Self-organising processes of task allocation

Zoethout, K.

Chapter 5
Task Dynamics in Self-Organising Task Groups[9]

abstract
Multi-agent simulation is applied to explore how different types of task variety cause workgroups to change their task allocation accordingly. We studied two groups, generalists and specialists. We hypothesised that the performance of the specialists would decrease when task variety increases. The generalists, on the other hand, would perform better in a high task variety condition. The results show that these hypotheses were only partly supported because both learning and motivational effects changed the task allocation process in a much more complex way. We conclude that although no variety leads to specialisation and high variety leads to generalisation, in general, performance is better when variety is low. Further, in case of no variety, specialists clearly outperform generalists. In highly dynamic situations, since there is no space for any development, the behaviour of specialists and generalists becomes more similar, and, consequently also their performance.

5.1 Introduction

Should one hire specialists of generalists to maximise group performance? This question still puzzles personnel managers and organisational scientists alike. Well known by practitioners and scientists is that group performance depends on many factors. Both task factors, such as the number of skills required, rotation schedules and variability in tasks, as well as personal factors such as expertise, learning, motivation and boredom have been found to affect group performance (for an overview see, for example, Yeatts & Hyten, 1998). Whereas the effects of separate variables – or limited combinations - have been empirically investigated by many researchers (e.g., Steiner, 1972; Hackman & Oldham, 1975; Wilke & Meertens, 1994; Tschan and von Cranach, 1996), it is difficult to derive empirical based conclusions on how the combination of these variables affects the performance of a team of experts versus generalists. Social simulation offers a methodology to systematically explore a large number of conditions, and thus may contribute to deriving such conclusions (e.g. Gilbert and

---

[9] This chapter is submitted to Autonomous Agents and Multi Agent Systems as: Zoethout, K., Jager., W, & Molleman, E., Tasks Dynamics and Flexibility of Self-Organising Task Groups: Expertise, Motivational, and Performance Differences of Specialists and Generalists

Troitzsch, 1999). In this chapter we explore how task and personnel factors jointly affect the performance of teams consisting of specialists or generalists.

The social simulation approach we follow here typically tries to explain group processes from a bottom-up perspective. This approach shows the influence of complexity theory and multi-agent simulation to describe higher order developmental and adaptation processes in terms of local interactions (Panait & Luke, 2005; Arrow, McGrath, and Berdahl, 2000; see also Axelrod and Cohen, 2000; Gilbert and Troitzsch, 1999). Although different disciplines meet in the broad area of group dynamics, there are still large differences between them. For instance, if we look at studies on self-organisation, some empirical work shows a solid theoretical basis (e.g. Arrow and Crosson, 2003). However, because of its methodology, these studies are often limited regarding their design. On the other hand, simulation studies may show an elaborate experimentally based description, but are limited with respect to the use of psychological theory regarding properties of individuals (e.g. Kitts, Macy, and Flache, 1999). In general most simulation models concerning processes within task groups or teams focus on abstract computational and mathematical descriptions (e.g. Pynadath and Tambe, 2002a; 2002b) or use a so-called bounded rationality approach (e.g. Carley, 1992). Psychological theory that for instance focuses on the influence of motivational related effects such as boredom, fatigue, etc. however, is less emphasised. Nevertheless, it is well known that motivation strongly influences processes within work groups (Hackman & Oldham, 1975; see also Wilke & Meertens, 1994).

In this chapter we will use a model that combines a multi-agent simulation approach with psychological theory on motivation. Our study deals with the question how a workgroup adapts to changes of the tasks it has to perform. There are a lot of approaches related to task descriptions and performance (Hunt, 1976; Weick, 1979; Tschan & Von Cranach, 1996). Our description is based on Wood (1986) who states that task changes are in fact a component of task complexity. The adaptation process of the group is described in terms of task allocation. Processes of task allocation are affected by task characteristics and group characteristics. With respect to the latter we focus on the expertise and the motivation of the individual team members. Following Ashby (1956) we studied the differences between groups of generalists and specialists, not only regarding the way in which they adapt to task changes, but also the implications for performance.

In the first section of the chapter we focus on the theories and models we use and their formalisation, which form the basis of WORKMATE, the simulation program that we developed to study self-organising processes of task allocation. WORKMATE is used to test a number of hypotheses concerning the relation between task dynamics and performance differences of groups of specialists and generalists. The second section describes the experimental design and the parameter settings. Next we will describe the results and we end up with conclusions and a discussion.

5.2 The model

WORKMATE III is a deterministic discrete event based simulation program for simulating self-organising processes of task allocation that is developed in DELPHI6. It is an elaborated version of the simulation program that we used for experiments on the emergence of job rotation (Zoethout, Jager, and Molleman, 2006a), and the relation between task variety and coordination time (Zoethout Jager, and Molleman, 2006b). In this section we shortly describe the theoretical framework WORKMATE III is based on.

5.2.1 Tasks and task dynamics

A task is considered as a set of actions in such a way that each action is related to a single skill (Hunt, 1976; Weick, 1979; Tschan & von Cranach, 1996). Each action has to be performed a number of times, i.e. cycles, before the whole task is finished. In this way, a task can be represented as a matrix of actions (what) and cycles (how often) (see also Figures 19a and 19b). This chapter deals with the relation between flexibility and performance under conditions of task dynamics. Task dynamics refer to the speed in which tasks change over time (Wood, 1986; see also Tschan & von Cranach, 1996). In our study task dynamics consists of two components, task variety and number of tasks and cycles. Task variety refers to the differences between every next task in relation to the former one with respect to its actions. Number of tasks and cycles refers to the size, i.e. number of cycles, of a single task and the number of times this tasks has to be performed. For instance, given a task variety of 1, performing 2 tasks of 100 cycles implies lower task dynamics than performing 8 tasks of 25 cycles. In the first example after 200 cycles one new skill has been used whereas in the second example after 200 cycles the agents has used seven new skills.

In the Figures 19a and 19b we see a task consisting of 3 actions, a, b, c and 3 cycles, i.e., 1, 2 and 3. Thus, the 3 actions need to be performed 3 times before the complete task is finished. The agents may perform the task in a number of ways, for instance cycle by cycle, action by action, or something in between. The possible ways a task can be allocated are bounded by two general allocation types, generalisation and specialisation. We define generalisation as the multi-functionality of agents, i.e. the agents use all their skills. Specialisation is defined as a clear preference of the agents for a subset of the skills necessary to perform a task (see Figure 19a and 19b).

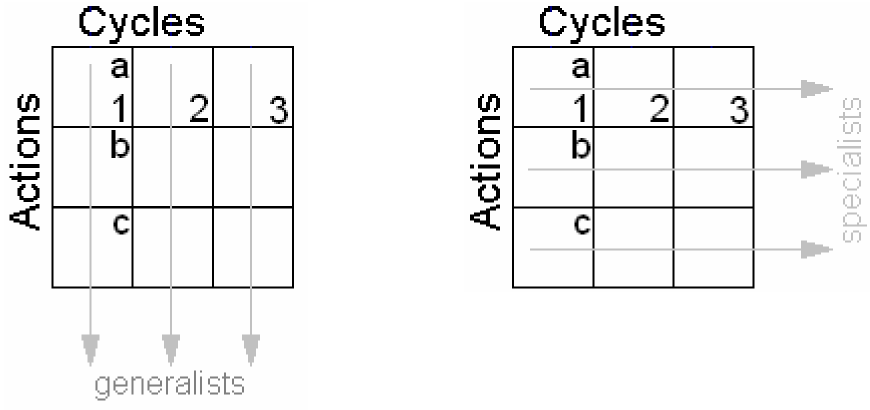

Figure 19a (left) and 19b (right): Representation of generalisation and specialisation.

Each arrow represents the task allocation process of a single agent. Figure 19a depicts that every agent starts with the first action and ends with the last. Thus, all the agents perform all the actions. Figure 19b depicts that each agent performs the same actions in each cycle. Because of the different ways in which the agents can allocate the actions a task consists of, we use the concept of round to describe the specific order in which the task is performed. A round refers to the timestep in which the agent performs a specific part of the task. For instance, the order in which the group of specialists (Figure 19b) performs the task is: at round 1, agent 1 performs action a1, agent 2 performs action b1, agent 3 performs action c1. At round 2 agent 1 performs action a2, etc.

5.2.2 The multi agent system

An agent is a simple model of a human being with properties that are necessary to perform tasks. The individual properties are represented as a set of skills. Each skill has two variable components: expertise and motivation, that are important components that determine group performance (Wilke and Meertens, 1994; see also Steiner, 1972). Skills are passive when they are not used and become active when they are needed for the performance of a task. In every round, each agent performs only one action. When activated, a threshold function determines whether the agents actually wants to perform a particular action. This function holds that only if both the expertise and the motivation are higher than their thresholds, the agent wants to perform the particular action. In this way every agent chooses a subset of actions he would like to perform. If the choices of all of the agents imply that there are more agents sharing the same preference as there are task-actions to perform, the agents negotiate. The negotiation process implies that the agents are trying to change the preferences of the other agents in such a way that the other agents will reach a complementary state with respect to their own (see also Zoethout, Jager, and Molleman, 2004). The influence of the agents is based on their expertise and motivation with respect to a particular skill, which implies that the agent with the highest expertise and/or motivation is more likely to get what he wants. The process ends as soon as the number of agents with a preference of a particular action is equal to the number of available actions. For instance, if we take a look at Figure 19b again, and we imagine that two out of three agents want to perform action a, in the first round, this will not be a problem. However, in the second round there is only one cycle of action a left, which means that they have to negotiate.

When the process of task allocation is being completed, the agents start performing the task. As a result of this, the expertise may change, i.e. the agents will increase the expertise of the skills they use and forget the skills they do not use. Furthermore, the motivation may change, i.e. the agents become bored after performing a particular action for a longer time and recover from it as soon as they stop. An important characteristic of most learning curves is that they reach a maximum asymptotically (Nembhart 2000). Therefore, we define learning by means of the relations among expertise (e) at a certain time (t), expertise in the future (t+1), the maximum expertise (e_{max}), and a parameter (λ, [0,1]) that determines the learning speed:

$$e_{(t+1)} = e_t + \lambda \frac{e_{\max} - e_t}{e_{\max}} \qquad (14a)$$

Forgetting is the inverse of learning, therefore:

$$e_{(t+1)} = \frac{(e_t - \mu) e_{\max}}{e_{\max} - \mu} \qquad (14b)$$

where μ [0,1] determines the forget speed.

In real life an enormous range exists between learning and forget speed of different tasks. Motor tasks such as bicycling or swimming are, once learned, never forgotten, whereas others, such as playing chess or billiards, need to be maintained. Therefore, in the experiments, the balance between learning and forget speed is chosen on rather practical grounds instead of being based on empirical evidence. This holds that the agents are able to forget with a speed that is high enough to produce interesting dynamics, whereas a skill that has not been used for a time is not immediately forgotten.

Motivation curves can be described by applying the same characteristics: a maximum that is reached asymptotically, and recovery as the inverse of boredom. This means that formula (14b) describes the motivational decrease related to boredom and formula (14a) represents the motivational increase related to the recovery from boredom. In this case the parameters μ and v respectively describe the recovery and the boredom speed.

The actual performance of the system is a function of expertise and motivation (Steiner, 1972; Wilke & Meertens, 1994). Both expertise and motivation are defined in terms of the time it takes to perform a task: the higher the degree of expertise or motivation, the sooner the task will be finished. Furthermore, we define a minimal time to complete an action, taction, which is equal to the actual time it takes to perform the action at a maximal rate of expertise and motivation. The actual performance time of a single agent, $t_{perf.}$, can therefore be expressed as:

$$t_{perf.} = \sum_{i=1}^{n} \frac{t_{action\_i}}{\lambda \frac{e_i}{e_{\max}} + (1-\lambda)\frac{m_i}{m_{\max}}} \qquad (15)$$

λ represents a parameter that determines the balance between expertise and motivation. In our experiments we assumed that expertise and motivation have the same effect on the performance time. This means that in the experiments λ is set on 0.5. In the present study, the agents perform the actions simultaneously. This means that the time it takes to perform the total task, tperf., is determined by the slowest agent.

As a consequence of the expertise and motivational changes, the initial preferences of the agents are likely to change, which implies that they may wish to re-allocate their task. We call this task rotation. The way in which the agents allocate and re-allocate their task and the frequency in which the task allocation may be adjusted depends on the group the agents are a member of. Our study makes use of two kind of groups: specialists and generalists.

5.2.3 Specialisation and generalisation

In both groups, the agents together have the abilities to perform the whole task. In both groups, the agents may only start with a new task when the former task has been finished. A group of specialists consists of agents that are all specialised in a particular part of the task. Although they do have the skills to perform the other actions as well, they have a clear preference to perform certain actions. Each agent has a different pattern of preferences. We could choose to let this group be a group of specialists in the strict sense, i.e. agents that specialise in only one skill. However, prior experiments have indicated that the performance would become very low because all the agents would become highly bored (Zoethout et al, 2006a). This would imply rather trivial results of our experiments. Therefore, we chose a setting in which the agents were free to self-organise task allocation whenever they want to, which opens the possibility of task rotation.

In the group of generalists, all the agents have skills with the same expertise and motivational values. Furthermore, the agents have to perform all the different actions that a task consists of instead of just one or two. This constraint implies that the group of generalists actually is not a self-organising group because they simply do not have the freedom to self-organise. However, as the results will show, the agents in the group of generalists do not have any incentive to change their allocation. Thus, although we constrained the agents to certain behaviour, they do not want it any either way.

5.2.4 Model and hypotheses

We study the expertise, the motivation and the performance of both groups in relation to the task dynamics. Figure 20 gives an overview of the model in relation to the experiments that we have conducted.

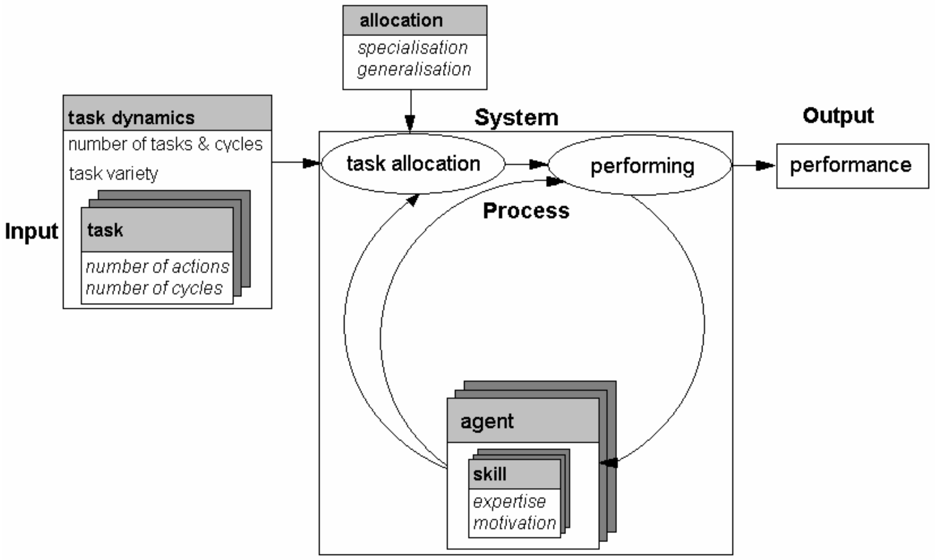

Figure 20: Model

The model is an input-process-output model. The input is a task. As we stated above, a task consists of a number of actions and a number of cycles. We studied different types of task dynamics by manipulating the number of tasks & cycles and the task variety. After the task has entered the system the process of task allocation starts. We manipulated the process of task allocation in such a way that we may speak of two groups, specialists and generalists. Both groups have different ways of allocating the task. The task allocation depends on two sets of variables, the task dynamics and the expertise and motivation of the agents. Expertise and motivation, being process variables, are mutually influenced by the process of performing the task, a process that also depends on the task allocation. The output of the system is the task that has been performed. The performance of the system, which is the dependent variable, indicates how well the group of agents has fulfilled the task.

Based on the classical principles of system theory formulated by Ashby (1956), we hypothesise a relation between task dynamics, specialisation/generalisation, and performance time as depicted in Figure 21:

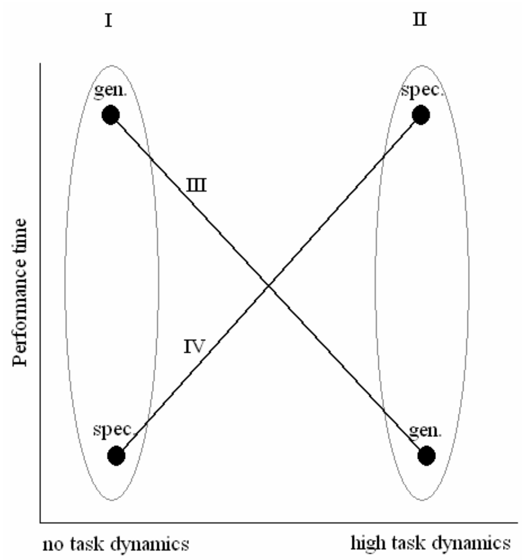

Figure 21: Overview of the hypotheses

The x-axis depicts two conditions of task dynamics, i.e. no task dynamics and high task dynamics. The y-axis depicts the performance time, i.e. the reverse of performance. Gen. and spec. respectively refer to the group of generalists and specialists. The hypotheses are depicted as roman numbers.

Hypothesis I: In a condition without task dynamics, the group of specialists will outperform the group of generalists

The rationale behind this hypothesis is based on the notion that specialisation leads to higher expertise, which implies better performance. Since the group of specialists is able to develop a certain level of task rotation, boredom will not have negative effects on performance.

Hypothesis II: In the condition of high task dynamics, the group of generalists will perform better than the group of specialists

Because the group of generalists is more flexible, they are able to adapt more quickly to task changes. As task changes occur more frequently, this flexibility will be more beneficial. Therefore:

Hypothesis III: The performance of the group of generalists will increase when task dynamics increase

Since the group of generalists benefits from their flexibility and the group of specialists benefits from its expertise, we also state:

Hypothesis IV: The performance of the group of specialists will decrease when task dynamics increase

5.3 Experimental Design

5.3.1 Variables and design

The experimental design is aimed at identifying the performance of groups of specialists versus generalists in varying conditions of task dynamics. We used 2 independent variables. The first independent variable is called task dynamics. This construct is split up into two components. The first component is called number of tasks and cycles. This component consists of 4 conditions in which the agents have to perform 200 cycles: in the first condition the agents have to perform 1 task of 200 cycles, in the second they have to perform 2 tasks of 100 cycles, in the third 4 tasks of 50 cycles, and in the fourth 8 tasks of 25 cycles. The second component is task variety, in which we used 4 conditions: no variety, low variety, which indicates a change of 1 action from one task to another. Moderate variety refers to a condition in which a new task requires 3 new skills, and high variety if a new tasks demands 5 new skills. Since all the tasks consist of 5 actions, high variety implies that every new task differs completely from the former one. Of course, these conditions, are not applicable to the condition in which the agents must only perform 1 task, because task variety is defined as a difference between multiple tasks.

The second independent variable concerns the group of agents. We used two kinds of groups, a group of specialists and a group of generalists. We measured two process variables, i.e. expertise and motivation and one dependent variables, i.e. performance time. Performance time refers to the time it takes to complete a task (see also Formula 15). We used the performance time of the slowest agent to indicate the performance time of the group. Since the agents perform together, the slowest agents indicates when the task is finished.

Table 6 summarises the independent variables into a research design:

Table 6: Research design for the group of specialists and the group of generalists

	Variety			
Tasks - Cycles	0 (no)	1 (low)	3 (moderate)	5 (high)
1 - 200	Condition 1	-	-	-
2 - 100	-	C2	C3	C4
4 - 50	-	C5	C6	C7
8 - 25	-	C8	C9	C10

The table shows all 10 conditions for both groups which means that there are 10 x 2 = 20 conditions. C2, C3, etc. refer to condition 2, condition 3, etc. By definition, the condition of 1 task of 200 cycle has no variety because variety is defined as a property of multiple tasks.

5.3.2 parameter values and initial settings of the agents

In the experiments we used the following parameter values:

The system consists of 5 agents

A task consists of 5 actions

All of the tasks together take 200 cycles

The initial values of expertise and motivation are equal

The maxima of both motivation and expertise are set on 25

The motivation – and expertise thresholds are set on 10

The learning speed is 100

The forget speed is 3

The boredom rate is 100, the recovery rate is 100

After conducting numerous trials, we found out that these values mark the parameter space in which the processes occurred that we want to study (see also Zoethout et al., 2006a).

In the condition of generalists, all the skills of one single agent have the same initial value, but as Table 7a shows, the agents have different skill values:

Table 7a: Initial values of the generalists

Skill	Agent 1	Agent 2	Agent 3	Agent 4	Agent 5
1	*14*	*15*	*16*	*17*	*18*
2	*14*	*15*	*16*	*17*	*18*
3	*14*	*15*	*16*	*17*	*18*
4	*14*	*15*	*16*	*17*	*18*
5	*14*	*15*	*16*	*17*	*18*

Since the initial values of expertise and motivation are equal, the values in Table 7a represent both. In the condition of specialists, the values of the skills of the agents are all different whereas each agent has another best skill:

Table 7b: Initial values of the specialists

Skill	Agent 1	Agent 2	Agent 3	Agent 4	Agent 5
1	*14*	*15*	*16*	*17*	*18*
2	*15*	*16*	*17*	*18*	*14*
3	*16*	*17*	*18*	*14*	*15*
4	*17*	*18*	*14*	*15*	*16*
5	*18*	*14*	*15*	*16*	*17*

After the agents start working, the values presented in Table 7a and 7b change due to learning and boredom and, therefore, these values only apply to the first task. New skills, i.e. skill 6, 7, etc., that are required to perform next tasks, all start with the value of 14: this value offers the agents enough expertise to perform the task, but it is nonetheless relatively low because the agents just start to use it.

5.4 Results

For every condition we analysed the performance time of both groups as well as the process variables expertise and motivation. But first we will discuss how the different conditions are related to the total performance time after the task. In this way we hope to find an answer to the question, what group performs the best under what conditions.

5.4.1 Total performance time

Figure 22a and 22b depict the total performance time of both groups in all conditions:

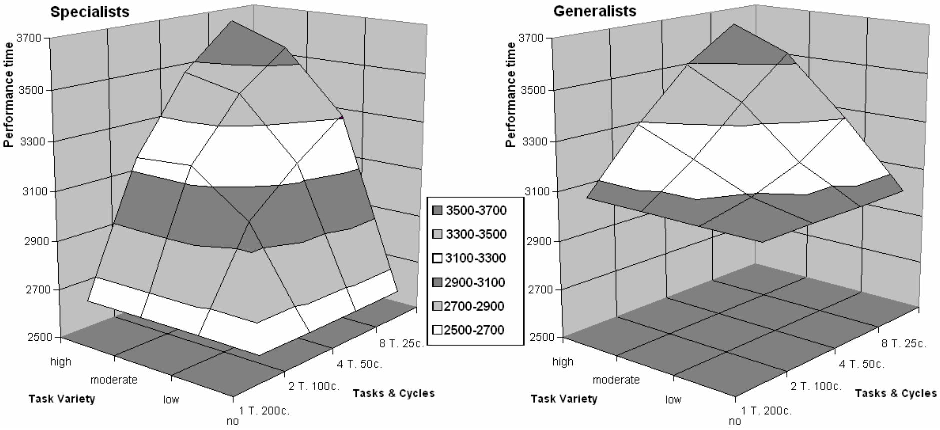

Figure 22a (left) and 22b (right): Total performance time of specialists and generalists in all conditions

1 T. 200 c. means 1 task of 200 cycles, etc. The performance time in the figures is the sum of the performance time of every cycle. We see that the performance time after performing 1 task of 200 cycles (bottom left line) is the same as the performance time in the condition of no task variety (bottom right line) for reasons we stated in section 5.3.1.

By looking at the condition of no variety (bottom lines), i.e. C1, we observe that the performance time of the specialists is much lower than the performance time of the generalists.

The reason for this comports with our intuition and common sense and is explained more elaborately in the next section about the underlying processes: When there is no task variety, driven by boredom, each specialist will alternately use his best two skills. In this condition, specialists will attain high levels of expertise for both these skills and, therefore, performance will be highest (i.e., performance time lowest). Generalists on the other hand will use all skills and reach a lower level of expertise, and, therefore, performance will be substantially lower (i.e. a higher performance time).

In the other conditions, the following holds for generalists and specialists: The higher the level of task dynamics, the more new skills are needed to complete new tasks that enter the system. When new tasks enter the system, generalists will start using all new skills that are needed for task completion. This will lower their average expertise and, therefore, their performance. The more new skill are needed to complete new tasks, the lower the average expertise level will become and, therefore, the lower the performance. Figure 22b shows that the relationship between variety and performance is more or less linear for generalists (see also Figure 22c). When a new task that requires novel skills enters the system, specialists prefer to continue using the two skills they already used. If one of these skills is not required anymore, they prefer to start using the skill for which they have the highest expertise. For none of the agents this will be the new skill. Therefore, they will postpone the use of the new skill until all the other actions are completed. At that moment all the agents start using the new skill. Their expertise is low and there are no opportunities for rotation anymore, which causes boredom to become high. As a result performance time increases. This increase in performance time is largest when we move from the condition without variety to a condition with low variety (see also Figure 22c).

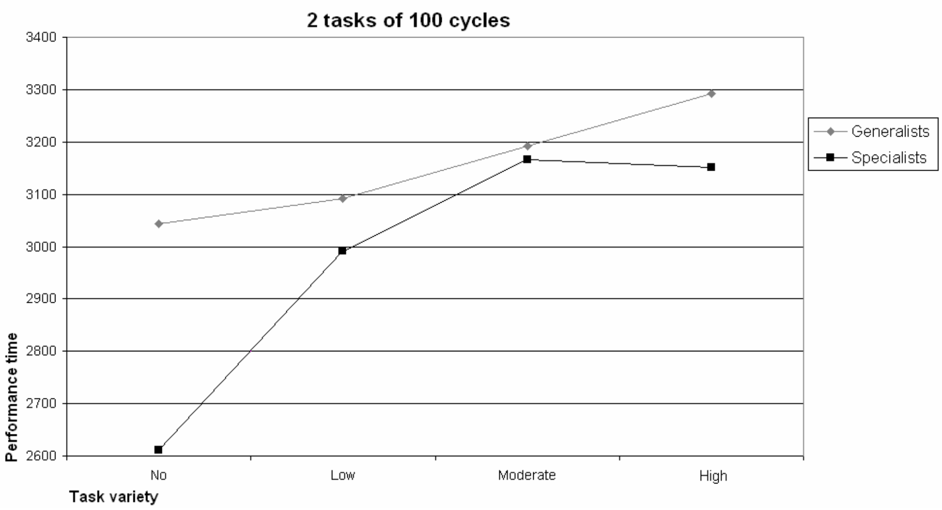

Figure 22c: Performance time differences of both groups

Figure 22c compares the specialists and the generalists at the condition of 2 tasks of 100 cycles. As we already stated, the relationship between task variety and performance time is more or less linear for the generalists. The specialists however, show a maximum at moderate variety: as task variety increases, new tasks that require more novel skills enter the system more often and more old skills are not needed anymore. This has a negative effect on both the specialists and the generalists. However, this gives the specialists less leeway to postpone the use of new skills, and, moreover, when more new skills come in, more opportunities to rotate will remain. This has a positive effect on the motivation. Hence, for both groups, expertise decreases when task dynamics increase. For the group of specialists, motivation increases as task dynamics increase. This implies that no task dynamics cause the highest expertise and high task dynamics cause the highest motivation. Therefore, somewhere between both conditions, i.e. moderate task variety, the agents profit the least from both benefits, which explains the curvilinear relationship task variety and performance time of the specialists.

5.4.2 Acceptance of the hypotheses

On the basis of these results, we will now accept or reject the hypotheses as formulated in section 5.2.4

Hypothesis 1: In the condition of no task dynamics, the group of specialists will perform better than the group of generalists

According to the Figures 22 a and 22b, in the condition with no task dynamics the group of specialists clearly performs better than the group of generalists. Therefore, hypothesis I is supported.

Hypothesis II: In the condition of high task dynamics, the group of generalists will perform better than the group of specialists

Although the dynamics of the group of specialists is different from the group of specialists, we may conclude that the performance of specialists in case of variety is still somewhat better than that of the generalists. Besides, we noticed that specialists are more vulnerable to task variety, especially if we move from no to a low level of variety. These findings do not support hypothesis II.

Hypothesis III: The performance of the group of generalists will increase when task dynamics increase

and

Hypothesis IV: The performance of the group of specialists will decrease when task dynamics increase

In both groups the increase of task dynamics led to a worse performance. This supports hypothesis IV but not hypothesis III.

5.4.3 Underlying processes

We already mentioned the differences in performance time of the specialists and generalists. To comprehend these differences in more detail, we now focus on the processes of task allocation in relation to performance time, expertise and motivation. We describe three conditions in detail that are typical for the description in the former section, i.e. condition 1, 3 and 10 (Table 6).

5.4.3.1 Condition 1: 1 task of 200 cycles, no variety

With regards to the group of specialists, in all conditions the expertise development can be characterised by the following steps. First the agent starts with its best skill, whereas the expertise of the other skills decreases. Second, because boredom increases, the agents start rotating between their best – and second best skill. Both skills more or less reach their maximum with respect to expertise. In the condition of no task variety, these two steps describe the whole process.

The group of generalists develop all skills evenly. Because the agents develop more skills at the same time, i.e. 5 instead of 2, their expertise develops slower, which cause a lower final expertise value (and, therefore, a lower performance).

In the group of specialists, according to the boredom function, skills that are actually used lead to a motivational decrease. When the agent stops using a particular skill, according to the recovery function, motivation increases again. But although the motivation stabilises, in the beginning the motivation is higher (17 and 18) than at the end of the task (round 15).

In all conditions, the motivation of the group of generalists remains the same. Since they use all their skills instead of just one they do not develop any boredom. Therefore, they do not have any incentive to rotate actions.

The performance time of the group of specialists during the process of task allocation and performance starts with an increase caused by the motivational decrease. Second, a slight decrease occurs, that is caused by task rotation, and increase of expertise: Task rotation causes the motivation not to decrease further, but to stabilise instead, whereas the increase of expertise causes the performance time to finally reach a minimum of about 12.5 (see Figure 23a):

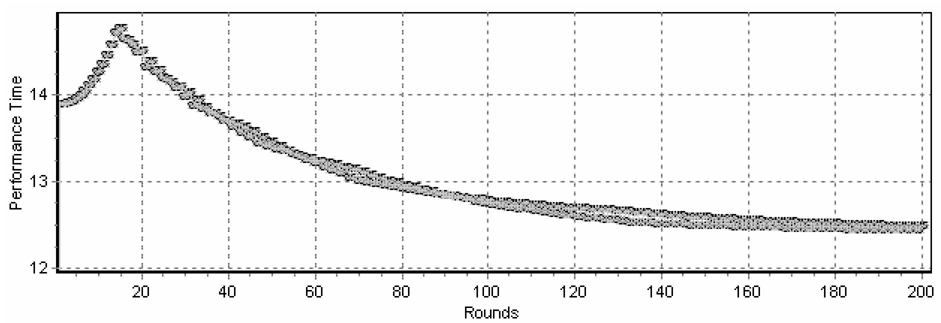

Figure 23a: Performance time of all agents in the group of specialists in condition 1

As a result of their initial values, the performance time of all agents is identical. Therefore, the Figure only depicts one curve. The maximum at the 16^{th} round represents the point from which the agents start rotating the task.

Since the group of generalists does not develop boredom, their performance time only decreases from the start as a result of the increase of expertise(see Figure 23b):

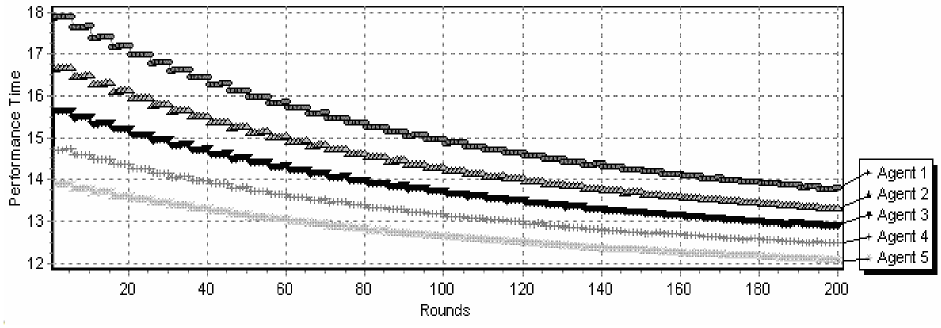

Figure 23b: Performance time of all agents in the group of generalists in condition 1

However, the performance time of the generalists decreases slower than the performance time of the specialists. This is the most important reason why the lowest value of the performance time of the generalists is higher than that of the specialists. For the sake of completeness, we mention one other reason: The performance time of both groups is determined by the worst agent. In the group of specialists, all agents perform equally whereas all agents only use their best two skills. In the group of generalists however, as Figure 23b depicts, all agents perform differently whereas all agents use all their skills. This implies that the performance of the specialists is determined by their best skills, whereas the performance of the generalists is determined by their worst skills. Thus, although Figure 23b may suggest that performance time may decrease further, it will never be lower than the performance time of the group of specialists.

Therefore, in this condition, the combination of expertise increase and motivational stability causes the group of specialists to be the best performing group. This clearly supports hypothesis I that states that in a condition without task dynamics, specialists perform better than generalists.

5.4.3.2 Condition 3: moderate variety, 2 tasks of 100 cycles

In the group of specialists, during the first task, i.e. the first 100 rounds, the process of expertise development is identical to condition 1: First the agent starts with its best skill, whereas the expertise of the other skills decreases. Second, because boredom increases, the agent starts rotating between its best – and its second best skill. Third, the second task starts, that requires 3 new skills. Depending on the specialisation of the agent, this may lead to three different situations: First, when the new task does not forces the agent to use any new skills, he simply proceeds rotating between the same ‘old’ skills. Second, when the new task forces the agent only to use new skills, he starts using them in a way similar to the first task: first he uses his best skill and when boredom reaches a certain level, he rotates between his best and second best skill. Third, when the new task permits the agents to use only one of his ‘old’ skills, the agent proceeds in using only this skill until boredom forces him to rotate to his second best skill. Figure 24a depicts the expertise development of agent 5 that resembles the third situation:

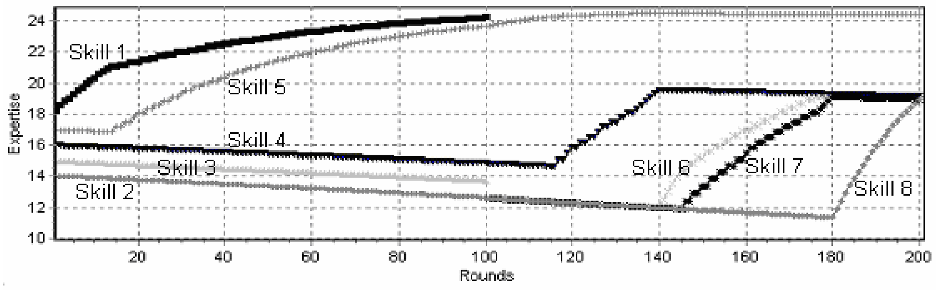

Figure 24a: Expertise of agent 5 of the group of specialists in condition 3

The x-axis shows the number of rounds. The y-axis shows the expertise values. When the second task enters the system at the 101st round, agent 5 does not use skill 1 anymore, because the part of the task that requires this skill is completed. Until the 116th round, he only uses skill 5 (originally his second best skill). Then he starts with his skill that is second best at that time. The fourth and last phase starts when either one of the actions the agent is working on has been completed. In this phase each agent consecutively picks up the remaining actions that correspond to their next best skills and finally they finish with their worst skill.

The expertise development of the group of generalists is much simpler (see Figure 24b): the first task shows an identical process as in the former condition. When the second task enters the system the 'old' skills of each agent develop further, identical to condition 1. The three 'new' skills (6,7, and 8) start at a lower value but then develop in the same way as the old skills do (see Figure 24b):

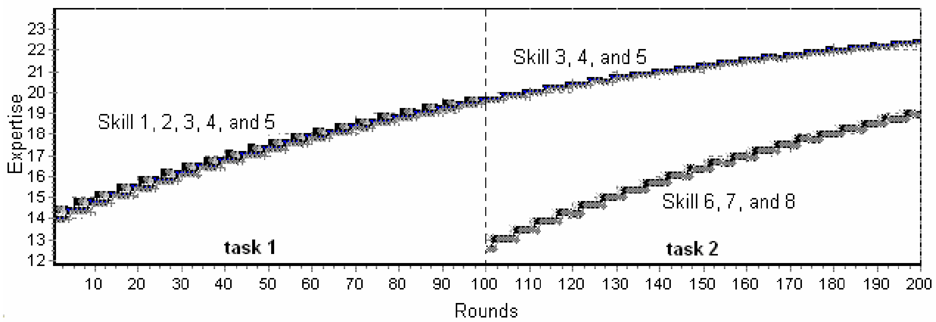

Figure 24b: Expertise of agent 5 in the group of generalists in condition 3

During the first task, the way in which the motivation of the group of specialists develops, is the same as in the former condition (condition 1). From the start of the second task, motivation decreases when agents only use one skill for a certain period (see Figure 24c):

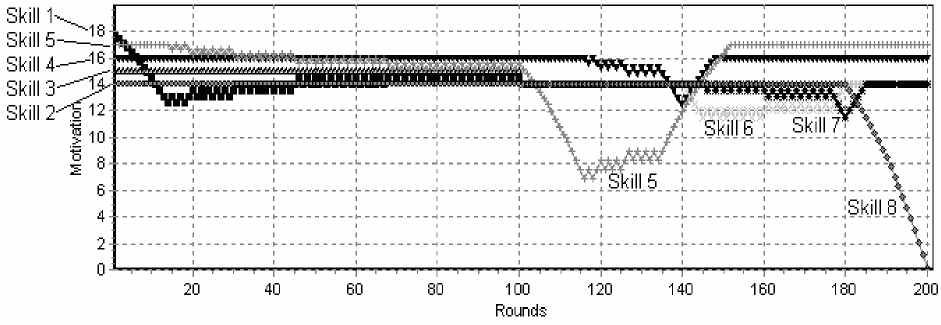

Figure 24c: Motivation of agent 5 in the group of specialists in condition 3

Figure 24c depicts the motivational development of agent 5. Skills that are not used are represented as a flat line because the motivation of these skills are not influenced. The first motivational decrease occurs right after the second task enters the system when the agent uses only one 'old' skill. This decrease continuous until the 116th round. From that point on , task rotation stabilises the motivation. From the 180th round, the agents work on their last action and do not have any possibilities left to rotate, which results in a dramatic motivational decrease of all agents.

This just mentioned motivational decrease has a strong effect on the performance time of the specialists resulting in an increase from the 180th round (see Figure 24d):

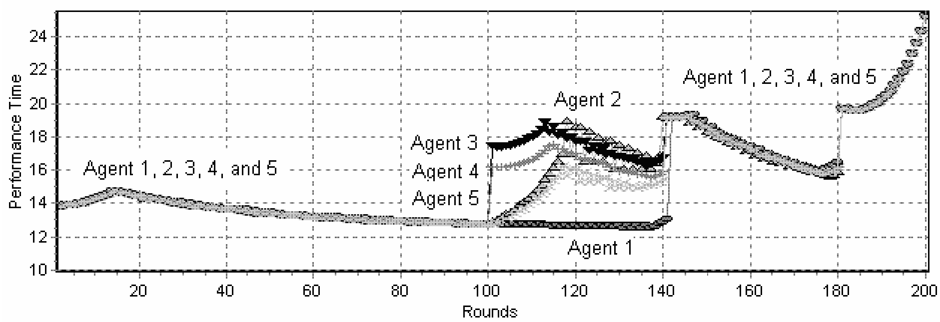

Figure 24d: Performance time of all agents in the group of specialists in condition 3

From the 101st round, performance time increases because some agents start using new skills, for which they have a lower expertise. These new skills cause the same performance time development as in the first task: at the beginning, boredom causes the agents to perform worse, until they start rotating their actions. This effect is repeated at the 140th round, when the agents start with a new action. From the 180th round, the agents only work on one action, with no opportunities for task rotation. This causes boredom and related to this, an increase of performance time.

In the group of generalists, the 'new' skills, that are related to the actions of the second task, cause local maxima for every 5 rounds in which the agents perform one cycle of the task: every agent uses 2 ‘old’ skills that cause low performance time, and 3 ‘new’ skills that cause high performance time (see Figure 24e):

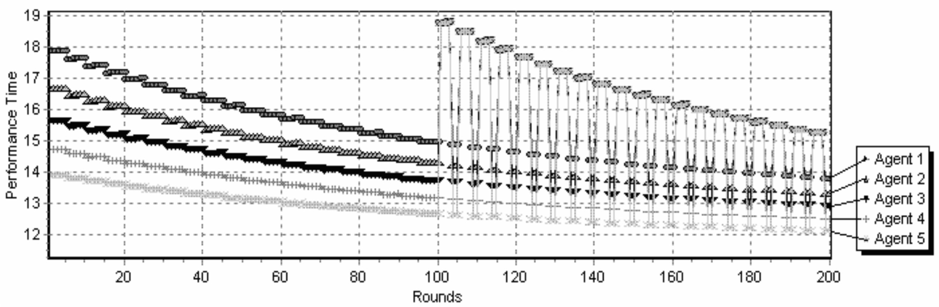

Figure 24e: Performance time of all agents in the group of generalists in condition 3

In Figure 24e we see that the expertise of the ‘new’ skills increases in the same way as the expertise of the ‘old’ skills.

Therefore, although the expertise of the specialists reaches a high level at an earlier stage than the expertise of the generalists, the motivational decrease of the specialists compensates for this. As a result, in these conditions the performance of both groups is about the same (see also Figure 22c). For the specialists, higher levels of task dynamics result in a lower motivational decrease, whereas lover levels of task dynamics result in a higher expertise development.

5.4.3.3 Condition 10 high variety, 8 tasks of 25 cycles

This condition represents the highest level of task dynamics in which the agents perform tasks that are all different. During the first task the agents demonstrate the same behaviour as in the other conditions. From the second task on however, all skills are new. Since we defined new skills as skills with low expertise and motivation, this results in groups of agents, both specialists and generalists, with identical skills. This implies that the specialists are not really specialists anymore. Nevertheless, since the specialists are free to allocate the task whereas the generalists are forced to use all their skills consecutively, both groups behave differently.

From the second task on the expertise development of the group of specialists during the execution of one task can be described in three phases: First, all agents have identical motivation and expertise for all their skills and, therefore, simply start with their first skill. At that moment the increase of expertise is just a bit smaller than the decrease in motivation, and, consequently, they immediately start task rotation after action 1. Second, as soon as action 1 and 2 have been finished, the agents start performing action 3 and 4 in the same way. Third, the agents perform action 5. This process is repeated for every task. Because the expertise of the skills decreases when the time in which the agents used them increases, every next task starts with lower expertise.

As regards the group of generalists, every task shows an increase of expertise, whereas every next task start somewhat lower than the former. Although this also holds for the specialists, the generalists perform all actions consecutively.

From the second task on the motivation of the group of specialists is stable, which is caused by task rotation. It ends in a decrease because during the performance of the last action there is no possibility to rotate anymore. Every next task develops in the same way. This process is comparable to the motivational process of condition 3 that also alternates motivational stability and decrease. The most important difference is that a task of 25 cycles does not offer that much time for boredom, which limits the motivational decrease. Therefore, the performance time of the specialists is mainly determined by the slow decrease of expertise (see Figure 25a):

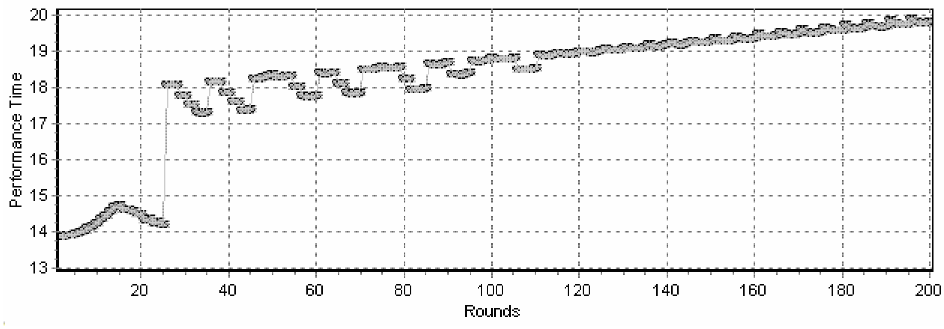

Figure 25a: Performance time of all agents in the group of specialists in condition 10

As a result of their initial values, the performance time of all agents is identical. Therefore, the Figure only depicts one curve. We see that in the beginning the performance curve clearly depicts the starting points of every new task. From the 100th round however, this diminishes. This is caused by a slight change of task allocation: because of the learning function (see Formulas 14a and 14b), agents with identical expertise values allocate their tasks differently when their expertise and motivation changes, because the change of expertise and motivation is a function of the 'old' expertise and motivation. For instance, when a skill with a value of 15 is used, it changes in a different way then a skill with a value of 20, which may lead to a different task allocation.

The performance time of the group of generalists follows the expertise development: every task shows a slight decrease, whereas every next task starts with a performance time that is a bit higher. This is caused by the decrease of expertise during the time in which the agents did not use the skills related to the new actions (see Figure 25b):

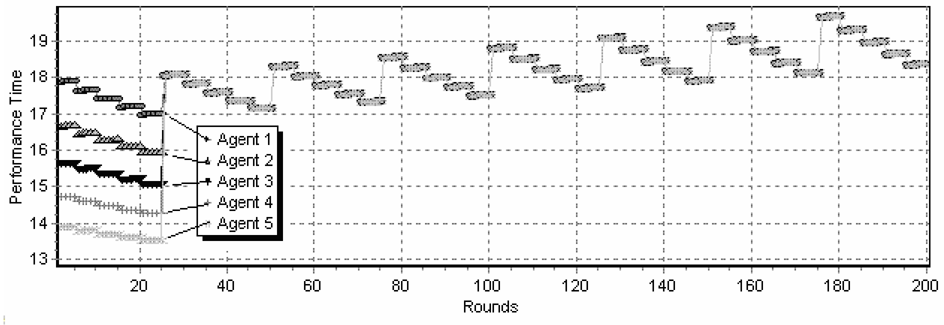

Figure 25b: Performance time of all agents in the group of generalists in condition 10

Although the process of task allocation is different, the performance of both groups is about the same. In both groups, the high level of task dynamics limits the possibility of boredom, but decreases the development of expertise dramatically. Therefore, these conditions cause the worst performance for both groups. This conclusion supports hypothesis IV that states that the performance of the specialists decreases when task dynamics increase, but it rejects hypothesis III that states the opposite for the generalists. In sum, in case of extremely high levels of task variety specialists have no time specialise and generalists have no time to become generalists, and, therefore, their processes and performance become quite similar.

5.5 Conclusion and Discussion

Although the total performance time, as depicted in Figures 22a and 22b, clearly indicates the differences between both groups, the underlying processes can explain how these results relate to the hypotheses. In general we found two opposite effects: The first effect concerns the expertise development that decreases in both groups when task dynamics increase. This effect is the main cause of the mountain-like shape of Figures 22a and 22b, with the highest task dynamics being on top. The second effect concerns the increase of motivation when task dynamics increase. This effect only holds for the group of specialists because the motivation of the generalists remains unaffected. This effect does not apply to the condition of no task dynamics.

As our description of condition 3 states, a combination of both effects leads to the curvilinear relationship between both groups as described in Figure 22c.

These results only support our hypotheses concerning the performance of the specialists: According to our study, in a condition with no task variety specialists perform better than generalists. Further, the performance of specialists decreases when task dynamics increase. As we stated, the main reason behind this is the decrease of expertise development that does not allow the specialists to maximise their performance. This reason comports with the rationale behind the hypotheses, which is based on the principles as formulated by Ashby (1956). This principle implies that specialists will benefit from their high expertise in situations with no task dynamics.

Generalists on the other hand should outperform specialists in conditions with higher task dynamics because they are more flexible than specialists and, therefore are better able to adapt to changing situations. In a way the flexibility of generalists compensates for the profit of high expertise that they lack in conditions with low task dynamics. Hence both specialists and generalists have their own specific benefits. Thus, we hypothesised that specialists should be better with no variety and generalists with high variety.

But why do our results not comport with the second half of this principle? Why does the performance of the generalists differ from what we hypothesised? The results show that with high task dynamics the generalists are as much the victim of the decrease of expertise development as the specialists. But the generalists do not compensate this decrease with their benefit, their flexibility, at all. On the contrary, their flexibility it is more or less shared by the specialists: in conditions of higher task dynamics, the specialists appear to behave more like generalists! This is partly caused by our assumption that new skills start with low expertise. During the first stage of our experiments the generalists and the specialists differ in two ways. The way expertise and motivation is distributed among the agents differs (see Table 7a and 7b) and the way they allocate tasks differs. The more new tasks differ from previous ones, the more the starting situation for both groups in terms of expertise and motivation become similar, due to our assumption of low expertise for new skills for both groups. Besides, a higher level of task dynamics implies the need for more new skills. Therefore task dynamics decrease opportunities to specialise, and, therefore, the performance of specialists become more or less similar to the performance of generalist in case of very high levels of task variety. But another cause for the decrease of specialisation can be found in the effect, that agents in the specialist group tend to start with their best skill, then their second best, etc. and finish with their worst skill. This effect shows up in all conditions with task dynamics. This implies that the specialists do not restrict themselves to their own speciality, but finally help each other to finish the team task. Actually, in our experiments the specialists were not true specialists because they were able to use all the skills, although their expertise was higher for some than for others. Moreover, due to boredom, they tend to specialise in two skills, and, therefore, probably can be better typified as 'minimal generalists' than true specialists. Previous studies have shown that a low level of multifunctionality mostly outperforms a high level of multifunctionality as well as a situation with no multifunctionality at all (i.e., when there are only 'true specialists'; Molleman & Slomp, 1999; Van den Beukel, 2003). Both reasons indicate that high task dynamics cause a group of specialists to self-organise into a group of generalists. Hence, this study does not confirm a classic relationship between performance and task dynamics with respect to generalisation and specialisation. Nevertheless, the results support the underlying proposition that a situation with high task dynamics asks for generalists and a situation with low and no task dynamics needs specialists.

An important difference between the specialists and the generalists in this study is the freedom to self-organise. The specialists are free to re-allocate the task whenever they feel the need to do so, which may result in a shift for specialisation to generalisation. The generalists on the other hand, do not have any freedom to self-organise, since they must simply perform all actions consecutively. Although this implies that the specialists are able to stabilise their motivation by task rotation, the generalists do not have to stabilise anything because they are already satisfied with their situation. This brings up one of the fundamental question of management science: should an organisation or team be forced to fit into a design or should we give it enough freedom to self-organise? Of course we cannot give a definite answer to this. Given the freedom to self-organise, no matter what the original team structure is, the team will manage to reshape itself according to the demands of its environment. But then again, as the group of generalists demonstrated, sometimes the absence of freedom leads to a well-ordered structure in which motivated workers perform quite well. In fact, our study indicates that the difference between design and self-organisation does not lead to spectacular performance differences, except for the condition of no variety. This suggests that we might put the importance of team design into perspective. All in all, we may conclude that it is difficult to catch group dynamical processes in a simple system theoretical description. However, the use of multi-agent simulation based on psychological theory may certainly help to understand how basic individual characteristics are related to complex group dynamics.

Further, we may ask ourselves whether or not our conclusions have real empirical value? First of all, we did not limit our experiments by using agents with cognitive properties only, but used a model in which we combined a simplified cognitive architecture with variable motivational states. Because of this, the specialists behaved differently from more traditional agent models: The agents developed task rotation and tasks that were highly repetitive cause a larger motivational decrease than tasks that were less repetitive. This effect appears in real life as well. The specialists were able to reduce their motivation loss by performing two actions instead of one and the generalists showed no motivation loss at all. In a way, this comports with the findings that workers performing a task as a whole feel more motivated (e.g. Hackman & Oldham, 1980).

The empirical value of the parameter values, i.e. learn, forget, boredom, and recovery speed, as well as the relation between motivation and expertise, we have used, can be questioned. We simply selected a parameter space that produced behaviour that we could study: For instance, a higher forget speed would result in a group of agents that is not able to perform anymore. In future research we could vary the boredom and recovery rate. But empirical studies that indicate such parameter values are yet to be done. This study may serve as a start to formalise group dynamical processes concerning expertise, motivational and performance development and relate its parameter values to specific types of tasks.

But the strongest contribution of our research concerns the question it evokes regarding group dynamical processes in work teams: how can differences between specialists and generalists be explained in terms of processes of task allocation? What processes should be altered to increase their performance? How is the motivation of workers related to traditional approaches concerning specialisation and generalisation? Although this study does not answer all these questions fully, it may elucidate our view and changes our perspective on traditional problems.