10.5465/AMBPP.2013.214

LEARNING THROUGH INACCURATE REPLICATION

HELMUT DIETL
Department of Business Administration
University of Zurich
Plattenstrasse 14, 8032 Zurich, Switzerland

DIRK MARTIGNONI
University of Zurich, Switzerland

MARKUS LANG
University of Zurich, Switzerland

ERIC LUCAS
University of Zurich, Switzerland

INTRODUCTION

Existing research points to the strategic importance of replicating organizational routines, business models, and best practices to ensure organizational growth and performance improvement (Nelson and Winter 1982; Teece, Pisano and Shuen 1997; Argote and Ingram 2000; Rivkin 2000; Winter and Szulanski 2001; Winter et al. 2011). While there is a strong consensus that accurate replication is hard if not impossible to achieve (Szulanski and Winter 2002; Rerup 2004), there is much less agreement on the performance implications of inaccuracies in the replication process (Winter and Szulanski 2001; Szulanski and Jensen 2004). In our study, we seek to enhance our understanding of why prior research has generated conflicting findings on the performance implications of replication errors. We demonstrate that contrary to popular managerial wisdom (Shoemaker 2011), any positive effect of inaccuracies in replication processes do not arise because they help firms to discover particular attractive solutions; instead, it is the fact that these inaccuracies make firms abandon particular unattractive solutions.

There is a large body of research on the performance implications of replication errors. On the one hand, considerable empirical evidence indicates that firms should seek to replicate a successful practice as accurately as possible. For example, Winter et al. (2011) find 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 error will be dysfunctional (Rivkin 2001). The enormous effort that firms such as Intel or McDonald's allocate to ensure replication accuracy (Reinhardt 1997; Iansiti 1998; McDonald 1998; Szulanski and Winter 2002) provides some further, indirect, evidence for the costs associated with replication errors.

On the other hand, mainly theoretical arguments suggest that there can also be benefits to inaccurate replication. For example, Becker, Knudsen and March (2006) portray replication errors as a source of novelty and variation. Szulanski and Winter (2002) and Winter (2005b) 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, Cattani and Dorsch 2007). Indeed, there seems to be a trade-off between accuracy and learning associated with the replication process, a trade-off that

10.5465/AMBPP.2013.214

is often labeled as "replication dilemma" (Szulanski and Amin 2001; Winter and Szulanski 2001; Szulanski and Jensen 2004).

The literature on adaptive and evolutionary systems, in contrast, is less circumspect: low levels of errors (Knudsen and Levinthal 2007), noise (Levinthal 1997, Denrell and March 2001), imperfections (Posen et al. 2013), perturbations (Winter, Cattani, Dorsch 2007), randomness (Fang and Levinthal 2009), mistakes (Rivkin 2000), or mutations (Nelson and Winter 1982) are thought to be beneficial in the long-run, while large errors can result in an "error catastrophe" (Kauffman, 1993). The popular management literature ascribes these positive effects of small errors to the fact that they may help firms to discover new and better solutions; indeed, these random errors are even sometimes depicted as "portals of discovery" (Shoemaker 2011, Chapter 5). Alexander Fleming's accidental discovery of Penicillin is a showcase for this logic (Horvitz 2002; Waller 2002; Shoemaker 2011). The academic management literature suggests that these positive effects can be explained by the structural characteristics of the underlying performance landscape. For example, according to Fang and Levinthal (2009), "…the introduction of noise, on average, favors the discovery of superior attractors whose bases of attraction are more extensive" (p.209).

The main contributions of our paper are, however, of a different kind: Using a standard NK performance landscape model (Levinthal 1997), we uncover the specific mechanisms through which replication errors may affect organizational performance. Moreover, we identify the necessary structural properties of the firm's task environment (performance landscape) for small replication errors to be beneficial. The central argument of our study is that if small replication errors improve long-run performance, this positive effect is primarily driven by the abandonment of particular unattractive solutions rather than the discovery of particular attractive solutions: a necessary condition for long-run performance to change through replication errors is that they induce a firm to abandon its current replication template or status quo (abandonment effect) and to discover a different and hopefully better practice (discovery effect).

We decompose the performance implications of replication errors into the discovery effect and the abandonment effect. Contrary to popular managerial wisdom, we demonstrate that small replication errors improve firm performance not because they induce a firm to discover high-performance practices but rather because they induce a firm to abandon low-performance templates. In fact, the performance of the practice that a firm discovers subsequent to an error is often still below average, yielding a negative discovery effect and a positive abandonment effect. We further show that the positive effects of small errors only materialize in task environments that exhibit what we call "high-performance resilience," i.e., good templates are more resistant to deviations than bad templates, or in more technical terms there is a positive correlation between the performance of a template and the size of its basin of attraction. Without this property or even a negative correlation ("low-performance resilience"), small errors do not have a positive average performance effect.

Given the importance of replication processes, it is not surprising that we are not the first to examine the implications of incomplete or inaccurate replication (or in more abstract terms, noise, random errors, mutations, or perturbations). Yet, much of our intuition on replication errors derives from studies on replication errors in populations of firms (Holland 1975; Aldrich 1979, Posen et al. 2013) or studies that assume repeated within-population learning (Csaszar and Siggelkow 2010). In our study, in contrast, we are interested in the average effect for one particular and isolated firm. 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

10.5465/AMBPP.2013.214

adverse effect on the firm that suffers from these errors (March 1991; Csaszar and Siggelkow 2010).

Most closely related to our work are Knudsen and Levinthal (2007) and Denrell and March (2001). They demonstrate that even an isolated firm can benefit from small (evaluation) errors. We complement and refine their findings by demonstrating that broader search or exploration induced through errors may not always be beneficial. Positive effects of small errors only arise if there is high-performance resilience, i.e. small errors are more likely to make a firm abandon a low performing alternative than an high performing alternative. In the case of task environments that exhibit low-performance resilience, small errors have negative long-run performance effects.

At first sight, our study may also seem at odds with some prior research on imitation and replication (Rivkin 2000; Rivkin 2001). Yet, these studies often assume that the template to be replicated is literally a "best practice" that cannot be improved anymore; i.e., 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). In addition, 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.

The remainder of the paper is structured as follows. Section 2 introduces our extensions of the standard NK landscape model. In Section 3, we analyze the performance implications of replication errors and uncover the specific mechanisms through which replication errors affect organizational performance. Finally, in Section 5 we briefly conclude by discussing our results and implications.

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 (task environment), (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. The (complex) performance landscape is a mapping of a firm's choice vector to performance. A firm is associated with a specific choice 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 or spatial replication is always possible.

The firm starts in period $t=0$ from a randomly determined choice 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 choice 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. This process may be interpreted as a search for better positions on a high-dimensional performance landscape ("hill climbing"). This kind of local search process implies that if the firm actually modifies its decision vector, performance will

10.5465/AMBPP.2013.214

always increase; a firm will never adopt a modification that decreases performance. Sooner or later, the firm will get stuck at a choice 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.

The local search process 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 choice vector in $t=R$ to which we refer as the "replication template" or simply "template" may not be replicated accurately. 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.

ANALYSIS

In our main experiment, firms are searching in performance landscapes of low complexity (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. This procedure ensures that the reported simulation results are statistically significant at the 1% level. In period t=200, the firm may or may not experience errors in replicating its practice. We observe the firms for 400 periods, which is sufficient to ensure that the model reaches steady state.

The solid line in Figure 2 displays average long-run performance (y-axis) for the full range of replication errors (x-axis), i.e., from accurate replication (zero errors) to completely inaccurate replication (ten errors). The dashed and dotted lines report the discovery effect and the abandonment effect, respectively, explained below.

Figure 1 about here

Consistent with the literature on adaptive and evolutionary systems (e.g., Denrell and March 2001, Knudsen and Levinthal 2007, Fang and Levinthal 2009, and Posen et al. 2013), the solid line shows that small replication errors can improve average long-run performance, although the inaccurately replicated template was already either a local or global optimum. A necessary condition for long-run performance to change 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 current template, long-run performance will not be affected. If it permanently abandons the replication template, the quality of the newly discovered practice relative to the abandoned template then determines whether performance improves or not through replication errors.

Hence, replication errors may affect long-run performance through two basic mechanisms: replication errors may make firms (1) discover particularly attractive (or unattractive) new practices ("discovery effect") and (2) abandon particularly unattractive (or

10.5465/AMBPP.2013.214

attractive) templates ("abandonment effect"). The relative strength of these two effects then determines whether replication errors have a positive or negative long-run performance effect. The discovery effect (dashed line) is measured as the average performance of the newly discovered practice relative to the average template, given that a replication error leads to the permanent abandonment of the replication template. The abandonment effect (dotted line) is measured as the performance of the average template relative to the average performance of those templates that are permanently abandoned. To compute the performance contribution, both effects are multiplied with the probability that a firm is permanently dislodged from its replication template.

Contrary to managerial wisdom, which depicts small errors as "portals of discovery" (Shoemaker 2011) for firms to discover new and better solutions, we find that the abandonment effect is always positive whereas the discovery effect is negative for the full range of replication errors on low complex landscapes (K=2). Hence, the positive effect of small replication errors is primarily driven by the abandonment effect rather than the discovery effect.

Our analysis shows that 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, for those firms that abandon their replication template, the inaccurately replicated template provides a worse starting position than any random position on the performance landscape. The positive effect of replication errors is driven by the fact that replication errors help firms abandon particularly unattractive templates (and subsequently discover below-average templates). In other words, average performance may improve with small replication errors because low performing firms improve from bad to below-average positions on the performance landscape.

Why do observe this positive abandonment and negative discovery effect? Given that replication errors do not affect the subsequent search process, the abandonment and discovery effects depend on the structural properties of the performance landscape. NK performance landscapes exhibit two important structural properties: a more or less strong correlation between the height of a peak and its basin of attraction and a more or less strong colocation of peaks, i.e. good peaks will be clustered around other good peaks (and bad peaks will be clustered around bad peaks). One can demonstrate that the abandonment effect can be linked to the correlation between the height of a peak and its basin of attraction whereas the discovery effect can be linked to the co-location of peaks in the landscape.

CONCLUSION

Our research suggests that prior research may have generated inconsistent findings on the performance implications of replication errors due to unobserved differences in the underlying task environments of the firms. We often tend to think that the complexity of the task environment is an important moderator (Szulanski and Winter 2002; Rivkin 2000; Winter et al. 2011): in complex environments such as Intel's manufacturing process (Reinhardt 1997; Iansiti 1998; McDonald 1998; Szulanski and Winter 2002), even small errors might have severe negative consequences while in less complex environments small errors might be a valuable source of exploration. Our analysis finds not support for this claim. Complexity only affects the

10.5465/AMBPP.2013.214

strength of this performance effect but never affects its *sign*. More recent research (Gavetti 2011) may also be interpreted as suggesting that the extent of peak clustering (co-location) in the task environment may moderate the performance implications of errors. Our study, in contrast, suggests that co-location has no effect on the performance implications of small errors.

Instead, we demonstrate that the primary moderator is performance-resilience. For example, even small replication errors can be costly in task environments that exhibit low-performance resilience. Studies that point to positive performance effects of small replication errors might be based on firms operating in task environments of high-performance resilience.

FIGURES

Figure 1: Decomposing the Effect of Replication Errors

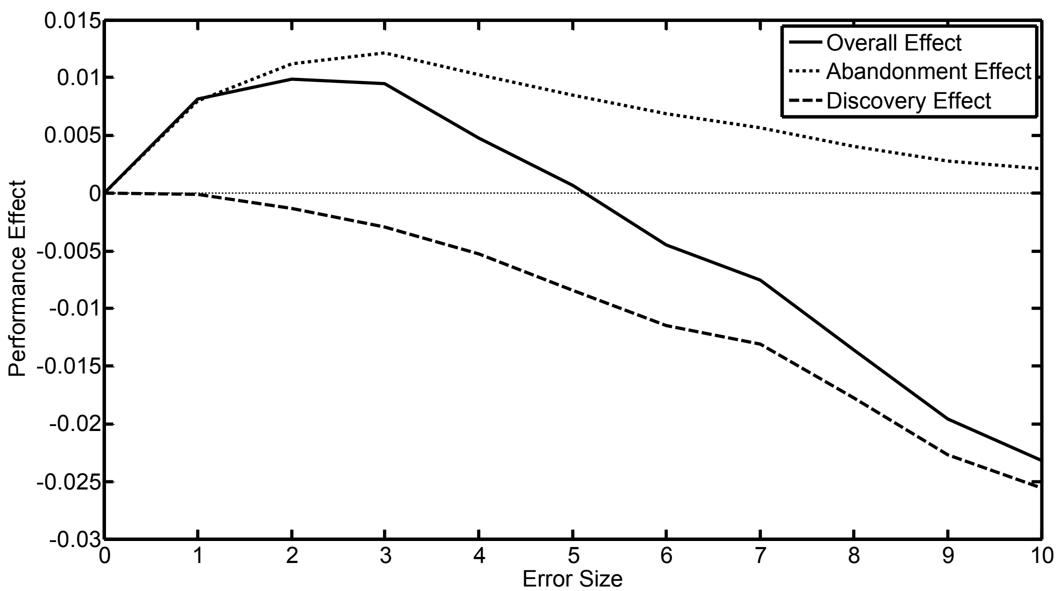

REFERENCES AVAILABLE FROM THE AUTHOR(S)