



Jennifer Kunz (2011)

Group-Level Exploration and Exploitation: A Computer Simulation-Based Analysis

Journal of Artificial Societies and Social Simulation 14 (4) 18

<<http://jasss.soc.surrey.ac.uk/14/4/18.html>>

Received: 23-Apr-2010 Accepted: 30-May-2011 Published: 31-Oct-2011



Abstract

Organisational research has studied the tension between exploration and exploitation for years. In essence, this body of research agrees on the necessity of a balance between explorative and exploitative processes to prevent an organisation from falling into a learning trap. Thus, to enhance the active management of this balance in organisations, a deeper theoretical understanding of the factors that influence the development of exploration and exploitation has to be gained. One of the recently discussed factors is the interplay between exploration and exploitation on different organisational levels. This paper picks up this discussion. It provides an in-depth, computer simulation-based analysis of the performance of organisational types with varying degrees of within-group and between-group exploration and exploitation in situations of different degrees of task complexity. The findings indicate that a high share of between-group processes as compared to within-group processes positively influences the organisational performance level and that dependent on task complexity the optimal share of exploration and exploitation varies.

Keywords:

Organisational Learning, Experience-Based Learning, Exploration, Exploitation, Knowledge Management, Genetic Algorithms



Introduction

- 1.1 After environmental shocks, like the recent financial and economic crisis, organisations have to perform considerable adaptation processes. In these situations they can either try to change the way they act in a rather incremental way, i.e. undergo a phase of exploitation, or they can try to change their behaviour more radically, i.e. perform a process of exploration. In most instances, organisations will exhibit a mixture of both adaptation strategies, as they are not one monolithic entity but they consist of various subunits on several organisational levels that can at least partly choose their own adaptation strategy. Here the question rises of what would be an optimal level of exploration and exploitation within and between the mentioned subunits to stabilise an organisation after such an environmental shock.
- 1.2 After several decades of research, literature on exploration and exploitation agrees on the necessity of a balance between explorative and exploitative processes to prevent an organisation from falling into a learning trap (March 1991; Babuji and Crossan 2004; Auh and Menguc 2005; Espedal 2008). Yet, the questions of how this balance can be achieved and of what actually would be an optimal equilibrium between exploration and exploitation still are topics of discussion (Auh and Menguc 2005; Gupta et al. 2006; Jansen et al. 2006; Greve 2007; Sidhu et al. 2007; Li et al. 2008; Fang et al. 2010).
- 1.3 Thus, to bring forward theory building and to enhance the active and effective management of exploration and exploitation in organisations especially after severe environmental shocks, a deeper theoretical understanding of the different possible explorative and exploitative processes within organisations and their respective effects on organisational performance still is necessary. Hence, recent research has started to investigate the influence of different organisational levels on the effective interplay between exploration and exploitation (e.g., Holmqvist 2004; Taylor and Greve 2006; Belderbos et al. 2010; Fang et al. 2010).
- 1.4 Since groups comprise the building blocks of organisations, the differences in performance outcomes of explorative and exploitative processes on this level are of special interest to an integrated theory of exploration and exploitation. This even more so as the group level operates as a mediator between explorative and exploitative processes on the individual and on the organisational level. However, beside the studies by Taylor and Greve (2006) and Fang et al. (2010) literature provides relatively little results regarding the specifics of group-level exploration and exploitation. The paper picks up this research gap. It analyses within-group and between-group exploration and exploitation under the condition of an environmental shock, as then an organisation is in a critical state, in which learning processes are crucial. Moreover, the study incorporates two further aspects: First, as previous research indicates (Thompson 1967; March 1991; Miller et al. 2006), task complexity might have an impact on the performance outcomes of different explorative and exploitative regimes. Second, both exploration and exploitation are always based on pre-existing knowledge. Thus, also the breadth of the existing knowledge pool on which these processes operate might have an impact on their outcomes. Therefore, the following analysis incorporates both contingent aspects.
- 1.5 Literature on ambidextrous organisations points to the fact, that within organisations some organisational units might be exclusively dedicated to exploration while others only perform explorative activities (Tushman and O'Reilly 1996; Benner and Tushman 2003). Consequently, it is reasonable to examine the performance of such isolated regimes. Nevertheless, in most instances, organisations are characterised by a mixture of exploration and exploitation. Therefore, the following study also is dedicated to the performance of organisations that have different degrees of within-group and between-group exploration and exploitation. In sum, the paper explores the questions
 1. of how within-group and between-group explorative and exploitative processes differ with respect to their performance outcomes and
 2. of how organisations that are characterised by different degrees of each of these processes perform under different degrees of task complexity and different breadth of knowledge pools after an environmental shock.
- 1.6 With respect to the aim of the paper, both the finally attainable performance levels at the end of a learning phase and the evolution of the performance levels over time during this learning phase are of interest. Therefore, computer-based simulation is applied as research method, i.e. the following study uses an agent-based simulation model carrying out a conceptual experiment (Axelrod 1997; Epstein 1999; Gilbert and Terna 2000).
- 1.7 The results of the simulation study add the following aspects to the existing literature: First, they point to significant differences in the isolated effectiveness of the *within-group and between-group explorative and exploitative processes*. The introduction of between-group processes changes the *achievable performance level* and the *evolution* of performance over time of both exploration and exploitation. Second, the findings exhibit a superiority of organisations that allow for a high degree of between-group processes as compared to organisations which are characterised by a low degree of them. Third, the results suggest different equilibria between exploration and exploitation dependent on task complexity.
- 1.8 The remainder of the paper is structured as follows. Section 2 consists of a brief literature review. Section 3 presents the main assumptions regarding the modelled groups and discusses different types of exploration and exploitation. In order to analyse the performance of these types, the model is transferred into a multi-agent simulation, described in Section 4. Section 5 presents the results of the simulation experiments. Section 6 is dedicated to the major findings and implications of the paper, and identifies future research possibilities.



Literature on exploration and exploitation

- 2.1 The discussion of exploration and exploitation was fostered by the seminal work of March (1991). Based on the conclusion that exploration is driven out by exploitation, March pointed out that "maintaining an appropriate balance between exploration and exploitation is a primary factor in system survival and prosperity" (March 1991, p. 71).
- 2.2 Subsequent research picked up this finding and came to the consensus that a balance of exploration and exploitation is of importance, but it has to be adapted to the situational circumstances of an organisation and that in practice, organisations can actually succeed to achieve this balance (Levinthal and March 1993; Siggelkow and Levinthal 2003; Auh and Menguc 2005; Rodan 2005; Hodgson and Knudsen 2006; Lavie and Rosenkopf 2006; Lazer and Friedman 2007; Uotila et al. 2009).

- 2.3 Regarding such situational circumstances research especially focuses on the importance of *environmental dynamic*. Lant and Mezias (1992) analyse the influence of environmental change on flexibility (exploration) and stability (exploitation). Sidhu et al. (2004) examine the influence of dynamic environments on exploration and exploitation exhibiting a positive impact of a more dynamic environment on exploration. The results of Jansen et al.'s (2006) study indicate a positive, moderating role of environmental dynamism on the relationship between exploratory innovation and financial performance, and a negative, moderating role on the relationship between exploitative innovation and performance. Sidhu et al. (2007) examine different search types linked to exploration and exploitation and find an impact of environmental dynamism on them.
- 2.4 Lately, research interest additionally is shifting towards the *importance of different organisational levels* regarding exploration and exploitation: Holmqvist (2004) provides a framework linking inter- and intra-organisational experiential learning with exploration and exploitation. Siggelkow and Rivkin (2006) find that in contrast to a generally accepted notion increased exploration at lower organisational levels can reduce organisational exploration and decrease performance in certain environments. Taylor and Greve (2006) analyse the impact of team compositional factors on variance-enhancing behaviour (i.e. exploration) and on mean performance-enhancing behaviour (i.e. exploitation) in the comic industry. Miller et al. (2006) extend March's model by several features, especially by a direct interpersonal learning dimension allowing for a further differentiation of March's results. Mom et al. (2007) focus on explorative and exploitative processes on the management level exhibiting different influences of top-down, bottom-up and horizontal knowledge inflows. In an empirical study Belderbos et al. (2010) show different effects of inter-firm and intra-firm exploration and exploitation on a firm's financial performance. The findings of the computer-based analysis by Fang et al. (2010) suggest that the division of an organisation in groups with moderate levels of linkages might result in the highest organisational performance.
- 2.5 As mentioned in Section 1, this paper focuses on exploration and exploitation on the within-group and between-group levels in a situation in which the studied organisations are confronted with an environmental change resulting in the need to adapt to new environmental requirements. Hence, the following study can be seen as part of the latter two research streams. As previously discussed, literature already offers results regarding single aspects of the performance of explorative and exploitative regimes on the different organisational levels. Yet, to date it lacks a coherent theory. The following study shall bring theory building forward by deriving in a structured manner stylised mechanisms that point out the differences between within-group and between-group explorative and exploitative regimes. On the one hand, it extends the literature by *separating different kinds of explorative and exploitative processes* and by studying them *in a controlled setting*. In this sense, in contrast to empirical studies, in which groups to some extent are always part of all processes, the study allows for an isolated examination of each of these regimes. This offers the possibility to derive the specific influence that each of them has on the development of performance outcomes. On the other hand, the following analyses study the performance outcomes of organisations that are characterised by *varying shares of these explorative and exploitative processes*. Thus they provide findings regarding the question of what are optimal equilibria between explorative and exploitative processes on the within- and the between-group level.
- 2.6 Literature provides a wide range of contingent aspects which could be incorporated into the study. For reasons of clarity, the analyses shall focus on only two of them. One major aspect of the successful creation of new knowledge and of learning is pre-existing knowledge. As in general, pre-existing knowledge determines the possibility to generate, integrate and apply new knowledge (Cohen and Levinthal 1990; Bower and Hilgard 1998). Therefore, the first aspect to be incorporated into the analysis is the *breadth of the knowledge pool* on which each group can operate. Furthermore, tasks can vary in *complexity* (Thompson 1967). When a task is relatively simple, a group can amass or generate the knowledge to do it alone. Yet, the more complex a task is, the more beneficial the possibility of knowledge sharing becomes. Hence, complexity of tasks shall be considered as the second contingent factor.



Model structure

Groups

- 3.1 The following analyses concentrate on explorative and exploitative processes within and between groups in an organisational context. Therefore, they focus on one organisation consisting of several organisational groups. These groups might be departments, communities-of-practise or project teams.
- 3.2 Each organisational group possesses a knowledge pool containing a set of separable knowledge units (Lee and Ahn 2007). Following the concept of bounded rationality (e.g., Simon 1977, 1991), that affects individuals and therefore also groups, these knowledge pools cannot contain to any possible situation a perfectly fitting knowledge unit, i.e. the knowledge pools are always limited to a subset of problem solutions. However, individuals and, hence, groups can perform knowledge creation processes to add a further solution to the existing knowledge pool. Yet, as another consequence of bounded rationality, these processes may not necessarily lead instantaneously to an optimal solution, but may require further learning processes. The learning concept used here is thus developed in the tradition of the literature on experience-based learning (e.g., Cyert and March 1963; March and Olsen 1975; Herriott et al. 1985; Levitt and March 1988; Carley 1992; Levinthal and March 1993).

- 3.3 To keep the results tractable, the analyses abstract from any strategic behaviour. Hence, it is assumed that the groups are willing to engage in explorative and exploitative processes and that they are aligned to organisational goals by an appropriate incentive system. This assumption is critical, for in practice strategic behaviour, like the reluctance to transfer expert knowledge in order to protect one's status, is an important issue. However, this paper focuses on the performance of different explorative and exploitative processes with respect to their *capacity to enhance performance improvements*, and not on their ability to cope with strategic behaviour.

Stylised learning regimes

- 3.4 In literature, exploration and exploitation are defined differently, as the following examples show: According to March (1991, p. 71), exploration incorporates "search, variation, risk taking, experimentation, play, flexibility, discovery, innovation," while exploitation is characterised by "refinement, choice, production, efficiency, selection, implementation, execution." Sidhu et al. (2007) focus on a more specific definition that cites the difference between local and non-local search. Jansen et al. (2006) differentiate between incremental improvements and radical innovations. The characterisation by Levinthal and March (1993), followed by Rothaermel (2001), describes exploration as "the pursuit of knowledge, of things that might come to be known" and exploitation as "the use and development of things already known" (Levinthal and March 1993, p. 105).
- 3.5 However, irrespectively of the differences in detail, most authors associate exploitation with a process of refinement, while they characterise exploration as a process of innovation. Transferred to the learning context that is of importance to this paper, this means, that exploitation operates on the existing knowledge pools of the groups, while exploration creates new knowledge units. Here, it is useful to introduce a more differentiated concept of the knowledge units. Actually, they comprise two components, that are closely linked to each other but that are also handled differently by exploration and exploitation. On the one hand, the groups possess a set of knowledge specifying how to do something, e.g. they can exhibit a range of different actions, like dealing with a customer in a very friendly or in a reserved way. On the other hand, the groups have *beliefs* regarding the appropriateness of each of these actions in a certain situation. Hence, each knowledge unit contains an action and a value linked to this action. The group can change either the value, i.e. after a bad experience it can reduce the expected value of a given action for a specific situation. In this case, the existing knowledge pool is refined in the sense of exploitation. Alternatively, the group can try an action that is new to the existing pool of actions. This is a process of innovation and thus exploration. In contrast to exploitation, this process comprises both, the introduction of a new action and the testing of the appropriateness of this action. This explorative creation process can operate on *different degrees* of innovation, as a new action might be only a bit different from existing ones or very different. Therefore, the following study incorporates two types of exploration, one that exhibits a medium degree of innovation and one that has a high degree.
- 3.6 Both, exploration and exploitation can occur within a group or through the interaction between groups. In the first case, group members use actions known to the group or try to find new solutions based on solely the ideas of the group members. In the second case, groups exchange existing actions with other groups or try to develop new solutions jointly with other groups. The former processes shall be called within-group exploitation and exploration, the latter between-group exploitation and exploration (cf. Figure 1).

Figure 1. Between-group and within-group exploration and exploitation (cf. for a similar classification with respect to the inter- and intra-firm level Belderbos et al. 2010)

- 3.7 In practice, *within-group* exploitation is applicable, if the need for any knowledge creation and knowledge transfer between groups is negligible and group members are perfectly prepared to cope with their portfolio of tasks. Alternatively, this process can occur, if the groups are prevented from exchanging knowledge by organisational barriers and at the same time they are reluctant to deviate from their existing solutions preventing them from exploration. In either case, groups only operate on the existing group-specific knowledge pool. This learning regime constitutes a stylised and extreme form of task performance often found in strictly regulated activities, like financial accounting. The group members and hence the group as a whole have to follow fixed procedures and rules. Knowledge creation to pursue innovative solutions for any given task is prohibited. Learning only happens within the range of regulations as process of refining the evaluation of the appropriateness of the existing actions in given situations.
- 3.8 However, exploitation can also take place between groups, when members of different groups communicate with each other and exchange knowledge. *Between-group exploitation* groups do not create actions that are completely new, but they improve their knowledge pools as best-practice spreads from group to group. This kind of learning regime can be observed, for example, when companies initiate benchmarking activities. Benchmarking projects identify best practices and lead to their diffusion within organisational units.
- 3.9 Through *within-group exploration*, groups enlarge their knowledge pool by activities that generate new knowledge. However, the groups operate in isolation, precluding knowledge transfer between groups. This type of learning might be typical among organisational groups that operate in a field, in which highly specialized knowledge is necessary and the transfer between groups is not useful.
- 3.10 Finally, exploration can also occur between groups. This *between-group exploration* is enhanced by a working situation in which groups not only exchange experiences but jointly try to come up with new solutions to existing problems, e.g. during product development activities that involve different departments.
- 3.11 These four learning regimes differ in the way how they handle and further develop the groups' knowledge pools. The exploitative regimes stick to the existing actions. However, in case of the between-group type at least the transfer of actions between groups is possible. Hence, in this case a group's action pool can change, while in the within-group version the action pool remains the same. Here, only the judgement of the appropriateness of the different actions regarding a specific problem can be altered. In contrast, the explorative regimes lead to changing action pools in the between-group and the within-group case. However, this change is different from the one induced by between-group exploitation, since exploration leads to actions that can be new to the whole organisation. Moreover, between-group exploration operates on a larger basis to construct these new actions than within-group exploration.



Simulation model

Elements of the simulation

Environment and task

- 4.1 The following simulation model is structured as an agent-based model, where the groups are represented by agents, i.e. "autonomous decision making entities" (Bonabeau 2002, p. 7280). Agent-based models are especially suitable to model situations in which emergent phenomena might occur (Bonabeau 2002). Therefore, this modelling approach seems to be a reasonable approach for the intended study. To model the learning processes the following simulation partly uses so-called genetic algorithms (Holland 1973; Holland 1995; Chattoe 1998; Dawid and Kopel 1998). They are selected as they had already been implemented to model learning in different areas of socio-economic research (Gilbert and Troitzsch 1999) and allow for the consideration of the derived theoretical assumptions.
- 4.2 In contrast to some other simulation-based studies in the research area of exploration and exploitation, like those of March (1991) and Miller et al. (2006), the groups here directly interact with the environment, as is the case in marketing or service departments. The focus of this study is how well an organisation containing several groups can adopt its actions to new environmental requirements or configuration after a shock. These requirements can be interpreted as an expected action.
- 4.3 The environmental configuration is implemented as a binary string made of n positions c_k with $n \in \mathbb{N}$, $k \in \{1, \dots, n\}$ and $c_k \in \{0, 1\} \forall k$, e.g. (101010) is one possible environmental configuration, where $n = 6$. As the focus of the analysis lies in groups of one organisation that faces one environment, the analysed groups are confronted with the same environmental requirements (i.e. the same configuration of this string).
- 4.4 As the task that each group is to perform is defined through the environmental configuration, the length of this string can be identified as the degree of task complexity. The longer this string, the more information bits have to be learnt by each group, and the more complex the task becomes.
- 4.5 The simulation can be run for a preset amount of periods. At the beginning of each period, each group has to apply a suitable action, which is codified through a bit string of the same length as the environmental one. At the end of each period, the environment compares these actions of the groups with its configuration and gives feedback to each group according to the similarity between its action and the environmental configuration. This feedback f is calculated as the quotient of the number of correct digits c divided by the number of all digits n :
- $$f = c / n. \quad (1)$$
- 4.6 This feedback lies between 0 and 1 and can be interpreted as either a degree of correctness or the probability of the group giving the correct answer. The first holds in cases, where task performance not only comprises right or wrong but also a scaling between this like a very good to a very bad task accomplishment. The latter fits situations in which only a correct or an incorrect answer is possible. In both cases, the higher the quotient, the more aspects of the environment are correctly known by the group. Additionally, the scaling of feedback of 0 to 1 allows for comparison of situations of different task complexity. This performance measure is similar to the one used in Ren et al. (2006), where the degree of performance quality is computed via the degree of knowledge which an agent has regarding her task.

Groups and organisation

- 4.7 To conduct the following analysis, the regimes are operationalised by defining groups as decentralised knowledge pools. These pools can be changed locally through the within-group regimes, or globally through between-group exploitation and exploration.
- 4.8 At the beginning of each *simulation run*, each group is provided with a set of y knowledge units containing actions that are implemented as binary bit strings of the same length as the environmental one and values assigned to these actions. The more actions a group has, the broader its knowledge pool is. The strings are generated randomly (i.e. the groups are heterogeneous regarding their initial knowledge pools). The number of actions remains constant during the whole simulation run (i.e. when a new action is generated, an old one is removed from the set). The old actions are forgotten. This process of forgetting aims at actions which had been less successful in the past.
- 4.9 In the following simulation runs, each organisation consists of 5 groups. For each simulation run, this number is constant.
- 4.10 The focus of the study lies in the development of organisational performance over time. However, since not the organisation as a whole but each group performs tasks and receives feedback from the environment, organisational performance has to be derived from group performances. Hence, organisational performance is calculated as an average of group performances following Ren et al. (2006).
- 4.11 The following analyses concentrate on both, the isolated performance of each regime and the performance of a combination of regimes. Regarding the isolated regimes six settings are analysed: 1) within-group exploitation, 2) between-group exploitation, 3) within-group exploration with a medium degree of innovation, 4) between-group exploration with a medium degree of innovation, 5) within-group exploration with a high degree of innovation and 6) between-group exploration with a high degree of innovation. In addition to these six settings, in which an organisation follows only one of the discussed regimes, further settings are simulated to examine the effect of a combination between different regimes on performance. To show both the effect of various degrees of exploration and exploitation, and the impact of different degrees of within-group and between-group processes the study comprises 12 further settings. In the study only explorative processes with the same degree of innovativeness are combined. Table 1 summarises the 18 organisational settings. (Settings Org 8 and Org 14 and settings Org 11 and Org 17 are identical, respectively. For reasons of clarity, they received two references.)
- 4.12 In order to allow for a structured presentation of the simulation results they are clustered in three experiments: The first experiment deals with the performance of the six pure regimes. The second experiment provides results of organisational types that exhibit different shares of within-group and between-group processes. The third experiment concentrates on organisational types with varying shares of exploration and exploitation.

Focus of analysis	Name	Abbreviation	Share of exploration/ exploitation (in %)	Share of within- group/ between- group processes (in %)	Degree of innovativeness of explorative processes
<i>Experiment 1:</i> Isolated regimes	Org 1	ET-WG-0	0/100	100/0	-
	Org 2	ET-BG-0	0/100	0/100	-
	Org 3	ER-WG-M	100/0	100/0	medium
	Org 4	ER-BG-M	100/0	0/100	medium
	Org 5	ER-WG-H	100/0	100/0	high
	Org 6	ER-BG-H	100/0	0/100	high
<i>Experiment 2:</i> Influence of different shares of within-group and between-group processes	Org 7	WG25-M	50/50	25/75	medium
	Org 8	WG50-M	50/50	50/50	medium
	Org 9	WG75-M	50/50	75/25	medium
	Org 10	WG25-H	50/50	25/75	high
	Org 11	WG50-H	50/50	50/50	high
	Org 12	WG75-H	50/50	75/25	high
<i>Experiment 3:</i> Influence of different shares of exploration and exploitation	Org 13	ER25-M	25/75	50/50	medium
	Org 14	ER50-M	50/50	50/50	medium
	Org 15	ER75-M	75/25	50/50	medium
	Org 16	ER25-H	25/75	50/50	high
	Org 17	ER50-H	50/50	50/50	high
	Org 18	ER75-H	75/25	50/50	high

Table 1. Simulated organisational types. The abbreviations have been defined as follows: ER = ExploRation; ET = ExploiTation, WG = Within-Group, BG = Between-Group, 0 = no innovativeness, M = Medium innovativeness, H = High innovativeness. In experiment 2, the organisations have the same degree of exploration and exploitation and the share of between-group processes equals 1-share of within-group processes. Hence, their abbreviation can be based solely on the degree of within-group processes and degree of innovativeness. An analogous definition is applied to experiment 3.

Steps of the simulation

Generation of random seed and pre-valuation

- 4.13 At the beginning of each simulation run, 18 organisations of the same number of groups are generated. Each organisation is assigned one of the 18 organisational types to which it sticks over the whole simulation run. The environmental configuration and the initial knowledge pools of the groups are generated randomly. Each organisation is endowed with the same pools of actions (i.e. groups within an organisation differ in their initial actions and values, but the organisations contain the same groups to make the results comparable). Since the study focuses on

adaptation processes after an environmental shock, before the actual learning phase starts, the actions are valued *randomly*, i.e. when the learning processes start, the groups do not know how well each action fits to the environmental configuration. The following paragraphs discuss the implementation of the four basic learning regimes.

Within-group exploitation

- 4.14 In the case of within-group exploitation, the groups have to work on the basis of the actions and their values in their own knowledge pools given at the beginning of the simulation run. In each period each group chooses the action from its pool with the highest value f and presents it to the environment. The environment values it according to (1). The groups thereafter perceive their values as feedback and assign it to the used action to revalue it. Hence, each group improves its performance by testing existing actions and trying to select the one that best fits to the environmental code. This learning is of the exploitative type, because it elaborates on existing actions without changing them. It is within-group, since there is no transfer of actions among groups. Figure 2 illustrates the selection process for Group 1 in an arbitrary period.

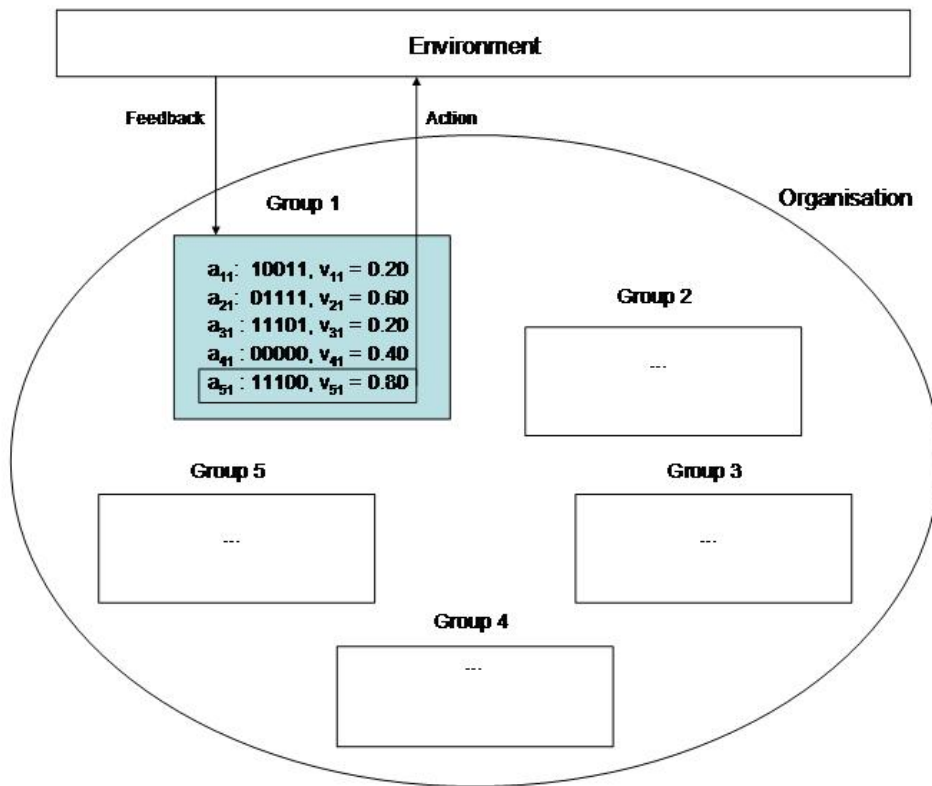


Figure 2. Selection of an action in case of within-group exploitation (a_{ik} = action i of group k , v_{ik} = value of action i of group k)

Within-group exploration

- 4.15 In case of within-group exploration, each group continuously changes its knowledge pool by creating, testing and introducing new actions. However, as in case of within-group exploitation, the groups operate in isolation from each other. Knowledge transfer does not occur.
- 4.16 This learning regime is implemented using genetic operators. This technique has been selected, because it mirrors important aspects of human learning process, which are relevant to this study. As stated earlier, the analysis is in the tradition of the experience-based learning literature. Hence, learning is conceptualised as the circle between creating and using actions, receiving environmental feedback and re-evaluating the used actions according to this feedback. This circle combines active (knowledge creation) and reactive (revaluation) elements of action taking and learning. Moreover, experience-based learning also comprises the fact, that knowledge creation always is based on existing knowledge. The balance between action and reaction on the one hand, and knowledge creation based on existing knowledge on the other, are both used in the concept of genetic algorithms.
- 4.17 Genetic algorithms themselves contain the interplay of exploration and exploitation (Holland 1973; Booker 1987). In this sense, explorative regimes are themselves not free of exploitative elements as a consequence of the path-dependent modelling of knowledge creation: new knowledge is always based in some sort on old knowledge. In this paper, however, the terms "exploitation" and "exploration" are reserved for the knowledge-handling activities concerning complete actions, not the provenance of their parts.
- 4.18 In essence, exploration is modelled as follows. At the beginning of each period, a group selects some actions from its knowledge pool as the basis to create new actions. Following the idea behind genetic algorithms, it selects the two best actions in terms of their feedback values f . When there are more than two actions with the same values, two actions are selected randomly from this subset. This is the parent population.
- 4.19 Knowledge creation happens via the genetic operators (one-point) crossing-over and mutation (Holland 1973, 1995; Dawid and Kopel 1998). In the first step, the actions undergo the crossing-over process, with a randomly-selected crossing-over point. Thereafter, each digit of every new created action mutates (is flipped into its opposite) with a probability p (mutation rate). The mutation rate characterises the degree of innovativeness. In the following study a mutation rate of 1% is applied in case of highly innovative exploration and a rate of 0.1% is used in case of medium innovativeness. These parameter values are somewhat arbitrary. However, when mutation rates become too large, the whole genetic algorithm will break down, because the spontaneous change rates destroy the improvement process. Therefore, literature suggests using mutation rates in the range that is also implemented in this model (e.g., Dawid and Kopel 1998).
- 4.20 The newly-generated child generation gets the expected value f of its parents according to the following equations (Gilbert et al. 1995; Kennedy and Eberhart 2001) (f_c = value of the child, f_1 = value of the first parent, f_2 = value of the second parent, m = crossing-over point, n = length of action, each child gets the first m digits from the first parent, and the $n-m$ digits from the second parent):

$$f_c = (m / n) * f_1 + ((n-m) / n) * f_2 \quad (2)$$

- 4.21 The child action with the highest expected value, then, is selected for presentation to the environment, which values it. Thereafter, the group introduces this action into its knowledge pool by replacing the action with the lowest value and assigning the environmental value to this new action. However, if this action already exists by chance in the knowledge pool, no other action is replaced; only the existing one gets the new value. This procedure is contrary to the classical genetic algorithms. Yet, it has been introduced to model the learning processes better, as exactly the same piece of knowledge only is "remembered once." Figure 3 illustrates the selection process for Group 1 in an arbitrary period.

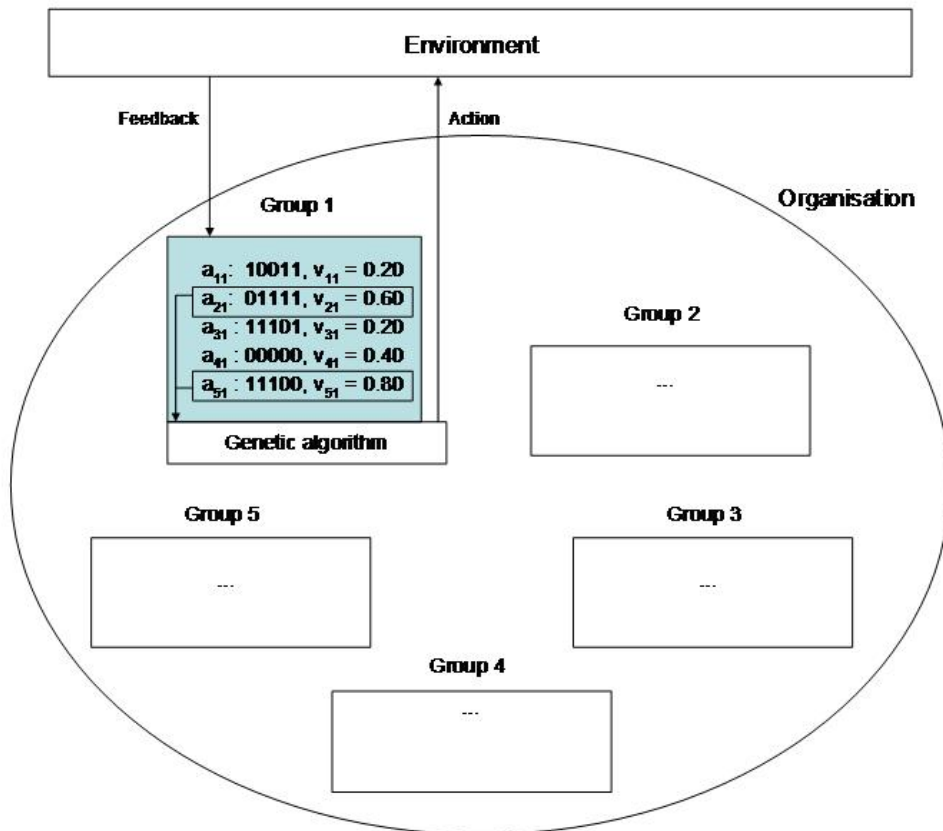


Figure 3. Selection of an action in case of within-group exploration (a_{ik} = action i of group k , v_{ik} = value of action i of group k)

Between-group exploitation

- 4.22 In between-group exploitation, the groups transfer knowledge among each other. There are many ways in which this connection can be modelled. In this study, one type will be selected, based on the results of interpersonal learning in the literature. People tend to prefer learning from people who are spatially near to them; this leads to spatial myopia (Levinthal and March 1993; Miller et al. 2006). Since groups are composed of individuals, it is reasonable to assume that they suffer from the same myopia. The influence of this tendency compared to more distant interpersonal learning has already been analysed by Miller et al. (2006) and is beyond the scope of this study. The present analysis therefore focuses on learning between adjacent agents, i.e. groups that are located close together.
- 4.23 To implement this assumption, the groups of an organisation are located on an edgeless grid as in Miller et al. (2006). However, here each group has two neighbours, not four, i.e. the groups are positioned along a ring structure. This grid structure is selected for the following reason: the analysis will deliver insights into the qualitative differences between within-group and between-group exploration and exploitation. In order to render the results of each regime comparable and to concentrate on the steps from within-group to between-group processes, the simplest grid allowing for an interaction between groups should be used. The simplest structure which enhances the bilateral transfer of knowledge allows each group to exchange knowledge bilaterally with its immediate neighbours. If each group only exchanges knowledge bilaterally with one neighbour, the knowledge cannot flow through the whole organisation, but one will get pairs of groups. In contrast, if each group only transfers unilaterally knowledge to, for example, its neighbour on the left, diffusion will be possible. However, this blocks a two-sided knowledge transfer.
- 4.24 Under this learning regime, at the beginning of each period, each group chooses its best action on the basis of the feedback values f . It then transfers this action to its neighbours and simultaneously receives one action from each of them. Thereafter, each group has a set of three actions: one action from its neighbour on the left, one action from its neighbour on the right and the action with the highest value from its own knowledge pool. From this set it selects the action with the highest value. The chosen action then is valued by the environment and thereafter introduced in the group's knowledge pool, replacing a weak action. Alternatively, if the action already exists, only the feedback value of this existing action is changed to the actual environmental feedback. Figure 4 illustrates the selection process for Group 1 in an arbitrary period.

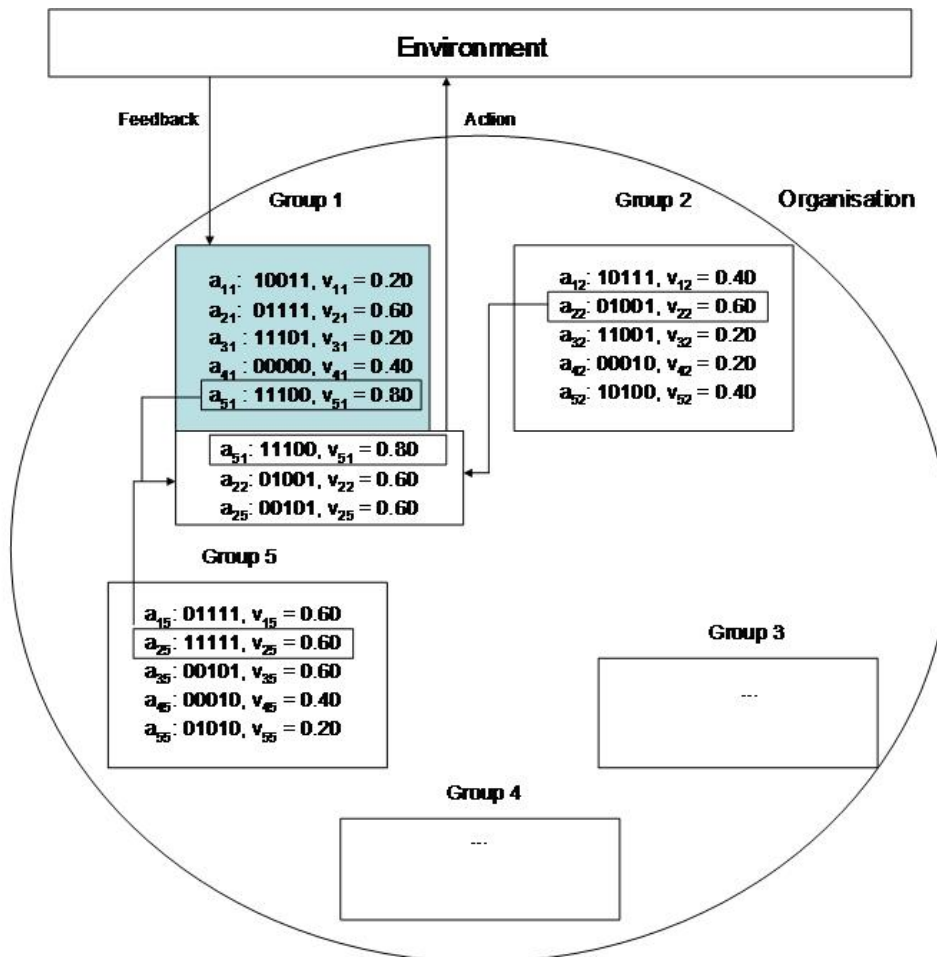
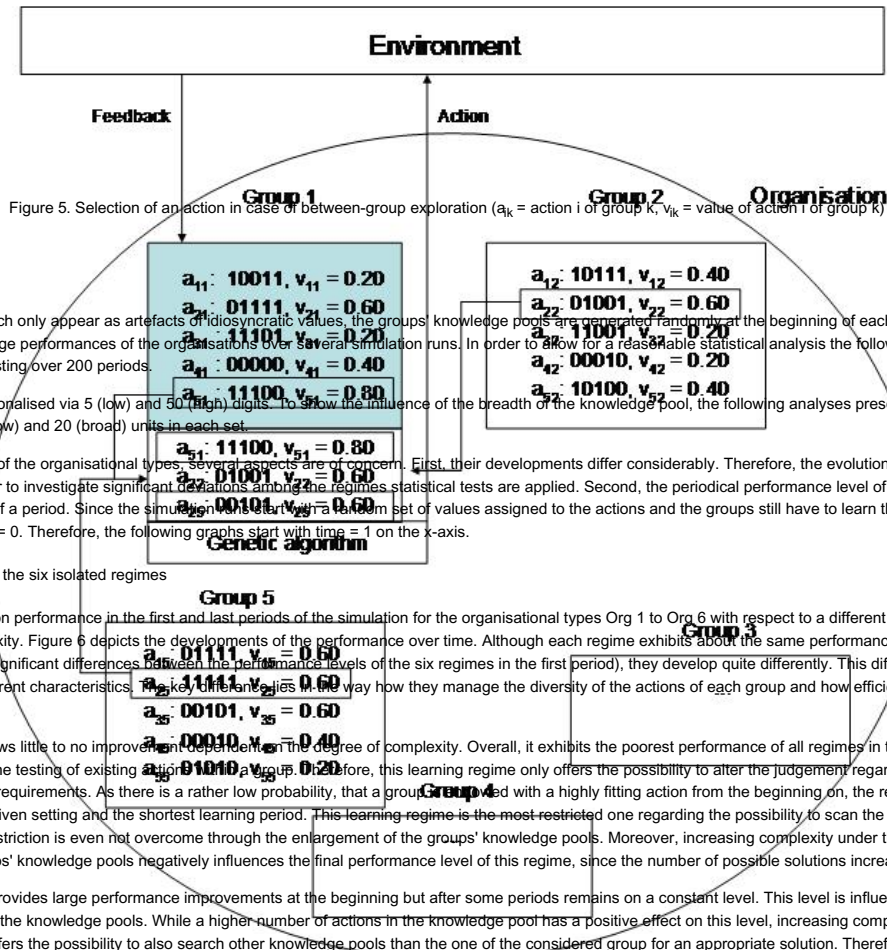


Figure 4. Selection of an action in case of between-group exploitation (a_{ik} = action i of group k , v_{ik} = value of action i of group k)

Between-group exploration

- 4.25 As in case of within-group exploration, between-group exploration is modelled with genetic algorithms to simulate the process of knowledge generation. Moreover, as in the case of between-group exploitation, the groups are structured along a circle and have two neighbours. The procedure of between-group exploration is as follows: Each group forms a team with its immediate neighbours. Each of them cedes its best action to a between-group knowledge set. Hence, this set contains three actions: one action from the neighbour on the left, one action from the neighbour on the right and the action with the highest value from the group's own knowledge pool. This set is used as basis of knowledge creation, analogous to the procedure used in within-group exploration. Consequently, under this learning regime each group receives a newly created action. And as the new actions are created by using actions from different groups, knowledge transfer occurs.
- 4.26 Analogous to the procedure with the other regimes, each group presents its action to the environment, which assigns a value to it. The new action either replaces the weakest action in the group's knowledge pool or, if it already exists, it is simply revalued with the environmental feedback. Figure 5 illustrates the selection process for Group 1 in an arbitrary period.



Results

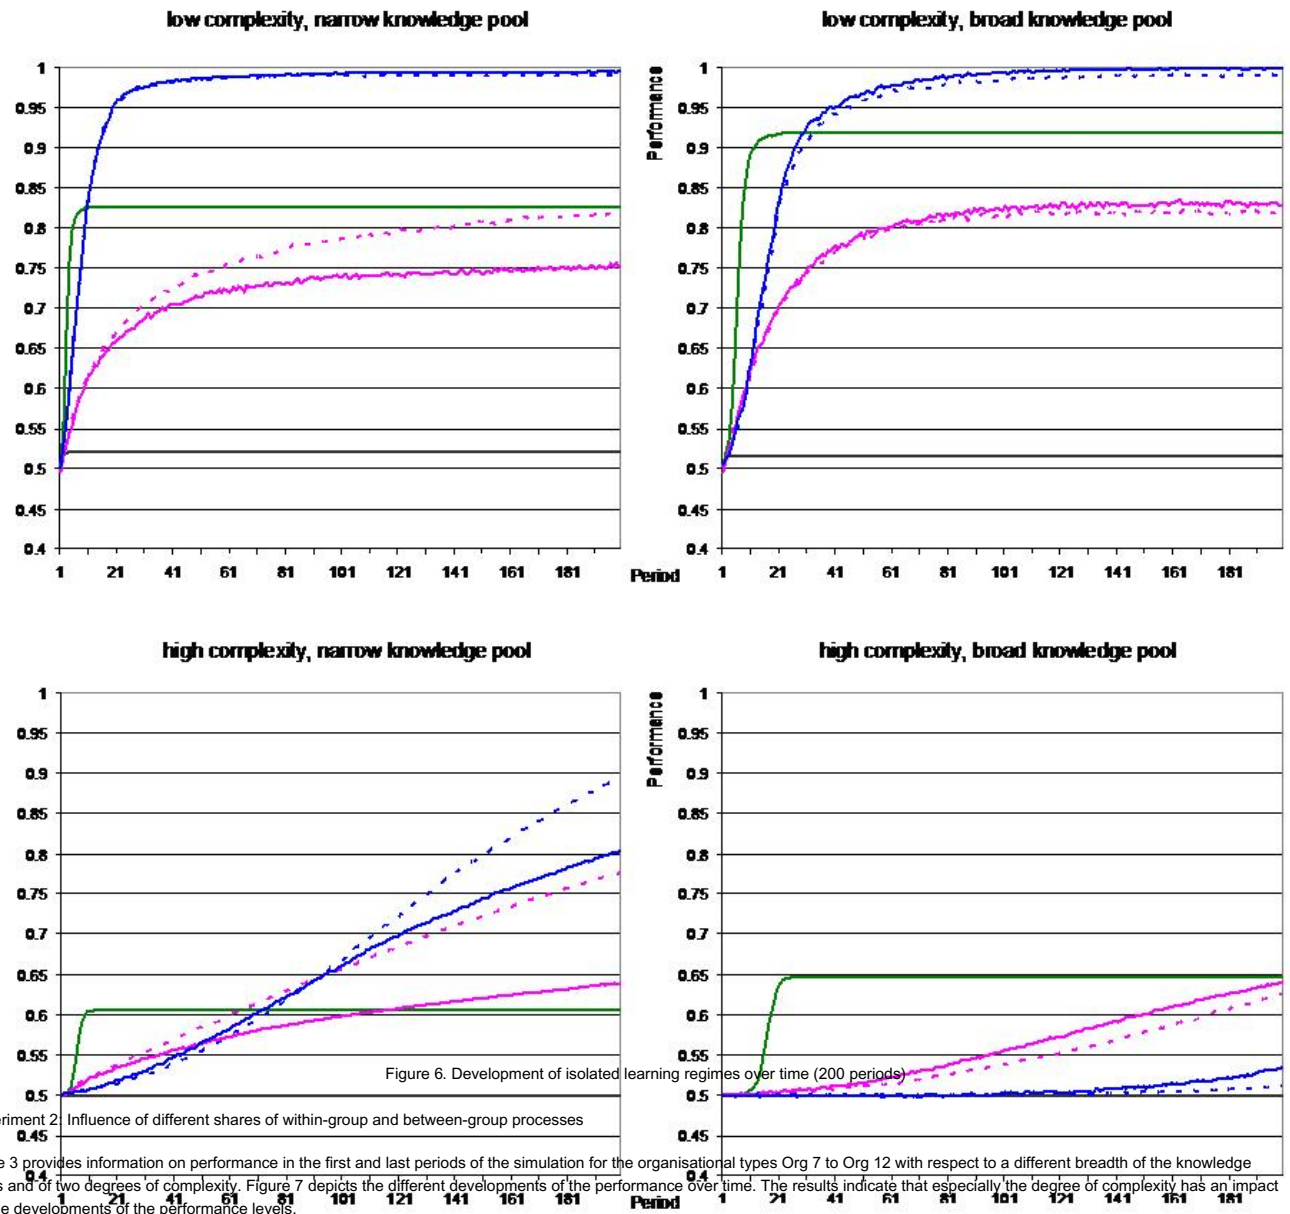
- 5.1 In order to prevent results which only appear as artefacts of idiosyncratic values, the groups' knowledge pools are generated randomly at the beginning of each simulation run. Hence, the following results are the average performances of the organisations over several simulation runs. In order to allow for a reasonable statistical analysis the following results are all based on 1000 simulation runs, each lasting over 200 periods.
- 5.2 The task complexity is operationalised via 5 (low) and 50 (high) digits. To show the influence of the breadth of the knowledge pool, the following analyses present the results of each simulation setting with 5 (narrow) and 20 (broad) units in each set.
- 5.3 In analysing the performance of the organisational types, several aspects are of concern. First, their developments differ considerably. Therefore, the evolution of their means is exhibited graphically. Moreover, in order to investigate significant deviations among the regimes statistical tests are applied. Second, the periodical performance level of each organisational type is always calculated at the end of a period. Since the simulation starts with a random set of values assigned to the actions and the groups still have to learn the correct values, there exist no performance level on time = 0. Therefore, the following graphs start with time = 1 on the x-axis.
- Experiment 1: Performance of the six isolated regimes
- 5.4 Table 2 provides information on performance in the first and last periods of the simulation for the organisational types Org 1 to Org 6 with respect to a different breadth of the knowledge pools and of two degrees of complexity. Figure 6 depicts the developments of the performance over time. Although each regime exhibits about the same performance level in the first period (as Table 2 shows, there are no significant differences between the performance levels of the six regimes in the first period), they develop quite differently. This differing evolution of each regime can be explained by their different characteristics. The key differences in the way how they manage the diversity of the actions of each group and how efficiently each group can scan the space of possible solutions.
- 5.5 *Within-group exploitation* shows little to no improvement dependent on the degree of complexity. Overall, it exhibits the poorest performance of all regimes in the four settings. Within-group exploitation concentrates on the testing of existing actions. Therefore, this learning regime only offers the possibility to alter the judgement regarding the fitting between a given action and the environmental requirements. As there is a rather low probability, that a group started with a highly fitting action from the beginning on, the regime exhibits the lowest performance level under the given setting and the shortest learning period. This learning regime is the most restricted one regarding the possibility to scan the space of possible solutions. In the presented settings, this restriction is even not overcome through the enlargement of the groups' knowledge pools. Moreover, increasing complexity under the assumption of a given number of actions in the groups' knowledge pools negatively influences the final performance level of this regime, since the number of possible solutions increases.
- 5.6 *Between-group exploitation* provides large performance improvements at the beginning but after some periods remains on a constant level. This level is influenced by both the degree of complexity and the breadth of the knowledge pools. While a higher number of actions in the knowledge pool has a positive effect on this level, increasing complexity affects it negatively. Between-group exploitation offers the possibility to also search other knowledge pools than the one of the considered group for an appropriate solution. Therefore, compared to within-group exploitation, the space of possible solutions can be more effectively scanned. However, also this regime suffers from the fact that it only can operate on existing actions. Therefore, after some time of fast learning it quickly reaches its final performance level. The broader the groups' knowledge pools are, the more possible solutions can be scanned. Therefore, increasing the knowledge pools has a positive effect on the reachable performance level. In contrast, increasing complexity under the assumption of a given number of actions in the groups' knowledge pools negatively influences the final performance level of this regime, since the number of possible solutions increases.
- 5.7 In three settings, the two configurations of *within-group exploration* exhibit a lower rate of improvement than between-group exploitation. They also develop differently from between-group exploitation, as in the fourth setting (broad knowledge pools and high complexity), they outperform between-group exploitation, while in the other three settings their performance is lower. Moreover, the developments of the two configurations in terms of innovation rates differ from each other. In case of narrow knowledge pools a higher rate of innovativeness results in a higher performance level, while in case of broader knowledge pools the opposite holds. The observed developments can be explained as follows: Within-group exploration allows for the generation of new actions. Thereby, compared to within-group exploitation it offers the possibility to scan more different actions. Therefore, it outperforms within-group exploitation in all settings. Moreover, in contrast to between-group exploitation, it leads to actions that are completely new to the whole organisation. This is of importance when task complexity is relatively high, since then the probability that a highly fitting action already is part of the initial knowledge pools is rather low. Therefore, in case of high task complexity within-group exploration can reach a higher performance level as between-group exploitation. However, interestingly, the final setting—high task complexity and broad knowledge pools—exhibits a different behaviour. Here within-group exploration is outperformed by between-group exploitation. In this case, the knowledge creation process is confronted on the one hand with a broader basis of actions whose environmental fitting is initially unknown and on the other hand with a high number of different possible solutions to the task. Therefore, the knowledge creation processes only slowly can enhance the performance level. The differences between the two configurations can be explained as follows: A higher innovation rate implemented through the mutation rate leads to the generation of more quite different solutions to the given task. However, through the processes of innovation, those parts of the actions that fit well to the task are destroyed more probably. Hence, in case of a narrow knowledge pool the positive effect of the possibility to test a broad range of different solutions is stronger than the negative effect of destruction. With respect to the broad knowledge pool the opposite holds.
- 5.8 *Between-group exploration* also allows for the generation of new actions and compared to within-group exploration bases its search on a larger knowledge pool, since also transfer between groups is possible. Due to these characteristics this regime offers the possibility to scan the most different solutions. This makes it the superior regime in three settings. However, the combination of high task complexity and a broad knowledge pool changes its development. In this setting, between-group exploration is outperformed by within-group exploration and between-group exploitation. In this case, the same aspect holds that was already discussed with respect to within-group exploration: Because of the combination of many different possible solutions (due to high task complexity) and broad knowledge pools with initially unknown environmental fittings, the learning process exhibits rather low performance increases. In case of between-group exploration this process is further slowed down compared to within-group exploration, because the groups do not search separately but are partly linked to each other through the process of knowledge transfer and hence, the solutions that the groups scan in the same period are more similar. Finally, the two configurations in terms of innovation rates differ similarly to the two configurations of within-group exploration due to the same reasons.
- 5.9 In sum, with respect to the influence of complexity and breadth of knowledge pools the following findings can be summarized:

- | Proposition | Proposition | Proposition |
|--|-------------|-------------|
| 1: Complexity negatively influences the performance of all six learning regimes. The breadth of knowledge pools positively influences the performance of between-group exploitation. The effect of the breadth of knowledge pools on the explorative learning regimes depends also on the degree of complexity. | | |
| 5.10 The findings exhibit differences of performance regarding exploration and exploitation on the two organisational levels indicating that a study of these processes on different levels fosters theory building. From the mentioned observations especially four interesting aspects emerge. | | |
| 5.11 First, in all cases, between-group exploitation performs best in the early periods. Additionally, the higher the degree of complexity is, the longer it takes until other regimes reach its performance level, if they reach it at all. This spread between between-group exploitation and other regimes increases, if the knowledge pool of the groups gets larger. In sum, between-group exploitation provides the superior strategy to come up quickly with a good solution, especially in complex situation. This result is in line with existing literature. However, the simulation experiments provide a more differentiated point of view, as they show that only <i>between-group</i> exploitation supplies good solutions in the short run while <i>within-group</i> exploitation remains on a quite low performance level. | | |
| 5.12 Second, while between-group exploitation exhibits the highest performance levels in the early periods of the simulations, between-group exploration performs best in the long run in three cases. In essence, this type needs more time to arrive at a good solution. However, in the long run, its solutions are better, as through knowledge creation improvements beyond existing knowledge are possible. Nevertheless, if one increases the breadth of knowledge pools in a situation of low complexity, the difference between the two mentioned regimes decreases. | | |
| 5.13 Third, in case of high complexity and broad knowledge pools, within-group exploration performs better than between-group exploration that increases its performance level very slowly. Again, the simulation experiments show that a differentiation of explorative processes along organisational levels provides a better perspective on the performance of exploration. | | |

- 5.14 Fourth, the innovation rate of the explorative process has an impact on the performance of both within-group and between-group exploration. This finding points to the fact that also exploration on a given organisational level might result in different outcomes dependent on its specific definition. So far, literature defines exploration rather broad as a process of innovation. Yet, innovation can occur with different degrees of intensity.

Breadth of knowledge pools	Complexity	Organisation	First Period		Last Period		Test 1	Test 2	Test 3
			Mean	Std.	Mean	Std.			
narrow	low	ET-WG-0	0.5053	0.1013	0.5205	0.0950	***	n	***
		ET-BG-0	0.5043	0.1172	0.8264	0.1163	***	n	***
		ER-WG-M	0.4957	0.1041	0.7523	0.0761	***	n	***
		ER-BG-M	0.5011	0.1303	0.9948	0.0297	***	n	
		ER-WG-H	0.4959	0.1011	0.8201	0.0637	***	n	n
		ER-BG-H	0.5066	0.1293	0.9908	0.0187	***		***
	high	ET-WG-0	0.4999	0.0320	0.4999	0.0320	n	n	***
		ET-BG-0	0.5008	0.0466	0.6053	0.0380	***	n	***
		ER-WG-M	0.5001	0.0320	0.6390	0.0277	***	n	***
		ER-BG-M	0.5006	0.0408	0.8048	0.0817	***	n	***
		ER-WG-H	0.4990	0.0325	0.7766	0.0394	***	n	***
		ER-BG-H	0.5026	0.0425	0.8947	0.1379	***		
broad	low	ET-WG-0	0.4986	0.0968	0.5157	0.0905	***	n	***
		ET-BG-0	0.5016	0.1005	0.9196	0.0970	***	n	***
		ER-WG-M	0.4952	0.1006	0.8315	0.0617	***	n	***
		ER-BG-M	0.5048	0.1289	0.9986	0.0126	***		
		ER-WG-H	0.5012	0.1010	0.8204	0.0626	***	n	***
		ER-BG-H	0.5039	0.1295	0.9898	0.0206	***	n	***
	high	ET-WG-0	0.4985	0.0311	0.4985	0.0311	n	n	***
		ET-BG-0	0.4995	0.0433	0.6467	0.0307	***	n	
		ER-WG-M	0.5007	0.0314	0.6397	0.0412	***	n	***
		ER-BG-M	0.5000	0.0427	0.5341	0.0801	***	n	***
		ER-WG-H	0.5018	0.0319	0.6264	0.0579	***		***
		ER-BG-H	0.4997	0.0415	0.5101	0.0596	***	n	***
Test 1: Significance of difference between performance levels in first and last periods (Wilcoxon-test)									
Test 2: Significance of difference between performance levels of the six organisational types in the first period, the performance levels are compared to the next higher level (Wilcoxon-test)									
Test 3: Significance of difference between performance levels of the six organisational types in the last period, the performance levels are compared to the next higher level (Wilcoxon-test)									

Table 2. Statistic of isolated learning regimes (*** p < 0.001, n = insignificant)



5.15 Table 3 provides information on performance in the first and last periods of the simulation for the organisational types Org 7 to Org 12 with respect to a different breadth of the knowledge pools and of two degrees of complexity. Figure 7 depicts the different developments of the performance over time. The results indicate that especially the degree of complexity has an impact on the developments of the performance levels.

5.16 In case of low task complexity, WG25-M and WG25-H exhibit the highest final performance levels. With an increasing share of within-group processes the organisational types perform worse. Moreover, the innovation rate has a significant, but in absolute values very small influence on the final performance level.

5.17 The picture changes when complexity is increased. The order of organisational types remains the same, only their improvements are slowed down. However, those organisational types that use a lower innovation rate (WG25-M, WG50-M, and WG75-M) exhibit considerably lower performance levels. Furthermore these organisational types have very similar final performance levels. These observations can be condensed to the following propositions:

Proposition

- 4: Organisations with higher innovation rates can better cope with higher degrees of complexity, independently of the share of between-group processes. In situations of low complexity, organisations with a high share of between-group processes perform best independent of their innovation rate. In situations of high complexity, organisations with a high degree of between-group processes in combination with high innovation rates perform best.

Proposition

5.18 Again, these differences stem from the different characteristics of the underlying processes. The between-group processes allow for the search of an optimal solution in a broader knowledge pool as they provide the possibility to transfer knowledge between the groups. Moreover, in case of high complexity, a higher innovation rate results in actions that are more different from each other providing an advantage in a setting of high complexity where many possible solutions for a given task exist.

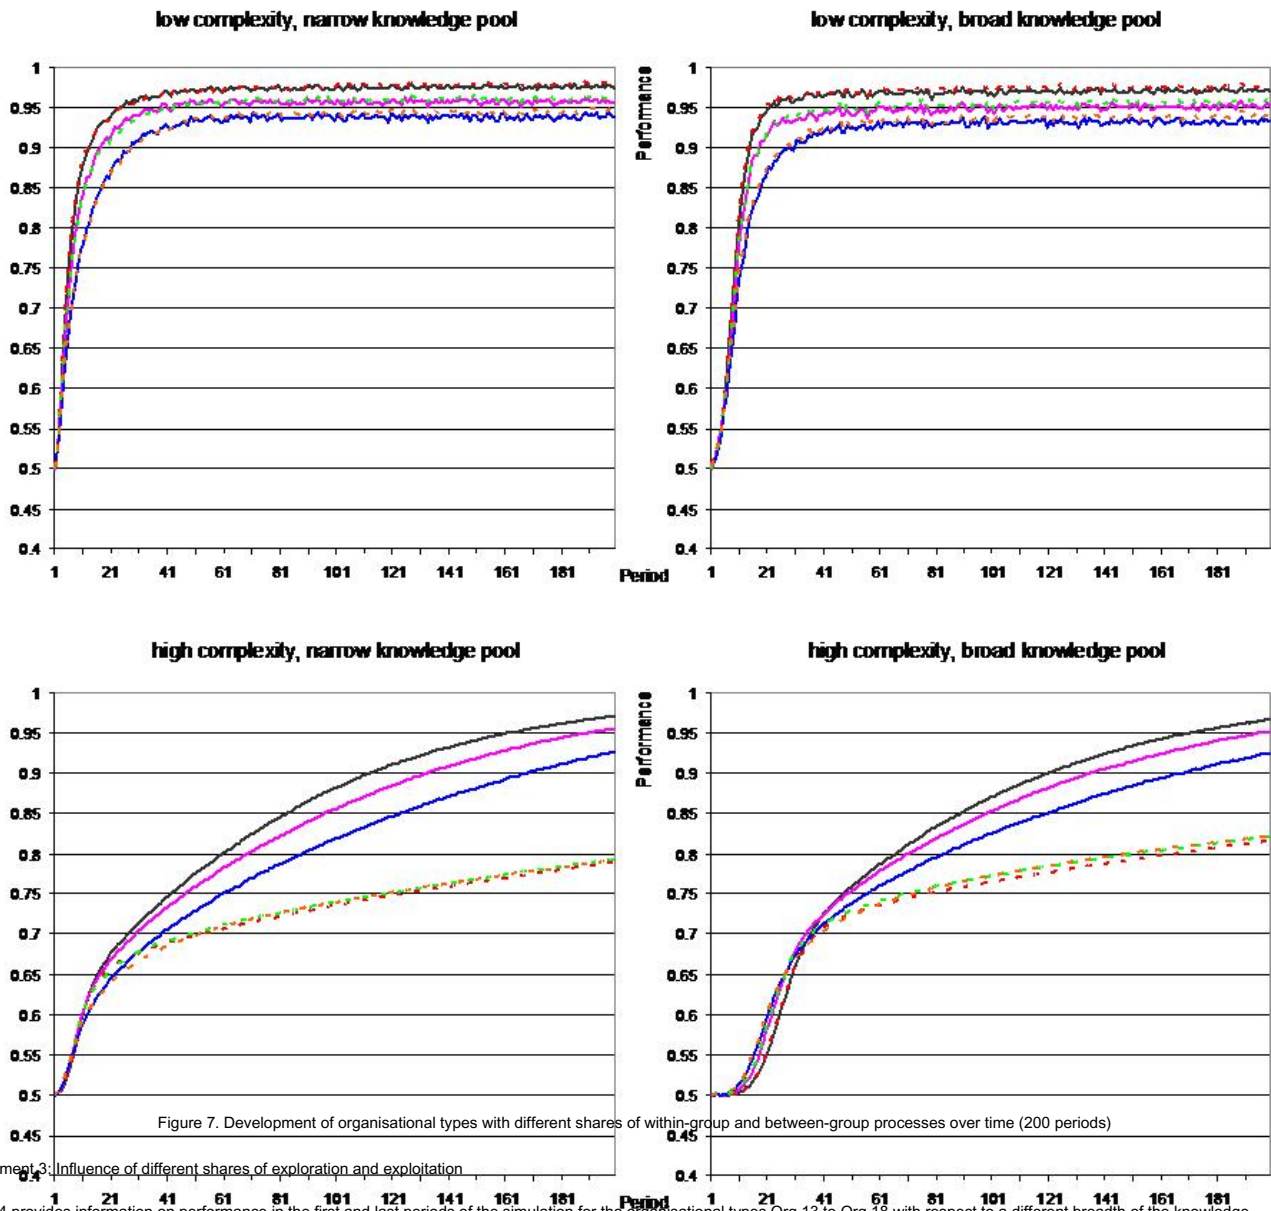
Breadth of knowledge pools	Complexity	Organisation	First Period		Last Period		Test 1	Test 2	Test 3
			Mean	Std.	Mean	Std.			
narrow	low	WG25-M	0.5028	0.1176	0.9800	0.0553	***		
		WG50-M	0.4985	0.1147	0.9626	0.0709	***	n	***
		WG75-M	0.5014	0.1077	0.9437	0.0808	***	n	***
		WG25-H	0.5015	0.1188	0.9750	0.0603	***	n	***
		WG50-H	0.4987	0.1108	0.9559	0.0771	***	n	***
		WG75-H	0.5025	0.1050	0.9377	0.0867	***	n	***
	high	WG25-M	0.5023	0.0406	0.7895	0.0431	***		n
		WG50-M	0.5010	0.0383	0.7931	0.0441	***	n	***
		WG75-M	0.5008	0.0349	0.7917	0.0447	***	n	n
		WG25-H	0.5021	0.0408	0.9711	0.0219	***	n	
		WG50-H	0.5011	0.0378	0.9563	0.0271	***	n	***
		WG75-H	0.5006	0.0346	0.9273	0.0330	***	n	***
broad	low	WG25-M	0.5039	0.1101	0.9751	0.0673	***	n	
		WG50-M	0.5007	0.1117	0.9571	0.0799	***	n	***
		WG75-M	0.5032	0.1035	0.9370	0.0862	***	n	***
		WG25-H	0.4984	0.1105	0.9715	0.0673	***	n	***
		WG50-H	0.5006	0.1102	0.9516	0.0852	***	n	**
		WG75-H	0.5092	0.1043	0.9316	0.0934	***		***
	high	WG25-M	0.5008	0.0394	0.8156	0.0414	***	n	**
		WG50-M	0.5006	0.0369	0.8207	0.0393	***	n	n
		WG75-M	0.5004	0.0334	0.8222	0.0398	***	n	***
		WG25-H	0.5003	0.0402	0.9661	0.0239	***	n	
		WG50-H	0.5005	0.0378	0.9523	0.0273	***	n	***
		WG75-H	0.5012	0.344	0.9256	0.0318	***		***

Test 1: Significance of difference between performance levels in first and last periods (Wilcoxon-test)

Test 2: Significance of difference between performance levels of the six organisational types in the first period, the performance levels are compared to the next higher level (Wilcoxon-test)

Test 3: Significance of difference between performance levels of the six organisational types in the last period, the performance levels are compared to the next higher level (Wilcoxon-test)

Table 3. Statistic of organisational types with different shares of within-group and between-group processes (*** p < 0.001, ** p < 0.01, n = insignificant)



5.19 Table 4 provides information on performance in the first and last periods of the simulation for the organisational types Org 13 to Org 18 with respect to a different breadth of the knowledge pools and of two degrees of complexity. Figure 8 depicts the different developments of the performance over time with respect to the organisational types. Again, the results indicate that especially the degree of complexity has an important influence on the developments of the performance levels.

5.20 In case of low complexity, on the one hand the introduction of a higher rate of innovation has hardly any material effect on the performance level, as the absolute differences between both respective organisational types is very small. On the other hand, increasing the share of exploitation compared to exploration positively influences the final performance level.

5.21 In contrast, with respect to the setting of high task complexity, a high innovation rate is favourable compared to a low one and a higher share of exploration leads to superior performance levels.

5.22 In sum, dependent on the degree of task complexity, different proportions of explorative and exploitative processes are advisable:

Proposition

7: In case of low complexity a high share of exploitation is favourable, while in case of high complexity a high share of exploration is advantageous.

5.23 Again, these differences stem from the different characteristics of the underlying processes. Exploitation—independently whether it is within-group or between-group—operates on existing actions. It offers the possibility to save favourable actions and it prevents the organisation from a considerable performance decrease as it does not allow for any innovative—and hence possibly very poorly fitting—solutions. In contrast, exploration—independently whether it is performed on a within-group or a between-group basis—provides the possibility to generate rather new solutions. Thereby it is more favourable in highly complex situations, in which many different possible solutions to a given task exist. A higher rate of innovation augments this positive effect.

Breadth of knowledge pools	Complexity	Organisation	First Period		Last Period		Test 1	Test 2	Test 3
			Mean	Std.	Mean	Std.			
narrow	low	ER25-M	0.4995	0.1090	0.9772	0.0599	***	n	
		ER50-M	0.4985	0.1147	0.9626	0.0709	***	n	***
		ER75-M	0.5033	0.1173	0.9426	0.0823	***		***
		ER25-H	0.5031	0.1106	0.9755	0.0615	***	n	*
		ER50-H	0.4987	0.1108	0.9559	0.0771	***	n	***
		ER75-H	0.5022	0.1153	0.9353	0.0863	***	n	***
	high	ER25-M	0.5005	0.0385	0.7358	0.0449	***	n	***
		ER50-M	0.5010	0.0383	0.7931	0.0441	***	n	***
		ER75-M	0.5006	0.0370	0.8350	0.0431	***	n	***
		ER25-H	0.5006	0.0388	0.8798	0.0414	***	n	***
		ER50-H	0.5011	0.0378	0.9563	0.0271	***		***
		ER75-H	0.5002	0.0364	0.9778	0.0171	***	n	
broad	low	ER25-M	0.5056	0.1075	0.9770	0.0644	***		
		ER50-M	0.5007	0.1117	0.9572	0.0799	***	n	***
		ER75-M	0.5003	0.1106	0.9366	0.0891	***	n	***
		ER25-H	0.5018	0.1031	0.9768	0.0612	***	n	n
		ER50-H	0.5006	0.1102	0.9516	0.0852	***	n	**
		ER75-H	0.5023	0.1144	0.9278	0.0952	***	n	***
	high	ER25-M	0.5018	0.0375	0.7721	0.0409	***		***
		ER50-M	0.5006	0.0369	0.8207	0.0393	***	n	***
		ER75-M	0.5017	0.0381	0.8521	0.0392	***	n	***
		ER25-H	0.5009	0.0378	0.8871	0.0393	***	n	***
		ER50-H	0.5005	0.0378	0.9523	0.0273	***	n	***
		ER75-H	0.5010	0.0354	0.9702	0.0214	***	n	

Test 1: Significance of difference between performance levels in first and last periods (Wilcoxon-test)

Test 2: Significance of difference between performance levels of the six organisational types in the first period, the performance levels are compared to the next higher level (Wilcoxon-test)

Test 3: Significance of difference between performance levels of the six organisational types in the last period, the performance levels are compared to the next higher level (Wilcoxon-test)

Table 4. Statistic of organisational types with different shares of exploration and exploitation (*** p < 0.001, ** p < 0.01, * p < 0.05, n = insignificant)

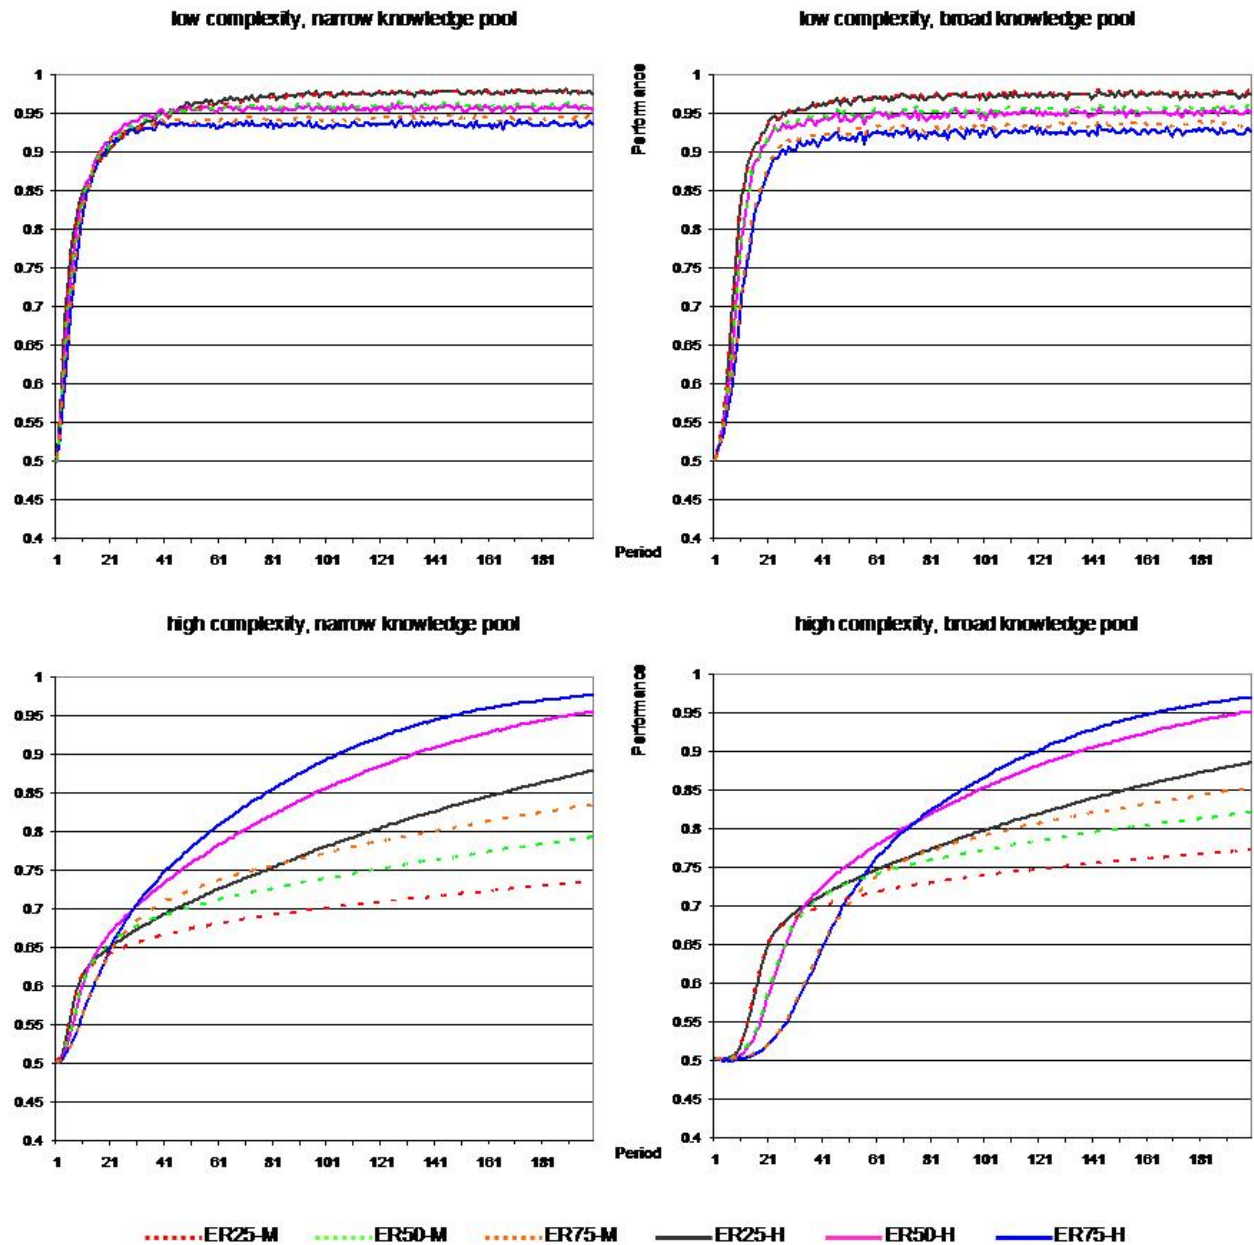


Figure 8. Development of organisational types with different shares of exploration and exploitation over time (200 periods)



Conclusions

Discussion of results and implications

- 6.1 The previous simulation experiments exhibited the influence of different ways of handling diversity of knowledge on the performance of an organisation that is separated into groups. Their results provide several implications for future research regarding an integrated theory of exploration and exploitation and for organisational practice.
- 6.2 First, the findings exhibit considerable differences between the isolated regimes with respect to the evolution and the final attainable performance levels. Therefore, the classification of both exploration and exploitation on the group level in between-group and in within-group processes proves to be a reasonable starting point for further research. It puts existing results in the literature into a new perspective: in general, empirical and theoretical research points to the importance of balancing exploration and exploitation and aligning it with environmental requirements (e.g., Rivkin and Siggelkow 2003; Siggelkow and Levinthal 2003; He and Wong 2004; Jansen et al. 2006). This has led to the claim for an active management of exploration and exploitation (Benner and Tushman 2003; He and Wong 2004). The results of the simulation point to the fact that the outcomes of explorative and exploitative processes on the group level are significantly affected by the possibility of interactions between groups. Hence, a concept of balancing exploration and exploitation, which fosters organisational performance, has to take into account this intergroup learning dimension. The following example shows the importance of this argument: Organisations often deliberately undergo periods of reorganisation to foster also processes of knowledge transfer and generation. Thereby, they often establish new structures especially on the group level. However, as Brown and Duguid (1991, p. 49) have pointed out with respect to communities-of-practice as nuclei of knowledge processes, the "reorganisation of the workplace into canonical groups can wittingly or unwittingly disrupt these highly functional uncanonical—and therefore often invisible—communities." These processes of reorganisation may not only cut invisible groups like communities-of-practice into pieces but they can also disturb relations between existing other informal organisational groups and thereby inhibit the evolution of *between-group* processes ending up with the (mainly poorer performing) within-group processes. As a consequence, empirical research and theory building should incorporate the possibility of *different* types of exploration and exploitation in order to deliver suggestions to practice that are better customised to the situational context. In particular, a balance of exploration and exploitation should be discussed with respect to the between-group perspective.
- 6.3 Second, as Fang et al. (2010) point out, literature on exploration and exploitation especially emphasises the tendency of organisations to concentrate on exploitation, because it delivers better results in the short run. The findings of Experiment 1 reveal a more differentiated picture.
- 6.4 On the one hand, while between-group exploitation actually outperforms the other regimes in the earlier periods, within-group exploitation remains on the lowest performance level. Hence,

the superiority of exploitation depends on the breadth of the knowledge pools on which it can operate. While it is beneficial to a whole organisation to stick to its existing solutions, this strategy might be quite demotivating to the members of a domain smaller than the whole organisation, as it provides rather low performance outcomes. Therefore, the question rises of whether the mentioned tendency of focusing on exploitation is actually a universal phenomenon that is independent from organisational levels or whether other motivational aspects might counteract it especially on lower organisational levels, like the individual and the group levels.

- 6.5 On the other hand, the results of Experiment 1 suggest that the difference in performance between between-group exploitation and exploration gets smaller with larger knowledge pools. Again, this result puts the discussion regarding the long-term superiority of exploration compared to exploitation into a new perspective. As explorative strategies in practice are relatively costly, since they can lead to many failures, in some cases the application of (between-group) exploitation might be the superior strategy, especially when the organisation already possesses as a whole a wide range of different solutions. Hence, the simulation results extend the current discussion of the usage of routines—which is a manifestation of exploitation—as an important source of *dynamic capability* (Feldman 2000; Feldman and Pentland 2003; Espedal 2006): Organisational routines are a depository of knowledge (Cohen and Bacdayan 1994), which in the older literature has often been seen as the cause of organisational inertia by blocking adaptation processes. Yet, exploitation as the application of existing organisational routines might enable fast reactions to environmental changes in a better way than exploration. The simulation results add to the previous discussion regarding the aspect of task complexity: with increasing task complexity the utility of such routines increases in comparison to exploration.
- 6.6 Third, Experiment 1 also exhibits a somewhat counterintuitive result. In case of high task complexity and broad knowledge pools, between-group exploration performs very poor compared to both within-group exploration and between-group exploitation. Here, the combination between knowledge transfer and the creation of new solutions on the basis of this transferred knowledge considerably slows down the process of improvement. Hence, if an organisation, that consists of groups which possess already a broad knowledge pool, e.g. because they contain well trained experts, has to recover after an environmental shock in an complex environment, the transfer of pre-existing knowledge and the isolated generation of new knowledge are more advisable strategies than a between-group explorative process.
- 6.7 These findings with respect to the isolated regimes are extended by the simulation results regarding the effect of different combinations of these regimes.
- 6.8 As the findings of Experiment 2 indicate, a relatively high share of between-group processes compared to within-group processes leads to considerably higher performance levels. These results are consistent with the empirical findings by Jansen et al. (2006): They found that connectedness between members of an organisational unit positively influences both explorative and exploitative innovation. Although their findings apply to another organisational level, they point into the same direction as the present simulation results do.
- 6.9 Moreover, Experiment 2 also exhibits that in case of high tasks complexity, the mentioned effect only holds, when the innovation rates of the explorative processes are relatively high. Hence, for organisations containing groups that are capable of generating very innovative new solutions the fostering of between-group processes is more promising than for organisations where the innovation rate is rather low. These findings further qualify the results in literature on the benefits of specific organisational structures with respect to explorative and exploitative processes. Lazer and Friedman's (2007) study, which does not specifically deal with groups, indicates that for intermediate timeframes a moderately connected system performs best with respect to balance exploration and exploitation. Fang et al. (2010) show, that an organisational structure of semiautonomous subunits with a limited number of links between these units fosters a balance between exploration and exploitation. In essence, the previous analysis is based on a similar structure, as the studied organisations contain several groups that in case of the between-group processes are capable of transferring some actions but that at the same time are autonomous units.
- 6.10 Furthermore, a comparison between the results of Experiments 1 and 2 indicates that in a situation of low task complexity a regime of pure between-group exploration provides the best solutions in the long run. However, in case of higher task complexity learning regimes that allow for a mixture between between-group and within-group exploration with a high degree of innovation and additionally a share of exploitation outperform pure between-group exploration. The results thereby provide a more differentiated perspective on the superiority of a balance between exploration and exploitation especially in complex situations (e.g., Hodgson and Knudsen 2006). In practice, groups are confronted with a whole bunch of different tasks ranging from a rather low degree of complexity to a higher one. The mentioned findings indicate that in such contexts fostering a mixture of all four regimes should lead to higher performances than the concentration on only one of them. Consequently, the necessity of ambidexterity that is discussed in an organisational context also holds for the group-level.
- 6.11 Additionally, the findings of Experiment 3 point to the fact that dependent on the degree of task complexity, different proportions of explorative and exploitative processes are advisable. While in case of rather complex tasks a higher share of exploration provides better performance, in case of tasks that are characterised by a low degree of complexity a higher share of exploitation results in superior performance. Hence, the findings indicate that the optimal balance of exploration and exploitation is not a fixed share of both processes but that it is affected by the situational context in which exploration and exploitation operate. The present results thereby further qualify findings of previous research (see e.g. the theoretical discussion of the influence of problem complexity on the optimality of exploration and exploitation in Fang et al. 2010).
- 6.12 Finally, it should be remarked that the results of the paper are based on a rather abstract modelling of knowledge handling processes, which does not only fit to the between-group versus within-group level but also to other organisational levels, like the group versus individual level. In this sense, the findings point to rather general aspects of knowledge handling between entities of any kind and within these entities.

Limitations and further research

- 6.13 The results point to important aspects for future theory building. Moreover, several interesting research directions emerge.
- 6.14 First, a more differentiated incorporation of the strength of connectedness between groups seems to be worthwhile. The strength of these connections was neglected to keep the analysis clear. The simulation was based on a binary setting. The groups either had the possibility to interact or to act in isolation. However, as the literature shows, the degree of connections also plays an important role in knowledge transfer processes (e.g., Reagans and McEvily 2003; Fang et al. 2010). Consequently, the binary experimentation setting should be substituted by a continuum, which allows for the characterisation of the quality of the linkage between groups and accounts for the fact that different organisational structures can be differently effective.
- 6.15 Second, in this paper the differentiation among six learning regimes has been proven to deliver insights into explorative and exploitative processes. This differentiated view might induce further analysis of other possibilities cutting the overall exploration and exploitation processes into underlying processes.
- 6.16 Third, the experiments abstracted from any costs of social interaction. Further research should elaborate on this aspect and analyse types and influences of these costs on the exploration and exploitation processes.
- 6.17 Finally, the model can be generalized to fit to other problems. It can be used to analyse other organisational levels, e.g. inter- and intra-firm exploration and exploitation. Moreover, some aspects of the model resemble concepts that have been analysed in the context of organisational network-structures and their effect on the balance between exploration and exploitation. Therefore, it offers starting points for a further development into this direction.



Appendix: Pseudo-Code

y = number of actions per group
n = length of environmental situation and actions

```
For (h = 1; h ≤ number of simulation runs; h++)
{
```

A. Initialization

1. Generate the environmental situation with n digits.
2. Generate one organisation with 5 groups, where each group has m randomly generated actions with n digits and add a randomly generated value to each action, where the value $\in [0, 1]$.
3. Copy this organisation 17 times to generate equal initial conditions for all 18 organisational types.

B. Simulation runs

```
For (i = 1; i ≤ number of periods; i++)
{
  For (j = 1; j ≤ 18; j++)
  {
```

4. Assign a learning regime to organisation j, where in case of those organisational types, that exhibit a mixture between different learning regimes, the learning regime is assigned via a randomized procedure that captures the respective probabilities of exploration and exploitation and of within-group and between group processes.

```
For (k = 1; k ≤ 5; k++)
```

- {
- 5. Group k performs the assigned learning regime, i. e. *Within-group exploitation*:
 - Group k selects an action with the highest value from the group's knowledge pool.
 - If more than one action has the highest value, one of these actions is selected randomly.
 - Group k presents this action to the environment.
- Between-group exploitation*:
 - The two direct neighbours of group k each select one action from their respective knowledge pools with the highest value and provide it to group k.
 - Group k compares the values of these two actions with the highest value of the actions in its knowledge pool and chooses the action with the highest value.
 - If two or more actions have all the highest value, one action is selected randomly.
 - Group k presents this action to the environment.
- Within-group exploration*:
 - Group k selects the two actions with the highest values from the group's knowledge pool.
 - Group k applies a genetic algorithm to these two actions containing
 - a one-point crossing-over and
 - a randomized mutation procedure, dependent on the innovation rate, the probability that a digit is switched to the opposite is either 0.01 (high innovation rate) or 0.001 (medium innovation rate).
 - The values of the two new actions are calculated on the basis of the values of the parent actions according to formula (2).
 - Group k selects the action with the highest value from the two newly generated actions.
 - If both actions have the same value, one action is chosen randomly.
 - Group k presents this action to the environment.
- Between-group exploration*:
 - The two direct neighbours of group k each select the action from their respective knowledge pools with the highest value and provide it to group k.
 - Group k compares the values of these two actions with the highest value of the actions in its knowledge pool and chooses the two actions with the highest values.
 - Group k applies a genetic algorithm to these two actions containing
 - a one-point crossing-over and
 - a randomized mutation procedure, dependent on the innovation rate, the probability that a digit is switched to the opposite is either 0.01 (high innovation rate) or 0.001 (medium innovation rate).
 - The values of the two new actions are calculated on the basis of the values of the parent actions according to formula (2).
 - Group k selects the action with the highest value from the two newly generated actions.
 - If both actions have the same value, one action is chosen randomly.
 - Group k presents this action to the environment.
- 6. Environment evaluates the presented action according to formula (1).
- 7. Group k either revalues the used action, if it is already part of the knowledge pool, or substitutes the action with the lowest value in its set by the used action and assigns the received value to it.

}
 }
 }
 }



References

- AUH, S and Menguc, B (2005) Balancing exploration and exploitation: The moderating role of competitive intensity. *Journal of Business Research*, 58. pp. 1652- 1661. [doi:10.1016/j.jbusres.2004.11.007]
- AXELROD, R (1997) Advancing the Art of Simulation in the Social Sciences. In Conte R, Hegselmann R and Terno P (Eds.). *Simulating Social Phenomena*. Berlin et al.: Springer. pp. 21-40 [doi:10.1007/978-3-662-03366-1_2]
- BABUJI, H and Crossan, M M (2004) From Questions to Answers: Reviewing Organizational Learning Research. *Management Learning*, 35 (4). pp. 397-417. [doi:10.1177/1350507604048270]
- BELDERBOS, R, FAEMS, D, LETEN, B and VAN LOOY, B (2010) Technological Activities and Their Impact on the Financial Performance of the Firm: Exploitation and Exploration within and between Firms. *Journal of Product Innovation Management*, 27. pp. 869-882. [doi:10.1111/j.1540-5885.2010.00757.x]
- BENNER, M J and Tushman, M L (2003) Exploitation, Exploration, and Process Management: The Productivity Dilemma Revisited. *Academy of Management Review*, 28 (2). pp. 238-256.
- BONABEAU, E (2002) Agent-Based modeling: Methods and Techniques for Simulating Human Systems. *Proceedings of the National Academy of Sciences*, 99 (10). pp. 7280-7287. [doi:10.1073/pnas.082080899]
- BOOKER, L (1987) Improving Search in Genetic Algorithms. In Davis, L (Ed.), *Genetic Algorithms and Simulated Annealing* (61-73). Los Altos: Morgan Kaufmann Publishers.
- BOWER, G H and Hilgard, E R (1998) *Theories of Learning*, 5th edition. New Jersey: Prentice-Hall.
- BROWN, J S and Duguid, P (1991) Organizational Learning and Communities-of-Practice: Toward a Unified View of Working, Learning, and Innovation. *Organization Science*, 2 (1). pp. 40-57. [doi:10.1287/orsc.2.1.40]
- CARLEY, K (1992) Organizational Learning and Personnel Turnover. *Organization Science*, 3 (1). pp. 20-46. [doi:10.1287/orsc.3.1.20]
- CHATTOE, E (1998) Just How (Un)realistic are Evolutionary Algorithms as Representations of Social Processes? *Journal of Artificial Societies and Social Simulation* 1 (3) <http://www.soc.surrey.ac.uk/JASSS/1/3/2.html>.
- COHEN, M D and Bacdayan, P (1994) Organizational Routines Are Stored as Procedural Memory: Evidence from a Laboratory Study. *Organization Science*, 5 (4). pp. 554-568. [doi:10.1287/orsc.5.4.554]
- COHEN, W M and Levinthal, D A (1990) Absorptive Capacity: A New Perspective on Learning and Innovation. *Administrative Science Quarterly*, 35. pp. 128-152. [doi:10.2307/2393553]
- CYERT, R M and March, J G (1963) *A Behavioral Theory of the Firm*. Englewood Cliffs: Prentice Hall.
- DAWID, H and Kopel, M (1998) On Economic Applications of the Genetic Algorithm: A Model of the Cobweb Type. *Journal of Evolutionary Economics*, 8. pp. 297-315. [doi:10.1007/s001910050066]
- EPSTEIN, J M (1999) Agent-Based Computational Models and Generative Social Science. *Complexity*, 4 (5). pp. 41-60. [doi:10.1002/(SICI)1099-0526(199905/06)4:5<41::AID-CPLX9>3.0.CO;2-F]
- ESPEDAL, B (2006) Do Organizational Routines Change as Experience Changes? *Journal of Applied Behavioral Science*, 42 (4). pp. 468-490. [doi:10.1177/0021886306291601]
- ESPEDAL, B (2008) In the Pursuit of Understanding How to Balance Lower and Higher Order Learning in Organizations. *Journal of Applied Behavioral Science*, 44. pp. 365-390. [doi:10.1177/0021886308319717]
- FANG, C, LEE, J, SCHILLING, M A (2010) Balancing Exploration and Exploitation Through Structural Design: The Isolation of Subgroups and Organizational Learning. *Organization Science*, 21 (3). pp. 625-642. [doi:10.1287/orsc.1090.0468]
- FELDMAN, M S (2000) Organizational Routines as a Source of Continuous Change. *Organization Science*, 11 (6). pp. 611-629. [doi:10.1287/orsc.11.6.611.12529]
- FELDMAN, M S and Pentland, B T (2003) Reconceptualizing Organizational Routines as a Source of Flexibility and Change. *Administrative Science Quarterly*, 48. pp. 94-118.

[doi:10.2307/3556620]

GILBERT, A H, Bell F and Valenzuela C L (1995) Adaptive Learning of Process Control and Profit Optimization Using a Classifier System. *Evolutionary Computation*, 3 (2). pp. 177-198. [doi:10.1162/evco.1995.3.2.177]

GILBERT, G N and Terna, P (2000) How to Build and Use Agent-Based Models in Social Science. *Mind & Society*, 1 (1). pp 57-72. [doi:10.1007/BF02512229]

GILBERT, G N and Troitzsch, K G (1999) *Simulation for the Social Scientist*. Buckingham, Philadelphia: Open University Press.

GREVE, H R (2007) Exploration and Exploitation in Product Innovation. *Industrial and Corporate Change*, 16 (5). pp. 945-975. [doi:10.1093/icc/dtm013]

GUPTA, A K, Smith, K G and Shalley, C E (2006) The Interplay between Exploration and Exploitation. *Academy of Management Journal*, 49 (2). pp. 693-706. [doi:10.5465/AMJ.2006.22083026]

HE, Z L and Wong, P K (2004) Exploration vs. Exploitation: An Empirical Test of Ambidexterity Hypothesis. *Organization Science*, 15 (4). pp. 481-494. [doi:10.1287/orsc.1040.0078]

HERRIOTT, S R, Levinthal, D and March, J G (1985) Learning from Experience in Organizations. *American Economic Review*, 75 (2). pp. 298-302.

HODGSON, G M, and Knudsen, T (2006) Balancing Inertia, Innovation, and Imitation in Complex Environments. *Journal of Economic Issues*, XL (2). pp. 287-295.

HOLLAND, J H (1973) Genetic Algorithms and the Optimal Allocation of Trials. *SIAM Journal of Computing*, 2 (2). pp. 88-105. [doi:10.1137/0202009]

HOLLAND, J H (1995) *Hidden Order: How Adaptation Builds Complexity*. Reading: Perseus Books.

HOLMQVIST, M (2004) Experiential Learning Processes of Exploitation and Exploration Within and Between Organizations: An Empirical Study of Product Development. *Organization Science*, 15 (1). pp. 70-81. [doi:10.1287/orsc.1030.0056]

JANSEN, J P, Van den Bosch, F A J and Volberda, H W (2006) Exploratory Innovation, Exploitative Innovation, and Performance: Effects of Organizational Antecedents and Environmental Moderators. *Management Science*, 52 (11). pp. 1661-1674. [doi:10.1287/mnsc.1060.0576]

KENNEDY, J and Eberhart, R C (2001) *Swarm Intelligence*. San Francisco: Morgan Kaufmann Publishers.

LANT, T K and Mezias, S J (1992) An Organizational Learning Model of Convergence and Reorientation. *Organization Science*, 3 (1). pp. 47-71. [doi:10.1287/orsc.3.1.47]

LAVIE, D and Rosenkopf, L (2006) Balancing Exploration and Exploitation in Alliance Formation. *Academy of Management Journal*, 49 (4). pp. 797-818. [doi:10.5465/AMJ.2006.22083085]

LAZER, D and Friedman, A (2007) The Network Structure of Exploration and Exploitation. *Administrative Science Quarterly*, 52. pp. 667-694. [doi:10.2189/asqu.52.4.667]

LEE, D J and Ahn, J H (2007) Reward Systems for Intraorganizational Knowledge Sharing. *European Journal of Operational Research*, 180. pp. 938-956. [doi:10.1016/j.ejor.2006.03.052]

LEVINTHAL, D A and March J G (1993) The Myopia of Learning. *Strategic Management Journal*, 14. pp. 95-112. [doi:10.1002/smj.4250141009]

LEVITT, B and March, J G (1988) Organizational Learning. *Annual Review of Sociology*, 14. pp. 319-340. [doi:10.1146/annurev.so.14.080188.001535]

LI, Y, Vanhaverbeke, W and Schoenmakers, W (2008) Exploration and Exploitation in Innovation: Reframing the Interpretation. *Creativity and Innovation Management*, 17 (2). pp. 107-126. [doi:10.1111/j.1467-8691.2008.00477.x]

MARCH, J G (1991) Exploration and Exploitation in Organizational Learning. *Organization Science*, 2 (1). pp. 71-87. [doi:10.1287/orsc.2.1.71]

MARCH, J G and Olsen, J P (1975) The Uncertainty of the Past: Organizational Learning under Ambiguity. *European Journal of Political Research*, 3. pp. 147-171. [doi:10.1111/j.1475-6765.1975.tb00521.x]

MILLER, K D, Zhao, M and Calantone, R J (2006) Adding Interpersonal Learning and Tacit Knowledge to March's Exploration-Exploitation Model. *Academy of Management Journal*, 49 (4). pp. 709-722. [doi:10.5465/AMJ.2006.22083027]

MOM, T J M, Van Den Bosch, F A J and Volberda, H W (2007) Investigating Managers' Exploration and Exploitation Activities: The Influence of Top-Down, Bottom-Up, and Horizontal Knowledge Inflows. *Journal of Management Studies*, 44 (6). pp. 910-931. [doi:10.1111/j.1467-6486.2007.00697.x]

REAGANS, R and McEvily, B (2003) Network Structure and Knowledge Transfer: The Effects of Cohesion and Range. *Administrative Science Quarterly*, 48. pp. 240-267. [doi:10.2307/3556658]

REN, Y, Carley, K M and Argote, L (2006) The Contingent Effects of Transactive Memory: When Is It More Beneficial to Know What Others Know? *Management Science*, 52 (5). pp. 671-682. [doi:10.1287/mnsc.1050.0496]

RIVKIN, J W and Siggelkow, N (2003) Balancing Search and Stability: Interdependencies among Elements of Organizational Design. *Management Science*, 49 (3). pp. 290-311. [doi:10.1287/mnsc.49.3.290.12740]

RODAN, S (2005) Exploration and Exploitation Revisited: Extending March's Model of Mutual Learning. *Scandinavian Journal of Management*, 21 (4). pp. 407-428. [doi:10.1016/j.scaman.2005.09.008]

ROTHAERMEL, F T (2001) Incumbent's Advantage through Exploiting Complementary Assets via Interfirm Cooperation. *Strategic Management Journal*, 22 (6-7). pp. 687-699. [doi:10.1002/smj.180]

SIDHU, J S, Volberda, H W and Commandeur, H R (2004) Exploring Exploration Orientation and its Determinants: Some Empirical Evidence. *Journal of Management Studies*, 41. pp. 913-32. [doi:10.1111/j.1467-6486.2004.00460.x]

SIDHU, J S, Commandeur, H R and Volberda, H W (2007) The Multifaceted Nature of Exploration and Exploitation: Values of Supply, Demand, and Spatial Search for Innovation. *Organization Science*, 18 (1). pp. 20-38. [doi:10.1287/orsc.1060.0212]

SIGGELKOW, N and Levinthal, D A (2003) Temporarily Divide to Concur: Centralized, Decentralized, and Reintegrated Organizational Approaches to Exploration and Adaptation. *Organization Science*, 14 (6). pp. 650-669. [doi:10.1287/orsc.14.6.650.24840]

SIGGELKOW, N and Rivkin, J W (2006) When Exploration Backfires: Unintended Consequences of Multilevel Organizational Search. *Academy of Management Journal*, 49 (4). pp. 779-795. [doi:10.5465/AMJ.2006.22083053]

SIMON, H A (1977) *The New Science of Management Decision*, Revised Version. New York: Prentice-Hall.

SIMON, H A (1991) Bounded Rationality and Organizational Learning. *Organization Science*, 2. pp. 125-134. [doi:10.1287/orsc.2.1.125]

TAYLOR, A and Greve, H R (2006) Superman or the Fantastic Four? Knowledge Combination and Experience in Innovative Teams. *Academy of Management Journal*, 49 (4). pp. 723-740. [doi:10.5465/AMJ.2006.22083029]

THOMPSON, J D (1967) *Organizations in Action*. New York: McGraw-Hill Companies.

TUSHMAN, M L and O'Reilly, C (1996) Ambidextrous Organizations: Managing Evolutionary and Revolutionary Change. *California Management Review*, 38 (4). pp. 8-30. [doi:10.2307/41165852]

UOTILA, J, Maula, M, Keil, T and Zahra, S A (2009) Exploration, Exploitation, and Financial Performance: Analysis of S&P 500 Corporations. *Strategic Management Journal*, 30. pp. 221-231. [doi:10.1002/smj.738]