanently abandon their replication template are lower than the average performance of the replication template.[19] Specifically, the long-run performance of newly discovered practices is a decreasing function of the size of error, i.e., the larger is the replication error the worse is the performance of the final practice that is newly discovered in the long run. In a sense, replication errors

[19]Note that for error size ε=0, firms never abandon their replication template.

seem to induce innovation (i.e. a change of the status quo) but not necessarily to particular beneficial innovation (compared to innovation that started from a random position).

Hence, the inversely U-shaped relationship between the size of replication errors and long-run performance differences *cannot* be attributed to the discovery of better practices through replication errors. Indeed, if a replication error leads a firm to permanently abandon its replication template, the new practice the firm will ultimately discover is (on average) worse than the average practice (i.e., the template) discovered through random search by period $t=R$. Put differently, replication errors might be a source of innovation in the sense that they lead firms to move to different positions on the landscape, yet they are not necessarily a source of "good" innovation: while firms might improve their current template, replication errors don't let firms discover particularly attractive positions on the performance landscape. Indeed, if firms start from a random position on the performance landscape, they will on average find better practices than those firms that inaccurately replicate a locally or globally optimal template. That is, the inaccurately replicated template provides a worse starting point than a random position on the performance landscape.[20]

4.4 Abandoning Low Local Peaks Through Replication Errors?

The size of the replication error does not only affect which practices are *discovered* but also which replication templates are (permanently) *abandoned*: Indeed, average long-run performance can improve if replication errors make firms discover good practices or abandon bad replication templates. As discussed above, the height of a peak is positively correlated with the size of its basin of attraction: the better is the performance of the local peak the larger is its basin of attraction. As a result, if firms are placed on a random position on the performance landscape, the probability that they are in the basin of attraction of the global peak is higher than the probability that they are in the basin of attraction of a low-performance local peak (Levinthal 1997).

Yet, this is just one side of the coin. The other side is the abandonment of a peak. The probability that a firm can escape the basin of attraction of its current peak is a function of the size of the peak's basin, which in turn, is an increasing function of its height. Given a certain size of error, the probability is higher to get dislodged from low-performance peaks than for high-performance peaks. Given a certain performance of the replication template, larger replication errors yield a higher probability of dislodging a firm from its current template than smaller replication errors.

In Figure 4a, we first analyze how the size of the replication error affects (i) the probability of abandoning the replication template and (ii) the performance of the abandoned template. The gray line

[20] Due to the clustering of local peaks around the global peak (Kauffman 1993), the firms are placed in an "unattractive" area of the landscape through the replication error.

shows the probability (left y-axis) that a firm is permanently dislodged from its replication template as a function of the size of the replication error (x-axis). The black line reports the average performance (right y-axis) of the replication template as a function of the size of the replication error (x-axis), conditional on that the template was permanently abandoned.

<Insert Figure 4 about here>

Both, the probability of abandoning the template and the average performance of the abandoned template are increasing in the size of the replication error. In the extreme case of $\varepsilon=10$, the probability approaches 100%, i.e., the template is almost always abandoned, regardless whether the firm is in a local or global peak by period *R*. For smaller replication errors, the probability of abandoning the template (or "exploration") is decreasing. The average performance of those templates that are abandoned is also increasing in the size of the replication error. The reason is that the replication template is an instrument of selective retention, i.e., replication errors in high-performance templates are easier to correct than in low-performance templates due to the positive correlation between the performance of a local peak and the size of its basin of attraction. Consider, for example, a firm that experiences a small replication error. If it has reached a high-performance replication template, it is associated with a large basin of attraction. Even if the replication error temporarily drives the firm several "steps away" from the template, the firm has a good chance to find back to the original template in the long run through the local search process because the "leap" generated by the replication error is not large enough to place the firm in the basin of attraction of another peak (Rivkin 2000). Now, let us assume that the firm has reached a low-performance template, which is typically associated with a comparatively smaller basin of attraction. In this case, the replication error drives the firm out of the basin of attraction of such a low-performance template.

In sum, given a certain performance of the template, firms that experience large replication errors are more likely to permanently abandon their replication template than firms with small replication errors. In addition, the templates that they abandon with large replication errors are, on average, better performing than those templates abandoned by firms that experience only small replication errors. It follows that a firm that experiences a small replication error only permanently abandons its replication template if, in general, it has a low average performance. Large replication errors, in contrast, do not distinguish between low- or high-performance replication templates and therefore let firms permanently abandon almost any template, regardless of its performance.

Figure 4b summarizes our findings on how replication errors affect the abandonment of the template and the discovery of new practices. Specifically, we report the average net performance effect, i.e., the difference in performance between the newly discovered final practice and the abandoned template, conditional on that the template was permanently abandoned. The black line in Figure 4b represents the average performance (left y-axis) of the replication templates that are permanently abandoned as a function of the size of the replication error (x-axis). The gray line in Figure 4b illustrates the average performance (left y-axis) of the newly discovered final practice.[21] The black dashed line represents the net effect (right y-axis), i.e., the difference between the gray and black line.

The net performance effect between the newly discovered final practice and the abandoned template is positive for small replication errors and decreases to negative levels for large replication errors. In other words, exploration that is triggered by small replication errors is, on average, improving performance while exploration triggered by large replication errors is decreasing performance.

4.5 The Combined Effects of Discovering and Abandoning

If we consider only the net performance effect of an explorative event, small replication errors should yield the highest long-run performance. Our first analysis (Figure 2), however, suggested that moderate replication errors lead to the highest long-run performance. Only if we also take into account the probability that a replication error triggers an exploration event, we can explain the inversely U-shaped relationship between size of replication error and long-run performance.

In Figure 5, we multiply the net performance effect between discovery and abandonment (black line, right y-axis) with the probability of exploration (gray line, left y-axis). The resulting line (black dashed line, right y-axis) reflects the average long-run performance effects of replication errors.

<Insert Figure 5 about here>

The black dashed line exactly matches the shape of the long-run performance effects of replication errors. The inversely U-shaped relationship between size of replication errors and long-run performance is the consequence of a monotonic increase in the probability of exploration in error size and a monotonic decrease in the performance gains from each exploration event. In other words, small replication errors have a low probability of triggering an exploration event (i.e., the replication template is permanently abandoned) but if the firm explores an alternative practice, this new practice will, on average, be better than the abandoned template. With large replication errors, in contrast, the probability of abandoning the replication template is very high. Yet, the net effect between the performance of the new practice and the abandoned template is at best zero if not negative. As a result, firms can hardly gain from large replication errors.

In sum, there are three important mediators in the relationship between replication errors and long-run performance: first, the probability that a replication error induces the firm to abandon its current (locally optimal) template. This probability is increasing in the size of the error. Second, the size of

[21]Note that the black lines in Figures 4a and 4b and the gray lines in Figure 3 and 4b correspond to each other,

the replication error determines the quality of the template that is abandoned – larger replication errors make firms abandon better templates. Third, the size of the replication error affects which new practice is discovered if the current template is abandoned – smaller replication errors make firms discover better practices. The interplay of these three mediators determines the long-run performance effects of replication errors.

It is important to note that although the replication dilemma between learning and accuracy is akin to the fundamental trade-off between exploration and exploitation (March 1991), there are some important differences. In the context of the NK landscape model, exploitation is often associated with "local search" and exploration with "long jumps" or "global search" (Kauffman 1993; Levinthal 1997; Rivkin 2005). Yet, neither exploitation nor exploration leads to a decrease in performance ("hill-climbing"). In our model, however, a high level of accuracy (no replication errors) means that the firm's performance remains flat, i.e., it neither decreases nor increases because the replication template is a local peak. On the other hand, learning that is induced through replication errors means that the firm can discover new positions on the landscape. However, unlike augmenting the search radius (long jumps), increasing the size of replication errors is not performance enhancing. Indeed, experiencing a replication error always leads to a decrease in average performance in the short run.

4.6 When are Long-run Performance Gains of Small Replication Errors most Pronounced?

4.6.1 The Moderating Effect of the Quality of the Replication Template

While small replication errors *on average* improve long-run performance, obviously this does not imply that replication errors *always* improve long-run performance. Replication errors may improve performance for some firms while other firms cannot benefit from replication errors. Not all firms are affected equally by small replication errors. Consider the two polar cases. At the one end of the continuum, there are firms that have already converged to the global peak by period $t=R$. Obviously, for such firms there is no upside to replication errors. Any replication error will lead the firm to a lower position on the landscape. At best, the firm will only temporarily abandon the global peak but sooner or later will find its way back to it. At the other end of the continuum, there are firms in the lowest local peak on the performance landscape. A replication error may dislodge these firms from this peak and, eventually, let them discover higher performing peaks. There is no downside to replication errors in this case. At best, they discover a better peak; at worst, they are only temporarily removed from their current local peak but ultimately return to it.

On a more abstract level, the long-run performance effect of replication errors depends on the quality of the template that is to be (more or less accurately) replicated. The better the replication template is, the lower is the potential upside to replication errors and the higher is the potential downside. This is exactly what we observe in Figure 6, which depicts the average long-run performance gains/losses (y-axis) for different sizes of replication errors between the final practice and the template as a function of the replication template's performance (x-axis).

<Insert Figure 6 about here>

On average, firms gain from replication errors of size ε=2 if they inaccurately replicate a template whose average performance is below 0.95. The contrary is true for firms that inaccurately replicate a template whose average performance is above this threshold. For replication errors of size ε=10 and ε=1 the critical thresholds are 0.93 and 0.95, respectively. The figure also shows that the gains from replication errors are more or less linearly decreasing in the quality of the template and eventually turn into losses depending on the size of error. The difference between small and large replication errors comes from the fact that the probability of abandoning the template increases with the size of error.

In sum, we observe that firms are affected differently by replication errors: if the replication template is of low performance, replication errors tend to improve long-run performance. If the replication template is of high performance or even the global peak, replication errors tend to decrease long-run performance.

4.6.2 The Moderating Effect of Complexity

As explained in Section 3.2, the positive correlation between the performance of a local peak and the size of its basin of attraction is particularly strong in low to moderately complex landscapes and dwindles with increasing complexity (Kauffman 1993, p.62). Given that both the abandonment and discovery effect discussed above are driven by this positive correlation, we should expect the positive (negative) long-run performance effects for small (large) errors to dwindle with increasing K as well.

Figure 7 provides a systematic analysis of the short-run and long-run performance implications (y-axis) of the full range of replication errors (x-axis) for different levels of complexity. K=0 represents a non-complex landscape, K=2 and K=4 a low and moderately complex landscape, respectively, while K=8 is a highly complex landscape. Panel (a) represents the difference in performance between the replication template in t=R and the performance of the more or less accurately replicated template in the short run (i.e., in period t=R+25). Panel (b) illustrates the long-run performance difference between the final practice to which the firm ultimately converges in period t=R+200 and the replication template in t=R.

<Insert Figure 7 about here>

In the short run, all types of replication errors have a negative impact on average performance for all levels of complexity. The negative performance effect increases in the size of errors for non to moderately complex landscapes, while the size of error has no influence in highly complex environments. Moreover, for small errors the negative performance effect is more pronounced in highly complex landscapes than in non to moderately complex landscapes, while the opposite holds true for large errors. Here, large errors induce a stronger performance drop in the short run if they occur in non to moderately complex landscapes. In the long run, replication errors have no effects in non-complex landscapes independent of their size. As discussed above, in low to highly complex landscapes, however, small errors have positive performance effects, where the effect is more pronounced for low to moderately complex landscapes than for highly complex landscapes. Large errors on the other hand, have negative effects in low/moderately complex landscapes and no effects in highly complex landscapes.

Our analysis confirms the intuition that the long-run performance effect of replication errors is weakened with increasing levels of complexity because the positive correlation between the performance of a local peak and the size of its basin of attraction dwindles. The relative performance effects are strongest for low and intermediate levels of complexity, i.e., K=2 and K=4. For high levels of complexity (K=8), the positive effects of small errors are comparatively small, while surprisingly large errors have no effect on performance.

For non-complex environments (i.e., K=0), the performance landscape has only a single peak whose basin of attraction is the entire landscape. Any replication error will only temporarily dislodge the firm from the global peak. Sooner or later, it will find its way back to the global peak. With increasing levels of complexity, the landscape becomes more rugged with numerous local peaks and hence the probability that the replication template is a global peak decreases.

In more complex environments (i.e., K>0), we observe differences in long-run performance effects of small and large replication errors. With large replication errors, a firm is likely to be dislodged from any peak, even if the replication template is the global peak. With small replication errors, in contrast, the firm is more likely to be dislodged from low-performance peaks and less likely from high-performance peaks. This is the result of the positive correlation between the performance of a local peak and the size of its basin of attraction – higher peaks have larger basins, which are harder to escape than the small basins of lower peaks. This positive relationship dwindles with increasing complexity (Kauffman 1993). With a negative correlation (which, however, cannot be generated by the standard NK landscape model), small replications errors would have negative long-run performance consequences, while large errors would still be neutral.

Even though replication errors decrease *average* performance in the short run, some firms may actually benefit from a replication error: by chance, the replication error may move them to a position on the landscape that has a higher performance than the replication template. We find that the probability to obtain a higher performance in t=R+1 compared to t=R increases with the size of error and the level of complexity: while no firm moves to a better position in t=R+1 through a replication error of size ε=1, this number changes for replication errors of size ε >1:[22] in low levels of complexity (K=1), 1% of the firms have a higher performance in $t=R+1$ compared to $t=R$ and this percentage increases up to 6% for high levels of complexity (K=9).

In sum, we find that in the short-run all types of errors have negative average performance effect, regardless of the level of complexity of the performance landscape. In the long run, however, we find a positive average performance effect for small errors, while large errors have negative average performance effects in low to moderately complex environments; otherwise, large replication errors are neutral in their long-run performance effect.

Finally, Table 1 provides a detailed summary of all our results in this study.

<Insert Table 1 about here>

5. Conclusions and Discussion

In our study, we sought to enhance our understanding of why prior research generated inconsistent findings on the performance implications of replication errors. The notion that replication errors (depending on, for example, their size) may have positive and negative implications has been already recognized more or less explicitly in the works of (Winter and Szulanski 2001).[23] They postulate the existence of a "replication dilemma" associated with the replication process, i.e. there is a trade-off between accuracy and learning. Accurate replication (or replication without errors) allows organizations to reap the economic benefits associated with the practice to be replicated. At the same time, however, accurate replication impedes any learning, i.e. it reproduces only the status quo and forgoes the opportunity to further improve the current practice. In our study, we use an NK simulation model to systematically examine these opposing forces and identify conditions under which replication errors may prove positive, negative, or neutral in their effect on performance.

Consistent with prior research, we found that all types of replication errors always have a negative average short-run performance effect which increases in the size of errors. In the long run, however, the performance effect of replication errors depends (i) on the size of error and (ii) on the level of complexity. Small replication errors always have positive long-run performance effects in the presence of complexity and their effects are more pronounced in low to moderately complex

[22] Recall that all firms are in local peaks in t=R and therefore cannot improve their performance through a one-bit change.

[23] In our analyses, we focus on the performance implications of replication errors. We also examined the effect of replication errors on a firm's innovativeness. In the context of our study, "innovativeness" might be understood as either the probability that the firm gets dislodged from its current practice and discovers a new practice (this probability is increasing in the size of the replication error, compare Figure 4a) or as how different the newly discovered practice is from the replication template (an increasing function of the size of the replication error. This analysis is not included in our study but it is available upon request).

environments than in highly complex environments. Large replication errors have a negative impact on performance in low to moderately complex environments, while they have no effects on long-run performance in the absence of complexity or the presence of high levels of complexity.

Thus, on a more abstract level, our study is an attempt to reconcile the inconsistent findings of prior research by identifying boundary conditions under which replication errors can be positive, negative, or neutral in their effect on performance, depending on two moderating variables: complexity of the organization's environment and the size of the replication errors. Studies, such as Winter et al. (2011) that find a consistent negative effect of replication errors on performance may have adopted a time horizon that was too short for the long-run positive performance effects to materialize. Alternatively, they may have focused on large errors in firms that operate in environments of low to moderately complexity. Studies that have identified a positive effect of replication errors such as Becker et al. (2006) may have focused on the long-run performance implications of small replication errors.

We observe consistent negative short-run performance effects for any type of replication error because the practice to be replicated is either a global or local optimum. Any deviation from such an optimum will have, on average, negative (short-run) performance effects. Even though the average performance decreases in the short run through replication errors, some firms actually benefit from replication errors because in the subsequent period they find a practice that yields a higher performance.

The intuition behind the long-run effects of replication errors is as follows. In the absence of complexity, the performance landscape has only one global peak. In such an environment, any replication error dislodges the organization only temporarily from the global optimum and thus in the long run, the firm will rediscover this global optimum. With increasing levels of complexity, the landscape becomes more rugged with numerous local peaks. In such complex landscapes, the long-run performance effects of replication errors depend on the error size. With large replication errors, a firm will be dislodged from any peak, even if it is the global peak. With small replication errors, in contrast, the firm is more likely to be dislodged from low-performance than from high-performance peaks. This is the result of the positive correlation between the performance of a local peak and the size of its basin of attraction. This positive correlation dwindles with increasing complexity. With a negative correlation, small replication errors can be expected to have negative long-run performance implications because they are more likely to dislodge a firm from high-performance than from low-performance peaks. Large replication errors, in contrast, would have positive performance effects in the long run on low to moderately complex environments.

In sum, our study suggests that replication errors only have a sustainable long-run performance effect if they are large enough to permanently dislodge the firm from its current practice. Another necessary condition is that they are not corrected immediately because, for example, they cannot be easily detected. Moreover, positive long-run performance implications of small replication errors only materialize if the replication template is an instrument of "selective retention", i.e., ceteris paribus, templates with a higher performance are harder to abandon than templates with a lower performance. This property is present in low to moderate levels of complexity in NK models, because then a positive correlation between the performance of a local peak and the size of its basin of attraction is present (Kauffmann 1993; Rivkin 2000).

Our findings contribute to the existing literature on replication in several ways. First, while the positive and negative effects of replication errors have been highlighted before, a systematic analysis of their combined effects has been missing. In our study, we quantify their relative strength for a variety of different settings. For example, while Rivkin (2000) focuses primarily on the case of replication practices being literally best practices (or global peaks), our study extends his analysis to settings in which the replication template is only a local peak. Second, our study also enhances our understanding what may constitute the core elements (or Arrow core) of a practice. Our study suggests that influential decisions constitute the Arrow core because replication errors in these decisions carry the strongest performance consequences (rather than decisions with a high impact or decisions that highly depend on other decisions). Third, our study also points to a potentially important characteristic of a replication template that has been neglected so far: the replication template as an instrument of selective retention. The value of a template is often understood in terms of the knowledge it embodies; replication templates may differ in their performance - a template may be associated with a high or low point on the performance landscape. This point or in the case of replication errors, the area around this point is then the starting point for the recipient organization of the replication template. Our study suggests that in the case of large replication errors, the inaccurately replicated template provides a worse starting point than a random position on the performance landscape. If a firms starts from a random position on the performance landscape, it will on average find a better practice than a firm that inaccurately replicates a locally or globally optimal template. From a dynamic perspective, however, it is also important whether a replication template is easy or hard to abandon. Ideally, a high-performance replication template should be hard to abandon while a low-performance template should be easier to abandon. Thus, the benefits of replicating an existing template (versus starting from scratch) depend on the performance of the template. Yet, the benefits also depend on the higher probability of high-performance templates to absorb and correct replication errors than low-performance templates.

Moreover, our study contributes to several related streams of research. In particular, our study sheds new light on the interplay between errors and learning. In the existing literature, the notion of errors and learning are discussed under a variety of different labels. For example, studies in the tradition of the behavioral theory of the firm (Simon 1955; Nelson and Winter 1982; Cyert and March 1992), often invoke the concept of trial-and-error learning. The efficacy of trial and error learning heuristic hinges on its ability to distinguish between good and bad solutions: "trial and error is not completely random or blind; it is in fact rather highly selective. The new expressions that are obtained by transforming given ones are examined to see whether they represent progress toward the goal. […] Problem solving requires selective trial and error" (Simon 1996: p.95-96). In our study, in contrast, learning through replication errors requires the adoption of an on average inferior practice, i.e., a practice that would have been abandoned by (selective) trial and error learning processes.

Related, but on a population rather than an individual level, research in the evolutionary perspective highlights the importance of a combination of blind variation and selective retention (Campbell 1965). Replication errors are a source of variation, which in turn, is the raw material for successful organizational adaptation and improvement (Holland 1975; March 1991). In populations, the potential benefits of replication errors are magnified by a selective retention mechanism on the population levels (Csaszar and Siggelkow 2010). Variations that improve performance are spread across the population while variations that decrease performance are contained to their source and, sooner or later, corrected by spreading of more successful variants. In our study, in contrast, we focus on the individual level and the benefits of spreading successful practices are not incorporated in our model. In our study, the benefits of errors as a source of variation are realized on the individual firm level: i.e., it does not require that the firm be embedded in a population of firms from which it learns.

Finally, there is also a large body of literature on learning from errors, which portrays the learning process as a process of learning to correct errors or learning to prevent errors to occur. Argyris (1982) even defines learning as detection and correction of error. However, in our study, learning from errors requires that errors cannot be corrected easily and cannot be prevented. If they are easy to correct, the firm will just return to its previous practice, there is no learning in the sense that firms will discover a new and potentially better practice.

FIGURES

Figure 1: Illustration of Simulation Model

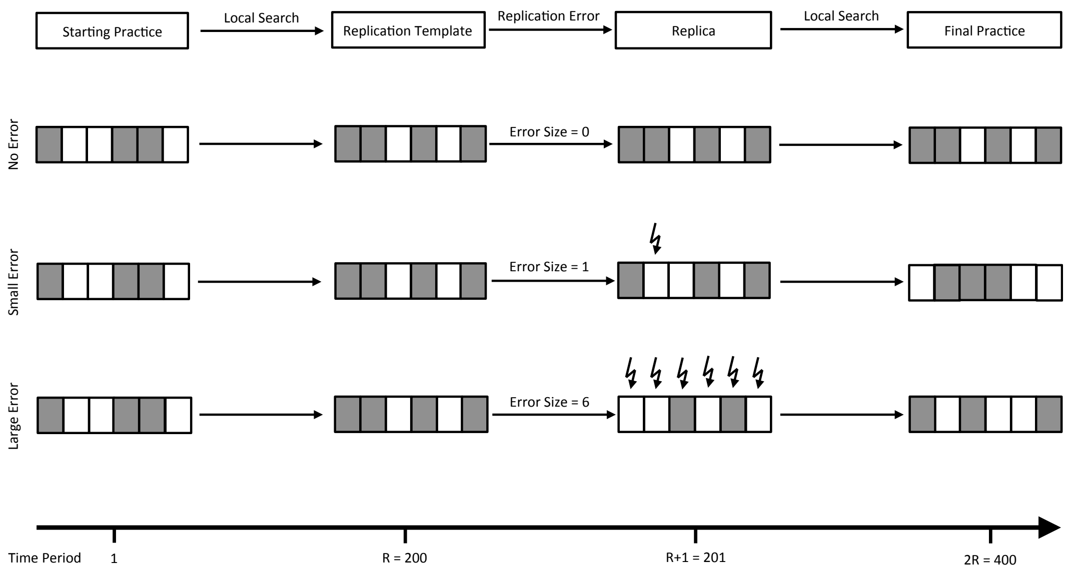

Figure 2: Short- and Long-Run Performance Implications of Replication Errors

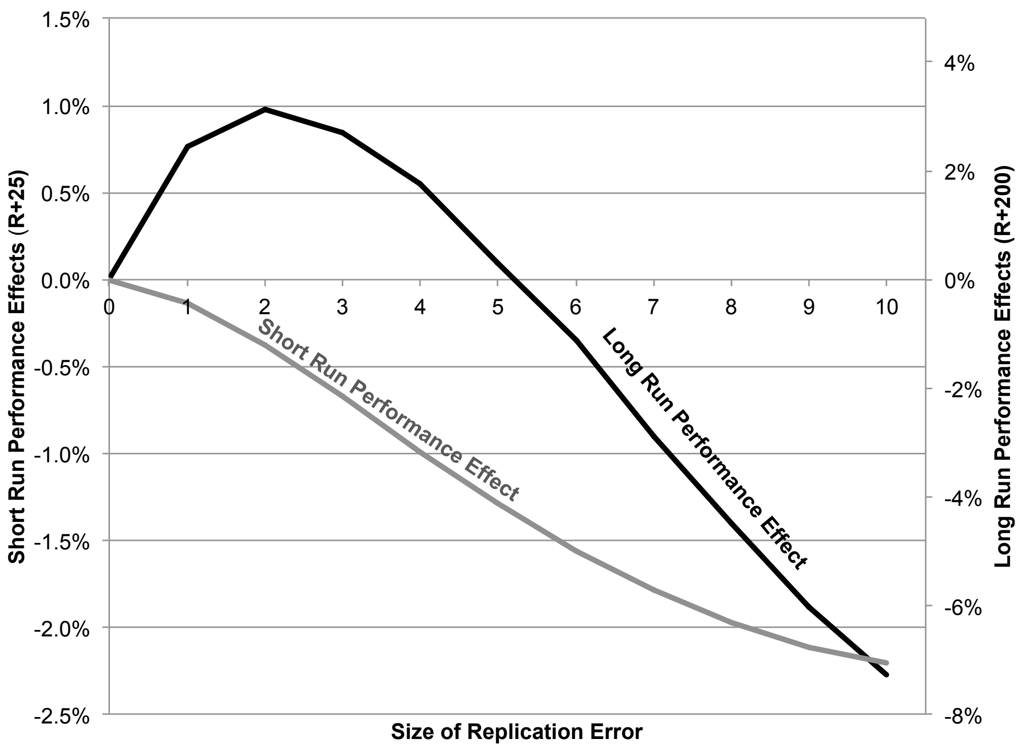

Notes: This figure reports the average short-run (t=R+25=225, left y-axis) and average long-run (t=R+200=400, right y-axis) performance effects of replication errors (x-axis), ranging from zero errors (left) to errors in all dimensions (right). The performance effect is measured as the difference to the performance of the replication template in t=200. The results are based on 100.000 landscapes (with N=10 and K=2).

Figure 3: Performance of Final Practice if Replication Template is Abandoned

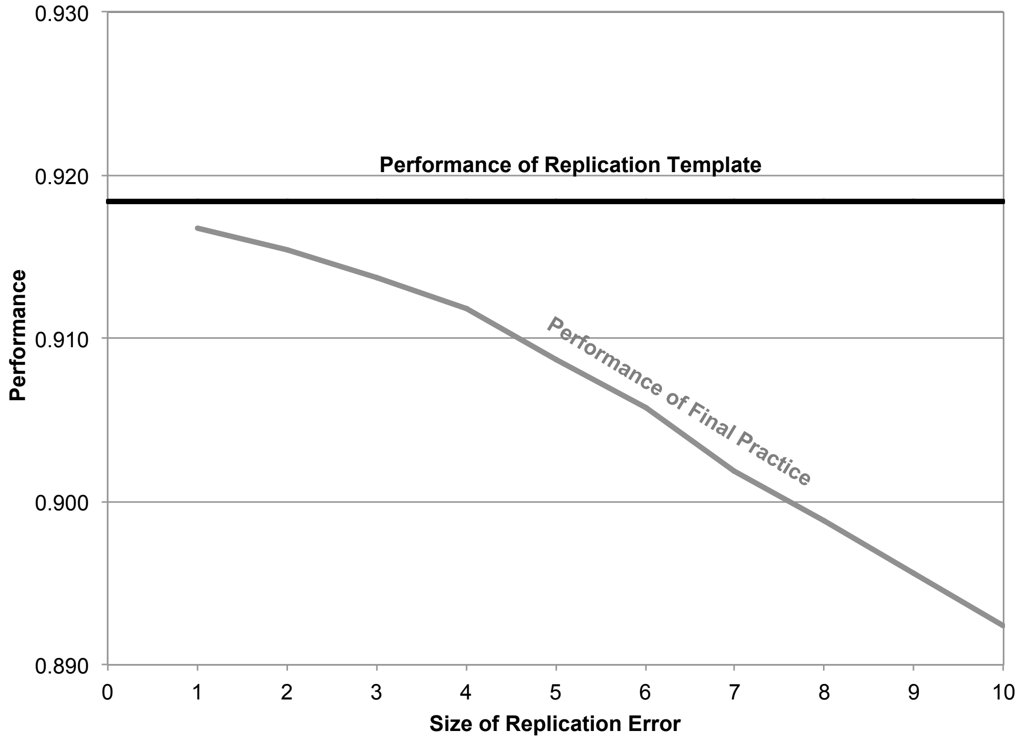

Notes: This figure reports the average long-run performance (y-axis) of the those firms that get dislodged from their template in *t*=200 through an replication error. We report the size of the error on the x-axis ranging from zero errors (left) to errors in all dimensions (right). With zero errors, there is no dislodging and therefore no new practice is found. For the sake of comparison, we also included the average performance of the replication template in *t*=200 (pre-replication). The results are based on 100.000 landscapes (with *N*=10 and *K*=2).

Figure 4a: Abandoning Current Peaks

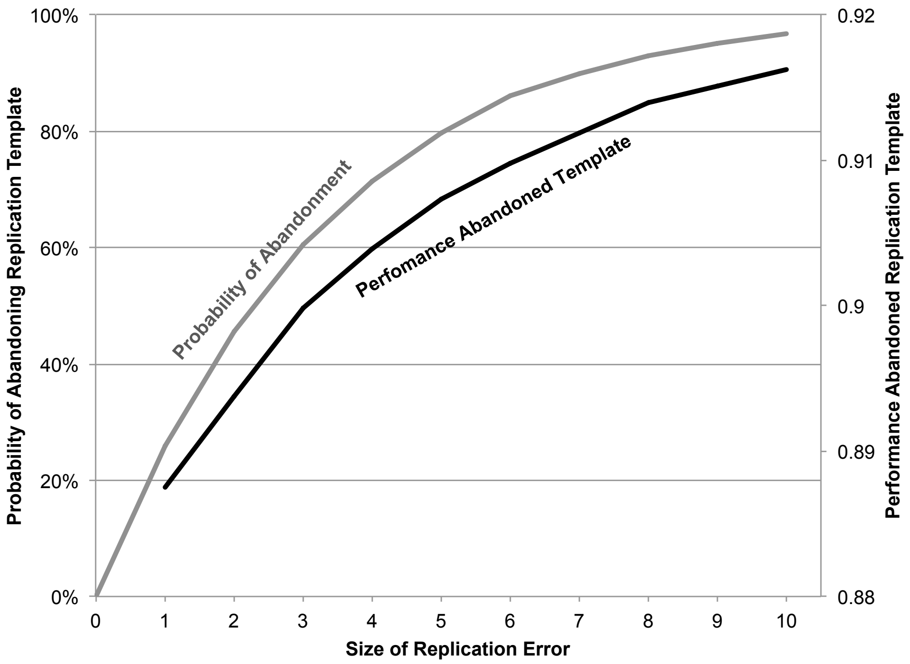

Notes: This figure reports the probability of abandoning (left y-axis) the replication template in t=200 as well as the average performance of the abandoned template (right y-axis) for the whole range of replication errors (x-axis) ranging from zero errors (left) to errors in all dimensions (right). With zero errors, the replication template is never abandoned. The results are based on 100.000 landscapes (with N=10 and K=2).

Figure 4b: Combined Effects of Discovery and Abandonment

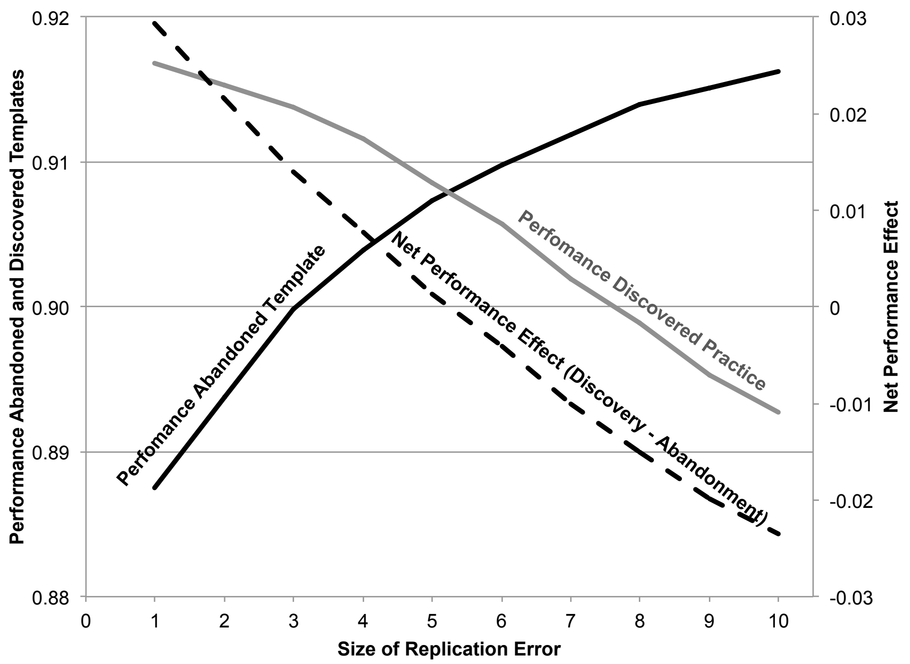

Notes: The figure reports the average performance of the abandoned template and the discovered practice (left y-axis) as well as the net effect (right y-axis), i.e., the difference between both lines for the whole range of replication errors (x-axis) ranging from zero errors (left) to errors in all dimensions (right). The results are based on 100.000 landscapes (with N=10 and K=2).

Figure 5: Probability and Performance Effect of Exploration Events

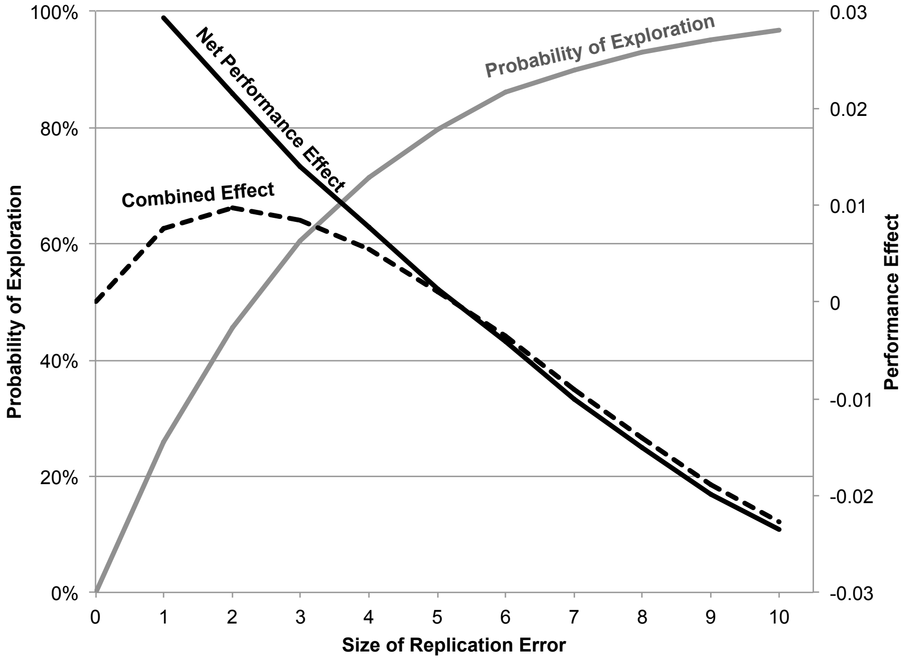

Notes: The figure reports the net performance effect between discovery and abandonment (right y-axis) as well as the probability of exploration (left y-axis) for the whole range of replication errors (x-axis) ranging from zero errors (left) to errors in all dimensions (right). The combined effect (right y-axis) is obtained by multiplying both lines reflects the average long-run performance effects of replication errors. The results are based on 100.000 landscapes (with *N*=10 and *K*=2).

Figure 6: The Moderating Effect of the Template's Quality

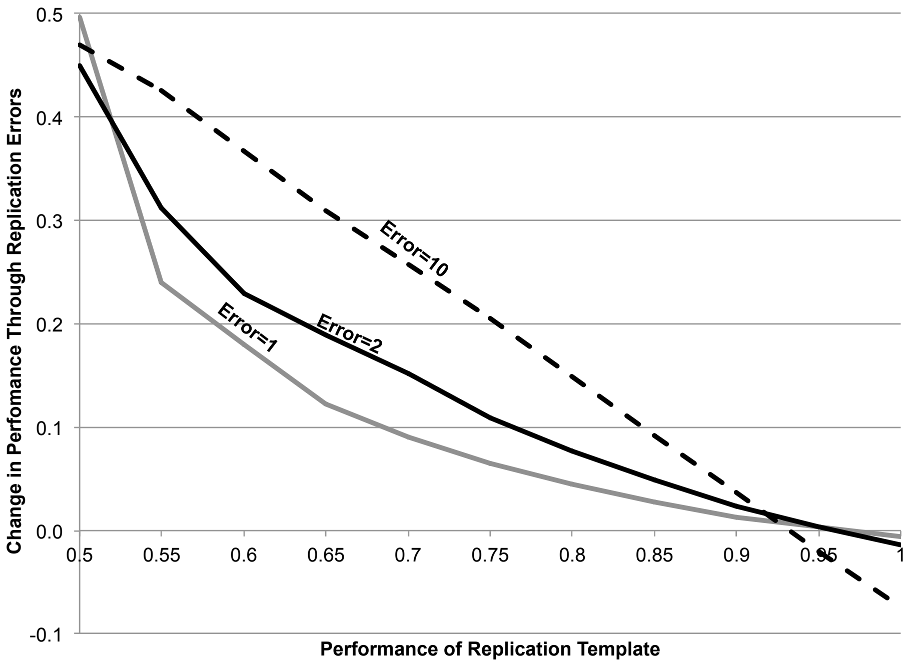

Notes: The figure reports the average long-run performance gains/losses (y-axis) for different sizes of replication errors between the final practice and the template as a function of the template's performance (x-axis). The results are based on 100.000 landscapes (with *N*=10 and *K*=2).

Figure 7a: Average Short Run Performance Effects

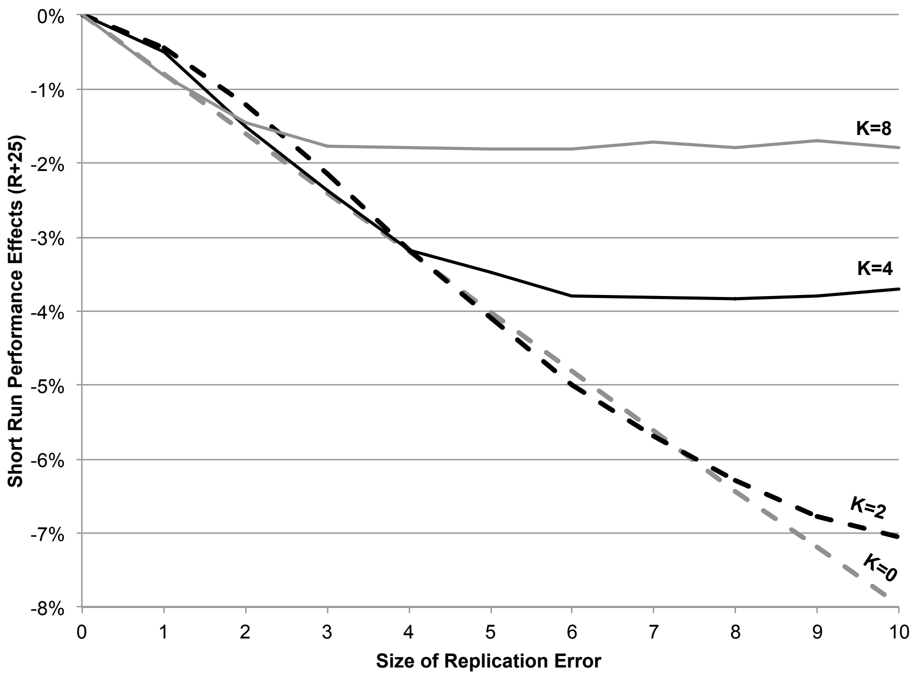

Notes: This figure reports the average short-run (t=R+25=225, y-axis) performance effects of replication errors (x-axis), ranging from zero errors (left) to errors in all dimensions (right) for different values of complexity. The performance effect is measured as the difference to the performance of the replication template in t=200. The results are based on 100.000 landscapes (with N=10).

Figure 7b: Average Long Run Performance Effects

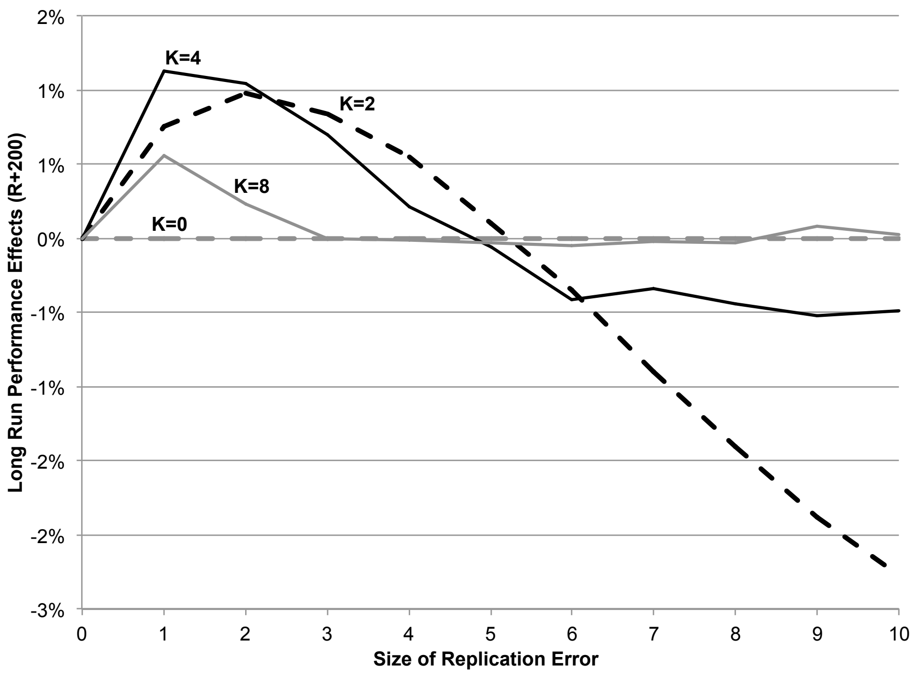

Notes: This figure reports the average long-run (t=R+200=400, y-axis) performance effects of replication errors (x-axis), ranging from zero errors (left) to errors in all dimensions (right) for different values of complexity. The performance effect is measured as the difference to the performance of the replication template in t=200. The results are based on 100.000 landscapes (with N=10).

Table 1: Summary of Results

Performance Effects of Replication Errors				
Time Horizon	**Short Run**		**Long Run**	
Error Size	**Small Errors**	**Large Errors**	**Small Errors**	**Large Errors**
Level of Complexity				
No Complexity	**-**	**- - -**	**0**	**0**
Low/ moderate Complexity	**-**	**- - -**	++	**-**
High Complexity	**- -**	**- -**	+	**0**

REFERENCES

Abernathy, W. J. (1978). The productivity dilemma: Roadblock to innovation in the automobile industry, Johns Hopkins University Press Baltimore, MD.

Aldrich, H. (1979). Organizations and environments, Stanford Business Books.

Almirall, E. and R. Casadesus-Masanell (2010). "Open versus closed innovation: A model of discovery and divergence." The Academy of Management Review **35**(1): 27-47.

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

Argyris, C. (1982). Reasoning, learning, and action: Individual and organizational, Jossey-Bass San Francisco.

Bartlett, C. A. and S. Ghoshal (1999). Managing across borders: The transnational solution, Taylor & Francis.

Baumann, O. and N. Siggelkow (2012). "Dealing with Complexity: Integrated vs. Chunky Search Processes." Organization Science.

Becker, M. C., T. Knudsen, et al. (2006). "Schumpeter, Winter, and the sources of novelty." Industrial and Corporate Change **15**(2): 353-371.

Benner, M. J. and M. L. Tushman (2003). "Exploitation, exploration, and process management: The productivity dilemma revisited." Academy of Management Review **28**(2): 238-256.

Campbell, D. T. (1965). "Variation and selective retention in socio-cultural evolution." Social change in developing areas: A reinterpretation of evolutionary theory: 19-49.

Csaszar, F. A. and N. Siggelkow (2010). "How much to copy? Determinants of effective imitation breadth." Organization Science **21**(3): 661-676.

Cyert, R. M. and J. G. March (1992). A behavioral theory of the firm, Wiley-Blackwell.

Denrell, J., C. Fang, et al. (2004). "From T-mazes to labyrinths: Learning from model-based feedback." Management Science: 1366-1378.

Denrell, J. and J. G. March (2001). "Adaptation as information restriction: The hot stove effect." Organization Science: 523-538.

DiMaggio, P. J. and W. W. Powell (1983). "The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields." American Sociological Review: 147-160.

Dorsch, J. (2000) "All eyes turn toward Southeast Asia: Semiconductor equipment vendors must chart a new market course as pure-play foundries gain power and influence." EDN DOI: http://www.edn.com/article/504318-All_eyes_turn_toward_Southeast_Asia.php.

Edmondson, A. C. (1996). "Learning from mistakes is easier said than done: Group and organizational influences on the detection and correction of human error." The Journal of Applied Behavioral Science **32**(1): 5-28.

El Akremi, A., K. Mignonac, et al. (2011). "Opportunistic behaviors in franchise chains: the role of cohesion among franchisees." Strategic Management Journal **32**(9): 930-948.

Fleming, L. and O. Sorenson (2001). "Technology as a complex adaptive system: evidence from patent data." Research Policy **30**(7): 1019-1039.

Frenken, K. (2000). "A complexity approach to innovation networks. The case of the aircraft industry (1909-1997)." Research Policy **29**(2): 257-272.

Ganco, M. and G. Hoetker (2009). NK modeling methodology in the strategy literature: Bounded search on a rugged landscape. Research Methodology in Strategy and Management. D. Bergh and D. Ketchen, Emerald Group Publishing Limited. **5:** 237-268.

Gavetti, G. (2005). "Cognition and hierarchy: Rethinking the microfoundations of capabilities' development." Organization Science **16**(6): 599-617.

Gavetti, G. and D. Levinthal (2000). "Looking forward and looking backward: Cognitive and experiential search." Administrative Science Quarterly: 113-137.

Holland, J. H. (1975). Adaptation in natural and artificial systems, University of Michigan press.

Iansiti, M. (1998). Technology integration: Making critical choices in a dynamic world, Harvard Business Press.

Katz, R. and T. 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**(1): 7-20.

Kauffman, S. A. (1993). The Origins of Order: Self Organization and Selection in Evolution, Oxford University Press, USA.

Kaufmann, P. J. and S. Eroglu (1999). "Standardization and adaptation in business format franchising." Journal of Business Venturing **14**(1): 69-85.

Knudsen, T. and D. A. Levinthal (2007). "Two faces of search: Alternative generation and alternative evaluation." Organization Science **18**(1): 39.

Kogut, B. and U. Zander (1992). "Knowledge of the firm, combinative capabilities, and the replication of technology." Organization Science: 383-397.

Levinthal, D. (2008). Explorations in the Role of Novelty in Organizational Adaptation: An Introductory Essay. Explorations in Organizations. J. March, Stanford University Press, USA**:** 98-105.

Levinthal, D. and H. E. Posen (2007). "Myopia of selection: Does organizational adaptation limit the efficacy of population selection?" Administrative Science Quarterly **52**(4): 586-620.

Levinthal, D. A. (1997). "Adaptation on Rugged Landscapes." Management Science **43**(7): 934-950.

Levitt, B. and J. G. March (1988). "Organizational learning." Annual Review of Sociology **14**(1): 319-338.

Lippman, S. A. and R. P. Rumelt (1982). "Uncertain imitability: An analysis of interfirm differences in efficiency under competition." The Bell Journal of Economics: 418-438.

Love, J. F. and A. W. Miller (1995). McDonald's: behind the arches, Bantam Books New York, NY.

March, J. G. (1991). "Exploration and exploitation in organizational learning." Organization Science: 71-87.

McDonald, C. J. (1998). "The evolution of IntelÕs Copy EXACTLY! technology transfer methodÕ." Intel Technology Journal **4**(1).

Meyer, J. W. and B. Rowan (1977). "Institutionalized organizations: Formal structure as myth and ceremony." American journal of sociology: 340-363.

Nelson, R. R. and S. G. Winter (1982). An evolutionary theory of economic change, Belknap press.

Polanyi, M. (1959). Personal knowledge: Towards a post-critical philosophy, University of Chicago Press Chicago, IL.

Porter, M. and N. Siggelkow (2008). "Contextuality within activity systems and sustainability of competitive advantage." Academy of Management Perspectives **22**(2): 34-56.

Reinhardt, A. (1997). "Intel's dreamers make room for a details man." Business Week: 125-128.

Rerup, C. (2004). "Imperfection, Transfer Failure, and the Replication of Knowledge: An Interview with Gabriel Szulanski." Journal of Management Inquiry **13**(2): 141-150.

Rivkin, J. and N. Siggelkow (2002). "Organizational sticking points on NK-landscapes." Complexity **7**(5): 31-43.

Rivkin, J. W. (2000). "Imitation of Complex Strategies." Management Science **46**(6): 824-844.

Rivkin, J. W. (2001). "Reproducing knowledge: Replication without imitation at moderate complexity." Organization Science: 274-293.

Rivkin, J. W. and N. Siggelkow (2003). "Balancing search and stability: Interdependencies among elements organizational design." Management Science **49**(3): 290-311.

Rivkin, J. W. and N. Siggelkow (2007). "Patterned interactions in complex systems: Implications for exploration." Management Science **53**(7): 1068-1085.

Ruef, M. (1997). "Assessing organizational fitness on a dynamic landscape: an empirical test of the relative inertia thesis." Strategic Management Journal **18**(11): 837-853.

Siggelkow, N. and D. A. Levinthal (2003). "Temporarily divide to conquer: Centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation." Organization Science **14**(6): 650-669.

Siggelkow, N. and J. W. Rivkin (2005). "Speed and search: Designing organizations for turbulence and complexity." Organization Science **16**(2): 101-122.

Simon, H. A. (1955). "A behavioral model of rational choice." The Quarterly Journal of Economics **69**(1): 99-118.

Simon, H. A. (1962). "The architecture of complexity." Proceedings of the American Philosophical Society **106**(6): 467-482.

Simon, H. A. (1996). The sciences of the artificial, MIT Press.

Szulanski, G. (2000). Appropriability and the Challenge of Scope: Banc One Routines Replication. The nature and dynamics of organizational capabilities. G. Dosi, R. Nelson and S. Winter, Oxford University Press, USA**:** 69.

Szulanski, G. (2000). "The process of knowledge transfer: A diachronic analysis of stickiness." Organizational Behavior and Human Decision Processes **82**(1): 9-27.

Szulanski, G. (2003). Sticky knowledge: Barriers to knowing in the firm, Sage Publications Ltd.

Szulanski, G. (2012) "Knowledge transfer: Use templates to pass on best practices, at least initially." INSEAD DOI: http://knowledge.insead.edu/contents/gabriel.cfm.

Szulanski, G. and K. Amin (2001). "Learning to make strategy: balancing discipline and imagination." Long Range Planning **34**(5): 537-556.

Szulanski, G. and R. J. Jensen (2004). "Overcoming stickiness: An empirical investigation of the role of the template in the replication of organizational routines." Managerial and Decision Economics **25**(6-7): 347-363.

Szulanski, G. and R. J. Jensen (2006). "Presumptive adaptation and the effectiveness of knowledge transfer." Strategic Management Journal **27**(10): 937-957.

Szulanski, G. and S. Winter (2002). "Getting it right the second time." Harvard Business Review **80**(1): 62.

Teece, D. J., G. Pisano, et al. (1997). "Dynamic capabilities and strategic management." Strategic Management Journal **18**(7): 509–533.

Terwiesch, C. and Y. Xu (2004). "The copy-exactly ramp-up strategy: Trading-off learning with process change." IEEE Transactions on Engineering Management **51**(1): 70-84.

Williams, C. (2007). "Transfer in context: Replication and adaptation in knowledge transfer relationships." Strategic Management Journal **28**(9): 867-889.

Winter, S. (1995). Resource-based and evolutionary theories of the firm: towards a synthesis. Four Rs of Profitability: Rents, Resources, Routines and Replication. C. Montgomery. Boston, Kluwer Academic Publishers**:** 147-178.

Winter, S., G. Szulanski, et al. (2011). "Reproducing knowledge: inaccurate replication and failure in franchise organizations." Organization Science: doi: 10.1287/orsc.1110.0663

Winter, S. G. (1975). "Optimization and evolution in the theory of the firm." Adaptive Economic Models: 73-118.

Winter, S. G. (2005). Developing evolutionary theory for economics and management. Great Minds in Management: The Process of Theory Development. M. Hitt and K. Smith, Oxford University Press**:** 209-546.

Winter, S. G. (2005). "Toward an evolutionary theory of production." The Evolutionary Foundations of Economics: 223-254.

Winter, S. G., G. Cattani, et al. (2007). "The Value of Moderate Obsession: Insights from a New Model of Organizational Search." Organization Science **18**(3): 403-419.

Winter, S. G. and G. Szulanski (2001). "Replication as strategy." Organization Science **12**(6): 730-743.