Auton Agent Multi-Agent Syst (2008) 16:75–94
DOI 10.1007/s10458-007-9022-9

Task dynamics in self-organising task groups: expertise, motivational, and performance differences of specialists and generalists

Kees Zoethout · Wander Jager · Eric Molleman

Published online: 27 October 2007

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 task variety leads to specialisation and high task variety leads to generalisation, in general, performance is better when task variety is low. Further, in case of no task variety, specialists outperform generalists. In case of moderate variety the opposite is true. With high task variety, since there is no space for any expertise and motivational development, the behaviour of specialists and generalists becomes more similar, and, consequently also their performance.

Keywords Multi-agents system · Specialisation-generalisation · Self-organisation · Psychological theory · Task allocation

1 Introduction

Should one hire specialists or generalists to maximise group performance? This question still puzzles personnel managers and organisational scientists alike. Well known by

K. Zoethout (✉) · E. Molleman
Faculty of Management and Organisation, University of Groningen, P.O. Box 800, Groningen 9700 AV, The Netherlands
e-mail: k.zoethout@rug.nl

E. Molleman
e-mail: h.b.m.molleman@rug.nl

W. Jager
Faculty of Economics, University of Groningen, P.O. Box 800, Groningen 9700 AV, The Netherlands
e-mail: w.jager@rug.nl

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 human factors such as expertise, learning, motivation and boredom have been found to affect group performance (for an overview see, for example, [27]). Whereas the effects of separate variables—or limited combinations—have been empirically investigated by many researchers, e.g. [1,9,21,22,25], 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 [8]. In this paper 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 originates from complexity theory and employs multi-agent simulation as a methodology to describe higher order developmental and adaptation processes in terms of local interactions [2,7,12,18], see also [4,8]. In general, most simulation models concerning processes within task groups or teams focus on abstract computational and mathematical descriptions, e.g. [19,20] or use a so-called bounded rationality approach, e.g., [13]. 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 [9], see also [25].

In this study we want to investigate the relationship between motivational processes and task performance by using a formalised psychological description of these variables and their interrelatedness. Because of the lack of such a formalised description, we developed WORKMATE, a computer simulation model that simulates self-organising processes of task allocation. WORKMATE is able to analyse the way a workgroup adapts to changes of the tasks it must perform. There are a lot of approaches related to task descriptions and performance [11,22,24]. Our description is based on Wood [26] who states that task changes can be described in terms of component complexity of tasks, i.e. changes of the contents of tasks in time. The adaptation process of the group is described in terms of task allocation, a process being affected by task characteristics and team member characteristics. With respect to the latter we focus on the expertise and the motivation of the individual team members.

We compared two groups, a group of generalists, i.e. agents using all skills necessary to perform a task, and a group of specialists, i.e. agents using only a small subset of skills necessary to perform a task. We focussed on the way they both adapt to task changes and on their performance.

In the next section of the paper, Sect. 2, we focus on the theories, models and their formalisation, which form the basis of WORKMATE. 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 third section describes the experimental design and the parameter settings. Next, in Sect. 4, we describe the results and we end up with conclusions (Sect. 5) and a discussion (Sect. 6).

2 The model

WORKMATE is a deterministic discrete event based simulation program for simulating self-organising processes of task allocation that is developed in DELPHI6 (see also [29,30]). In this section we shortly describe the theoretical framework WORKMATE is based on.

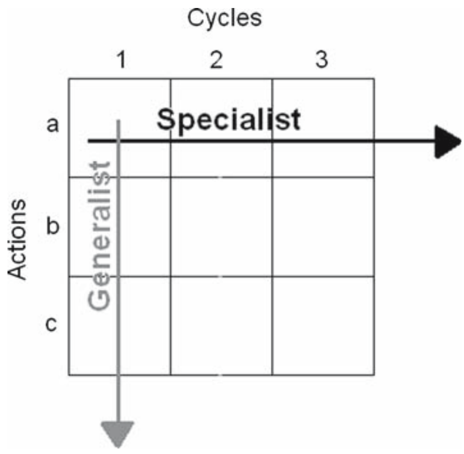

Fig. 1 Representation of a generalist and specialist performing a task

2.1 Tasks and task dynamics

A task is considered as a set of *actions*, i.e. subtasks, in such a way that each action requires a single skill [11,22,24]. Each action must 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 Fig. 1). This paper deals with the relation between task allocation and performance under conditions of task dynamics. The use of the concept of *task dynamics* implies that agents must perform multiple tasks in a serial way. According to our model the performance of multiple tasks in a serial way means that agents may only start with a new task if the former task has been finished.

Task dynamics refer to the speed in which tasks change over time [26], see also [22]. 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 the actions required to complete the task. *Number of tasks and cycles* refers to the size, i.e. number of cycles, of a single task and the number of tasks that must be performed. For instance, given a task variety of 1, performing 2 tasks, each consisting of 100 cycles, implies lower task dynamics than performing 8 tasks, each consisting 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 have used seven new skills.

Figure 1 represents a task consisting of 3 actions, a, b, c and 3 cycles, 1, 2 and 3. Thus, 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, which in combination are sufficient to complete the task. Specialisation is defined as a clear preference of the agents for a subset of the skills, i.e. one or two skills which do *not* cover all the skills necessary to complete the whole task (see Fig. 1).

The horizontal arrow represents the task allocation process of a single agent being a specialist. He performs all cycles of a single action, needing only one single skill. The vertical arrow represents the allocation process of a generalist. Each cycle he performs all actions, needing all skills necessary to perform one whole task.

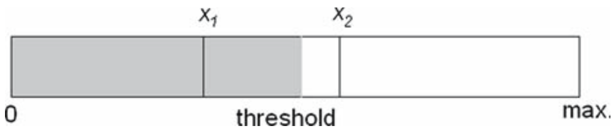

Fig. 2 Calculation of the I-node and You-node

2.2 The multi agent system

2.2.1 Agents

Important components that determine group performance are expertise and motivation [25], see also [21]. These components can be considered as components of individual agents. According to our description 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 consisting of two variable components: expertise and motivation. Analogous to human memory, skills are represented as either passive or active [16]: The so-called Long-Term Memory store (LTM) consists of a large set of passive skills. Once an agent starts to perform a task, a subset of these skills corresponding to the actions which the task consists of, is activated. The set of active skills corresponds to the Short-Term Memory store (STM). This implies that when performing a task, the motivation and expertise of *the agent* is determined by the active skills (see also [29]). On the basis of these active skills the agents individually decide first whether to use these skills or not. Second, on the basis of this decision, they allocate the task. We will describe both steps.

First, to every activated skill applies that the values of expertise and motivation determine whether or not that skill will be actually used to perform the corresponding action. This is represented by using thresholds. If the expertise exceeds its threshold a skill is ‘good enough’ and if the motivation exceeds its threshold a skill is ‘nice enough’. The decision whether a particular skill is sufficient to use depends on the following rule:

IF Expertise > ExpertiseThreshold **AND** Motivation > MotivationThreshold
THEN I DO **ELSE** YOU DO

This decision rule means that an agent wants to perform a particular action (I DO) if both the expertise and the motivation related to the skill to perform that action are sufficient. The logical reverse of this rule is that, if either the expertise or the motivation is insufficient, the agent does not want to perform that particular action (YOU DO). In this way every agent chooses a subset of actions he would like to perform.

The decision whether or not to perform an action is depicted by means of two nodes, an I-node and a You-node, each having values between 0 and 1. The values of these nodes are functions of expertise, motivation, their thresholds, their maxima, and a parameter α that determines the balance of expertise and motivation. The initial choice depends on these nodes according to the rule that I DO means that the value of the I-node is bigger than the value of the You-node, and You DO means the other way around.

The I- and You-nodes are calculated as follows: expertise and motivation have values (x) between 0 and a maximum with a threshold somewhere in between (see Fig. 2). A value of x between 0 and the threshold (x_1 in Fig. 2) refers to the You-node and a value of x between the threshold and the maximum (x_2 in Fig. 2) refers to the I-node. The I-node will be 1 if the x is maximal, and will approach 0 if the x reaches the threshold. The You-node will be 1 if

x is 0, and will approach 0 if x reaches the threshold. Therefore, for the I-node, the height of the threshold is subtracted from x. For the You-node, x is subtracted from the threshold. Then, these values are divided by their maxima to make up the I- and You-values between 0 and 1.

We can distinguish three situations. In the first situation there is insufficient expertise, which will consequently lead to a YOU DO choice. In this situation, no matter how high the motivation is, the skill will not be influenced since the agent cannot perform that particular action anyway. The values of I and You can simply be described as:

$$I = 0 \tag{1a}$$

$$You = 1 \tag{1b}$$

The second situation refers to expertise and motivation both exceeding their thresholds, which leads to the initial choice of I DO. In this situation the I-node is a function of the expertise (e) and motivation (m), their thresholds (th_e, th_m), their maxima (e_{max}, m_{max}), and a parameter α [0,1] that indicates 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 α is set to 0.5.

$$I = \alpha^* \frac{e - th_e}{e_{\max} - th_e} + (1 - \alpha)^* \frac{m - th_m}{m_{\max} - th_m} \tag{2a}$$

$$You = 0 \tag{2b}$$

The third situation refers to sufficient expertise and insufficient motivation. According to the decision rule, this situation would lead to an initial choice of YOU DO. Here the I-node$=0$ and the You-node is determined by motivation (m), the threshold (th_m) and the balance parameter α [0,1]. However, this situation may not provide a sound basis to start the allocation process. If we look at the initial choice of the other agent, again, we may distinguish three situations: insufficient expertise leading to YOU DO, sufficient expertise and motivation leading to I DO, and sufficient expertise and insufficient motivation. In the first and second situation the choice is clear. In the third situation both agents have insufficient motivation and sufficient expertise. It will not be plausible to determine the choice of these agents solely by their motivation, as the decision rule suggests. Therefore, we have defined the I-node as a function of expertise rather than labelling it zero, and the You-node as a function of motivation:

$$I = \frac{e - th_e}{e_{\max} - th_e} \tag{3a}$$

$$You = \frac{th_m - m}{th_m} \tag{3b}$$

2.2.2 Interaction

Now that we presented a formal description how agents individually choose a subset of actions they would like to perform, we continue to describe how the self-organising process of task allocation takes place. Self-organisation refers to the process in a system leading to the emergence of global order within this system without the presence of another system dictating this order [6,10]. This implies that the agents allocate the tasks by means of local interaction. Agents interact when the individual preferences of all agents imply that there are more agents sharing the same preference regarding a particular action than the number of cycles these actions can be performed. The interaction process implies that the agents

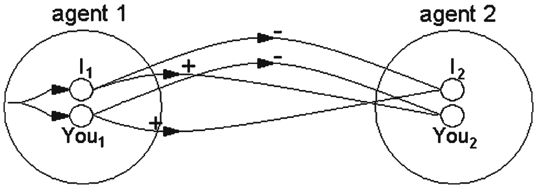

Fig. 3 Excitation (+) and inhibition (−) of two agents

are trying to influence the other agents in such a way that the other agents will reach a complementary state with respect to their own (see also [29]).

We based the process of influence on a neural network analogy by using excitatory and inhibitory connections. Excitatory connections increase the tension of the node that is affected and inhibitory connections decrease its tension (see Fig. 3).

Figure 3 depicts how agent 1 influences agent 2. The I-node (I_1) of agent 1 inhibits the I-node (I_2) and excitates the You-node (You_2) of agent 2. The You-node (You_1) of agents 1 inhibits the You-node (You_2) and excitates the I-node (I_2) of agents 2. As a result of this influence, the values of the I- and You-nodes change. These changes depend on the initial values of the I- and You nodes of both agents: the higher the value of the sending agent, the higher the potential influence. This potential influence is limited by the values of the I- and You node of the receiving agent. Excitation will decrease as the values approach their maxima and inhibition will decrease as the values approach 0. Therefore, excitation can be described as follows:

$$y_1 = y_0 + x_0(1 - y_0) \tag{4a}$$

y_0 represents the old value of the receiving node, y_1 the new value, and x_0 the old value of the 'sending' node. x_0 is multiplied by $(1–y_0)$ to make sure that the value of y_1 does not exceed the maximum value of 1. Inhibition can be described as:

$$y_1 = y_0 - x_0 y_0 \tag{4b}$$

x_0 is multiplied by y_0 to make sure that the value of y_0 does not exceed the minimum of 0.

The agents simultaneously influence each other. If two nodes with exact the same value influence each other, it is likely that nothing will happen. This means that the difference between two nodes should be included into the excitatory as well as the inhibitory functions:

$$Diff_{II} = Abs(I_1 - I_2) \tag{4c}$$

$$Diff_{YouYou} = Abs(You_1 - You_2) \tag{4d}$$

I_1, I_2, You_1, and You_2 correspond to Fig. 3. $Diff_{II}$ equals the absolute value of $(I_1 - I_2)$ and $Diff_{YouYou}$ equals the absolute value of $(You_1 - You_2)$. We combine the Eq. 4a–d to describe all excitatory and inhibitory connections as shown in Fig. 3:

$$I_2 := I_2 - \iota I_1 I_2 Diff_{II} : \quad I_1 \text{ inhibits } I_2 \tag{5a}$$

$$You_2 := You_2 - \iota You_1 You_2 Diff_{YouYou} : \quad You_1 \text{ inhibits } You_2 \tag{5b}$$

$$You_2 := You_2 + \varepsilon(1 - You_2) I_1 Diff_{II} : \quad I_1 \text{ excitates } You_2 \tag{5c}$$

$$I_2 := I_2 + \varepsilon You_1 (1 - I_2) Diff_{YouYou} : \quad You_1 \text{ excitates } I_2 \tag{5d}$$

ε and ι [0,1] represent parameters that can set the height of the excitation (ε) and the inhibition (ι) (for an elaborate description, see [28,29].

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.

2.2.3 Task performance

When all actions are being allocated, the process is being completed and the agents start performing the task. As a result of this, their 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. Besides, learning a new skill when performing a task can be motivating in itself and high motivation can increase learning speed (e.g. [5]).

An important characteristic of most learning curves is that they reach a maximum asymptotically [15]. 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 β [0,1] that determines the learning speed:

$$e_{(t+1)} = e_t + \beta \frac{e_{\max} - e_t}{e_{\max}} \tag{6a}$$

Since motivation positively affects the learning speed, Eq. 6a can be elaborated as:

$$e_{(t+1)} = e_t + \left(\beta + \gamma \frac{m - th_m}{m_{\max} - th_{m\ I}} \right) \frac{e_{\max} - e_t}{e_{\max}} \tag{6b}$$

when the motivation exceeds its threshold and

$$e_{(t+1)} = e_t + \left(\beta - \gamma \frac{th_m - m}{th_m} \right) \frac{e_{\max} - e_t}{e_{\max}} \tag{6c}$$

When the motivation does not exceed its threshold, γ [0,1] refers to the parameter that determines the impact of motivation on the learning speed, the fraction refer to the Eqs. 3a and 3b.

Forgetting is the inverse of learning. Since forgetting only applies to skills that are not used, motivation does not play a role here. Therefore, forgetting can be described as the inverse of formula (6a):

$$e_{(t+1)} = \frac{(e_t - \delta) e_{\max}}{e_{\max} - \delta} \tag{6d}$$

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

In real life an enormous range exists between learning and forgetting speed of different tasks. Motor tasks such as truck driving are, once learned, never forgotten, whereas others, such as sorting need to be maintained. Therefore, in the experiments, the balance between learning and forgetting 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 (6b) describes the motivational decrease related to boredom and formula (6a) represents the motivational increase related to the recovery from boredom. In this case the parameters β

and δ respectively describe the recovery and the boredom speed. Furthermore, learning speed influences motivation. According to (6a) learning speed (L) can be defined as:

$$L_t = \beta \left(\frac{e_{\max} - e_t}{e_{\max}} \right) \tag{7a}$$

where (L_t) equals e_{t+1} minus e_t. The final equation that describes the motivational processes is then:

$$m_{(t+1)} = m_t + \zeta L_t + \beta \frac{m_{\max} - m_t}{m_{\max}} \tag{7b}$$

(L_t) is added with a parameter ζ [0,1] to determine the amount of influence of expertise to motivation.

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, t_{action}, 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:

$$t_{perf.} = \sum_{i=1}^{n} \frac{t_{action\_i}}{\lambda \frac{e_i}{e_{\max}} + (1 - \lambda) \frac{m_i}{m_{\max}}} \tag{8}$$

In the present study, the agents perform the actions simultaneously. This means that the time it takes to perform the total task, $t_{perf.}$, is determined by the slowest agent.

2.2.4 Specialisation and generalisation

In our experiments we use two groups, a group of specialists and a group of generalists. Each group has the skills to perform the whole task. The group of *specialists* consists of agents that are all specialised in a particular part of the task, only using one or two skills. 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 are being specialised in only one skill. However, prior experiments have indicated that the performance would become very low because all the agents would become highly bored [29]. 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. For example, 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. This is called *task rotation* [30].

Whereas the group of specialists answers to the description we made in the previous sections, the properties of the *generalists* are being defined as agents that *must* perform all actions a task consists of consecutively. This constraint implies that although the expertise and motivational changes of the individual agents is subject to the same processes as the specialists, the group of generalists itself is not a self-organising group because they simply do not have the freedom to self-organise. Although such top-down rotation systems are not unrealistic [17], we realise that our operationalisation is a rather strict one. By defining generalists in this sense the generalist will not have the opportunity to develop themselves into specialists so that we are able to study differences between both groups over time.

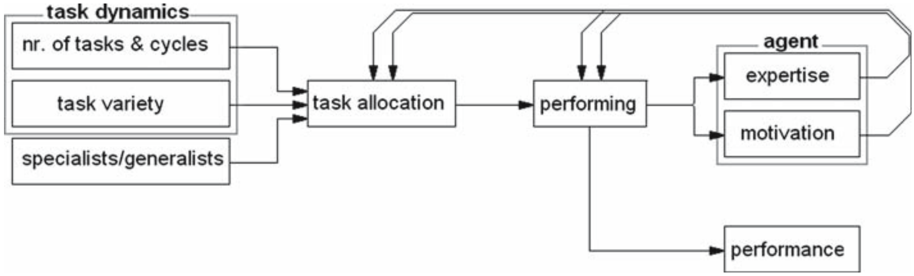

Fig. 4 Model

2.3 Model and hypotheses

We study the expertise, the motivation and the performance of both groups in relation to the task dynamics. Figure 4 gives an overview of the model in relation to the experiments that we conducted.

The independent variables are *number of tasks and cycles* and *task variety* (together representing the *tasks dynamics*), and the formalisation of agents as *specialists* respectively *generalists*. The latter refers to the two groups that we compare. These groups differ on *task allocation*: either all agents perform all actions (generalists), or they perform a small portion of it (specialists). When the agents actually are *performing* tasks, their *expertise* and *motivation* change, which may result in changes of the task allocation process and, again, in performing tasks. The *performance* of the system, which is the dependent variable, indicates how fast the group of agents has fulfilled all the tasks.

Based on the classical principles of system theory formulated by Ashby [3], we hypothesise a relation between task dynamics, specialisation/generalisation, and performance time as depicted in Fig. 5.

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.

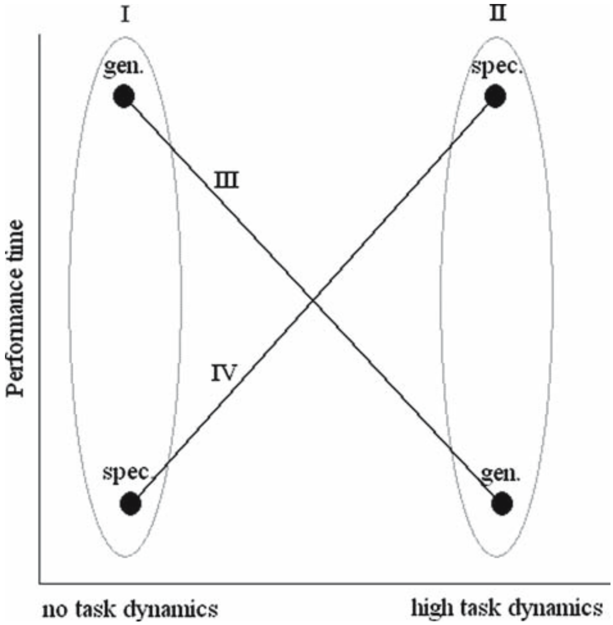

Fig. 5 Overview of the hypotheses

The group of specialists benefits from its expertise. However, this will become a disadvantage if tasks change more frequently and require new skills. Therefore we hypothesize:

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

3 Experimental design

3.1 Variables and design

Following the model (Fig. 4), the experimental design is aimed at identifying the performance of groups of specialists versus generalists in varying conditions of task dynamics. *Number of tasks and cycles*, being its first component, consists of four conditions in which the agents must perform 200 cycles: in the first condition the agents must perform 1 task of 200 cycles, in the second 2 tasks of 100 cycles, in the third 4 tasks of 50 cycles, and in the fourth 8 tasks of 25 cycles. *Task variety*, being the second component, has four levels: no variety; low variety, indicating a change of 1 action from one task to another. Moderate variety refers to a new task requiring 3 new skills, and high variety if a new task demands 5 new skills. Since all 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 perform only 1 task, because task variety is defined as a difference between multiple tasks (see Table 1 for an overview of the conditions).

Table 1 shows all 10 conditions for both groups which means that there are $10 \times 2 = 20$ conditions. C2, C3, etc. refer to condition 2, condition 3, etc. In condition C3, for example, the team has to complete two tasks, each with 100 cycles. After 100 cycles, the first task is

Table 1 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

completed and the team starts working on the second tasks that requires 2 old skills and 3 new skills. By definition, condition 1, the condition of 1 task of 200 cycles, has no variety because variety is defined as a property of multiple tasks.

We measured two process variables, i.e. expertise and motivation, and one dependent variable, i.e. performance time. Performance time refers to the time it takes to complete a task (see also Formula 8). We used the performance time of the slowest agent to indicate the performance time of the group. Since the agents perform together, the slowest agent determines when the task is finished.

3.2 Parameter values and initial settings of the agents

In the experiments we used the following parameter values:

1. The system consists of 5 agents
2. A task consists of 5 actions and each action requires one unique skill
3. All the tasks together take 200 cycles

This small number of agents and actions will help us to comprehend the principles and mechanisms behind the task allocation and task performance of both groups. Each condition runs for 200 cycles, which is sufficient to simulate the processes we want to study.

4. The initial values of expertise and motivation are equal

In all experiments the expertise and motivation of the agents starts with the same value. By varying these values, we could have created a wide variety of conditions. For the aim of this study, however, this was not necessary.

5. The maxima of both motivation and expertise are set on 25
6. The motivation—and expertise thresholds are set on 10

Since we defined learning and motivational processes by means of a maximum and a threshold, we need a value for both. The actual values are not that important, as long as they create enough space for the agents to develop the processes that we want to study.

7. The learning speed is 75%
8. The forget speed is 3
9. The boredom rate is 75%, the recovery rate is 75%

A learning speed of 75% offers the possibility for further increase due to positive motivational effects (see also 11), while it still produces noticeable learning effects. A forget speed of 3 has clear effects on task allocation without causing the agents to lose skills that have not been

Table 2 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

used for a period of time. Since expertise and motivation have equal effects on performance time, we have chosen a boredom rate equal to the learning speed. In analogy to the initial values of expertise and motivation, we have chosen for equal values for the boredom- and recovery rate.

10. The amount of influence of expertise on motivation (see formula 7b) is 100%
11. The amount of influence of motivation on learning speed (see formula 6b and c) is 100%

Tests indicated that variation of the amount of influence of expertise on motivation does not lead to changes with respect to the performance differences of specialists and generalists. Variation of the amount of influence of motivation on learning speed has a small effect on specialists only: more influence results in a better performance. This means that the height of these parameters does not matter that much. However, their effects on the development of motivation and expertise itself will be more visible when their values are higher.

In the group of generalists, all skills of the single agents have the same initial value of 16. In the group of specialists, the values of the skills of the agents are all different whereas each agent has another best skill. Table 2 shows the initial values of the group of specialists.

After the agents start working, their initial values change due to learning and boredom and, therefore, these values only apply to the first cycle of a task. New skills, i.e. skill 6, 7, etc., that are required to perform next tasks, all start with the value of 16.

4 Results

For every condition we analysed the performance time of both groups as well as the process variables expertise and motivation. First we will describe the processes of task allocation in relation to performance time, expertise, and motivation. Then we will describe the results of these processes by discussing the total performance time of both groups *after* all the tasks have been completed.

4.1 Development of expertise, motivation and performance

In all conditions, the task allocation process of the group of *specialists* can be characterised by the following steps. First, the agents start with their best skill, which implies a decrease of their expertise of the other skills. While each agent continues using this skill, expertise increases. Motivation remains stable at first, because of the simultaneous opposite effects of learning and boredom. As long as the increase in expertise is larger than the decrease in motivation, performance time decreases (i.e. performance improves). Second, because the increase of expertise decreases and motivation decreases further, motivational decrease exceeds the increase of expertise and performance time increases (i.e. performance becomes worse). This decrease continues until the third step in which the motivational decreases causes the agents to start rotating between their best—and second best skill. Both skills more or less reach their maximum with respect to expertise, whereas the motivation stabilises. This causes the performance time to slowly decrease further.

With regards to the group of *generalists*, all conditions by definition imply that each agent performs all actions and develops all skills evenly. This implies an increase in expertise, although noticeable smaller than for the specialists, because the specialists develop only two skills, whereas the generalists use all five. This increase in turn causes a slower decrease in motivation. 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 slow increase of both expertise and motivation causes the performance time to decrease from the start.

These processes describe what happens during the performance process of one single task. In conditions with multiple tasks, these processes are more or less repeated, with one important difference: dependent on the level of task variety, the agents use skills that are partly 'old', i.e. they were used in the previous task, and partly 'new', i.e. they are used for the first time. For the agents in the group of *specialists*, this may lead to three different situations: First, when the new task does not force the agent to use any new skills, the agent simply proceeds rotating between the same 'old' skills. Second, when the new task forces the agent to use only new skills, the agent starts using them in a way similar to the first task: first the agent uses his best skill and when boredom reaches a certain level, the agent 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 the agent to rotate to the second best skill.

In the condition of no variety (see also Table 1) only the first situation holds, whereas in the conditions of high variety only the second situation applies. A combination of the three situations occurs in the remaining conditions of low and moderate variety.

The group of *generalists* behaves much simpler. When a new task enters the system the 'old' skills of each agent simply develop further. The 'new' skills start at a lower value but then develop in the same way as the old skills do.

Although this description holds for each condition, we want to elaborate on the condition of high variety. Because the agents only use 'new' skills for new tasks, this condition has some distinctive characteristics. First of all the *specialists* are free to allocate the task whereas the *generalists* are forced to use all their skills consecutively, which results in different behaviours of both groups.

From the second task on, during the execution of a task, the expertise development of the group of *specialists* can be described in three phases: First, all agents have identical motivation and expertise for all their skills. This means that they simply start with the first skill and then rotate between their first and second skill. Second, as soon the actions related to these skills have been finished, the agents start performing the third and fourth action in the same way. Third, the agents perform the last action. This process is repeated for every task. An important implication of this process is that the specialists are not really specialists anymore since they use all their skills, being all identical at the beginning. Nevertheless, they are still free to allocate the task. On the other hand, the *generalists* are forced to use all their skills consecutively. This implies that every task shows an increase of expertise. Just like the specialists, every next task starts somewhat lower than the former, but both groups behave differently.

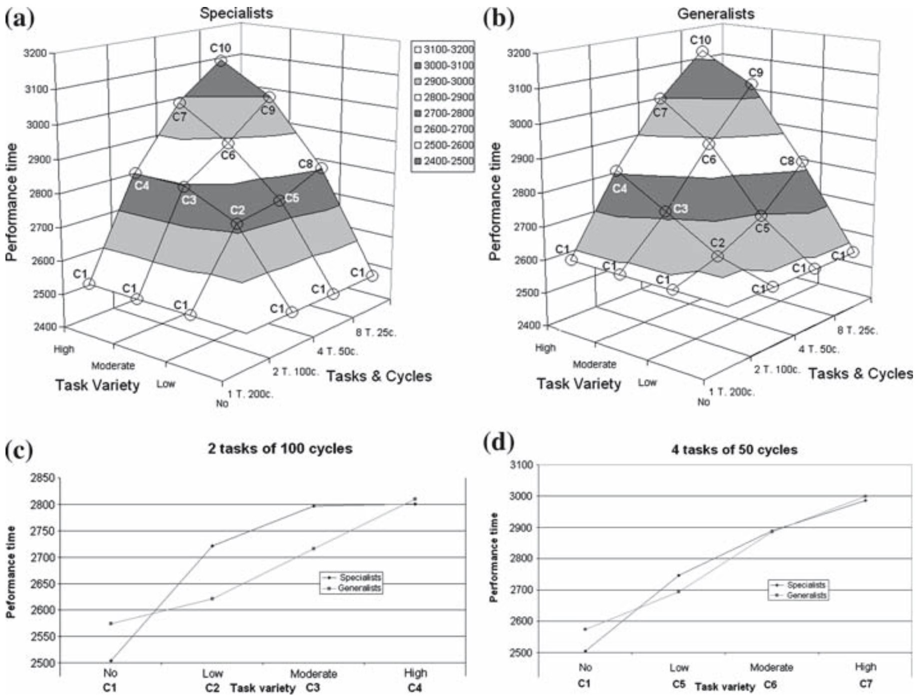

Fig. 6 (**a** left) and (**b** right): Total performance time of specialists and generalists in all conditions as represented in Table 1; (**c**) Performance time differences of both groups in the condition of 2 tasks of 100 cycles; (**d**) Performance time differences of both groups in the condition of 4 tasks of 50 cycles

4.2 Total performance time

Now that we have described the changes of expertise, motivation, task allocation and performance time during the process of task performance, we will continue to describe the total performance time of both groups in all conditions (see Fig. 6a and b).

In both figures, “1 T. 200c.” means 1 task of 200 cycles, “2 T. 100c.”, means 2 tasks each with 100 cycles, etcetera. The x-axis, Task Variety and the y-axis, Task & Cycles, correspond with the rows and columns of Table 2 respectively. The conditions that Table 1 describes are depicted in the figure as C1, C2, etc. The z-axis indicates the performance time, being the sum of the performance time of every cycle. We see that the performance time after performing 1 task of 200 cycles, i.e. C1 (bottom left line) is the same as the performance time in the condition of no task variety (bottom right line) for reasons we already stated in Sect. 3.1.

When comparing the Fig. 6a and b we observe three characteristics: first, in the condition of no variety (bottom lines), i.e. C1, the performance time of the *specialists* (2,504) is lower than the performance time of the *generalists* (2,574). This finding supports Hypothesis 1: 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 lower (i.e. a higher performance time).

Second, performance time increases in both groups when either task variety or the number of tasks & cycles increases. As tasks become shorter, both groups have less time to increase their expertise. As tasks become more different, both groups must use more ‘new’ skills for which the learning process still has to start.

Third, the group of generalists shows a relation between task dynamics and performance that is more or less linear, based on the principle that we described above. However, the group of specialists shows a relation that is not linear, causing a performance time that is *higher* than that of the generalists in conditions of low and moderate variety (see Fig. 6c).

Figure 6c depicts the performance time of both groups in the condition of 2 tasks and 100 cycles. We see that in the condition of no variety (C1), the group of specialists outperforms the group of generalists. In the condition of high variety (C4), both groups are more or less equal, but in between (C2 ‘low’ and C3 ‘moderate’) the specialists perform worse than the generalists. To comprehend these differences we will take a closer look at the underlying processes regarding the group of specialists.

When a new task requiring novel skills enters the system, *specialists* prefer to continue using the two skills they already used, because of their high expertise. If one of these skills is not required anymore, they prefer to start using the skill for which they have the highest expertise. In general, this will not be the skill related to the new actions. Therefore, they will postpone the use of these skills. In conditions of low task dynamics (C2), as Fig. 6c depicts, this implies less opportunity for task rotation at the end of the task. Because of this, boredom will not be compensated and performance time will increase. In the condition of moderate (C3) variety, but especially in the condition of high (C4) task variety, the leeway of the specialists to postpone the use of new skills decreases and, moreover, when more new skills come in, more opportunities to rotate will remain. This has a positive effect on their motivation. Hence, for both groups, expertise decreases when task dynamics increase. For the group of *specialists*, motivation *increases* as task dynamics increase because of the possibilities for task rotation. This implies that no task dynamics cause the highest expertise, high task dynamics cause the highest motivation. Therefore, somewhere between both conditions, the agents profit the least from both benefits, which explains the curvilinear relationship between task variety and performance time for the specialists.

Figure 6d depicts the performance time differences of both groups in the conditions of 4 tasks of 50 cycles (C1 ‘no task variety’; C5 ‘low’; C6 ‘moderate’; C7 ‘high’).

We notice that the differences between both groups are much smaller: Because there is less leeway to postpone the use of new skills, the motivational decrease of the specialists is much smaller. Hence, the performance differences of both groups decreases. With 8 tasks of 25 cycles, i.e. conditions C8, C9, and C10, both groups follow about the same linear curve, whereas the specialists perform slightly better than the generalists.

4.2.1 Acceptance of the hypotheses

On the basis of these results, we now accept or reject the hypotheses as formulated in Sect. 2.4.

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

According to the Fig. 6a–d, in the condition with no task dynamics the group of specialists 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 differ from the group of specialists, we may conclude that the performance of both groups in case of high task dynamics is more or less equal. 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 Conclusion

Figure 6a–d, clearly shows that the difference in performance between the generalists and specialists depend on task variety. The underlying processes can explain how these results exactly relate to the hypotheses. In general we found three 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 Fig. 6a and b, with the highest task dynamics being on top. The second effect concerns the delay of using new skills when there is task variety. The third effect pertains to the increase of motivation when task dynamics increase. These last two effects only hold for the group of *specialists* because they are allowed to postpone the use of new skills and because the motivation of the *generalists* remains unaffected in case of task dynamics. For the generalist there is only one force effective, that is less development of expertise when task variety increases. Therefore, for the generalists, the relationship between task variety and performance is more or less linear. For the specialists all three effects are present. The first effect, the lower increase in expertise when task variety increases, results in a higher performance time for specialists too. The second effect, the delay of the use of new skills, is especially there when task variety is low to moderate. The third effect, the increase in motivation, becomes stronger when task variety increases. The three effects together result in the curvilinear relationship between task variety and performance for specialists. Task dynamics are created by the number of new skills (task variety) and the number of cycles per task. The curvilinear effect is most apparent if a task has many cycles (Fig. 6c). As task dynamics increase due to a decreasing number of cycles per task, the relationship between task variety and performance for specialists becomes close to that of generalists. The findings indicate that generalists outperform specialists if there are tasks dynamics, but only if these dynamics are small.

According to our study, in a condition with no task variety specialists perform better than generalists. 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 that is based on the principles formulated by Ashby [3]. 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.

But why are our results not in line with this principle when tasks dynamics are high? Why does the performance of the generalists differ from what we hypothesised? The results show that under high task dynamics both generalists as specialist show about an equal decrease in expertise. However, the generalists do not compensate this decrease with their flexibility. Instead, the specialists are forced to start behaving more like generalists by performing a larger portion of the actions sequentially.

Three causes may explain this. First of all, the more new tasks differ from previous ones, the more the starting situation for both groups becomes similar in terms of expertise and motivation, due to the assumption of identical expertise and motivation for new skills for both groups. Second, 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 becomes more or less similar to the performance of generalist in case of very high levels of task variety. The third 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 have to help each other to finish the team task. Actually, in the 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' [14,23].

Whereas the first cause can be considered as an artefact of the model, the other 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 a low to moderate level of task dynamics asks for generalists and a situation with no task dynamics specialists are preferred.

6 Discussion

On the basis of this study the question raises to what extent the results can be generalised to different settings. To answer this question we will focus on the different aspects of the model, the simulation platform we used and the implications of the results. We will end this section with a description of possible future research lines.

Concerning the model, we supposed that the most important factors that determine group performance are expertise and motivation. We did not take the concept of *coordination costs*, which is based on the interaction between the agents, into account, although this component certainly affects group performance [25], see also [21]. For example, coordination costs will decrease in case of specialisation and increase when task dynamics increase [29]. However, although more dynamics imply higher coordination costs, it is plausible that groups may learn how to deal with these dynamics by letting allocation rules emerge. The emergence of these rules can only be simulated if the underlying cognitive architecture of the agents offers the substratum to let them to. This means that the use of coordination costs as a variable to describe self-organising processes of task allocation would only be useful when it could be related to a more elaborate cognitive description of the agents. The parsimonious model that we use in the present study does not contain this architecture because we were mainly interested in expertise, motivation and task allocation being components that determine group performance. Therefore, we did not study the influence of coordination costs.

Second, 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 from pure specialisation towards a moderate type of generalist. 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 specialists are able to stabilise their motivation by task rotation, generalists do not have to stabilise anything because they do not experience any motivation problem.

This relates to a fundamental question in the area 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. On the other hand, as the group of generalists demonstrated, sometimes the absence of freedom implies a well-ordered structure in which motivated workers perform quite well. In fact, our study indicates that, except in the condition of no or low variety, the difference between design and self-organisation does not lead to spectacular performance differences. But then again, a design should limit coordination costs, which means that a definite answer to this question must be found in the balance between low coordination costs and flexibility.

A third discussion point is the number of agents that a group consists of. We used only five agents and one could wonder whether larger groups would generate other dynamics. First of all, we state that the use of such small numbers is based on the assumption that we first wanted to comprehend the basic mechanisms before continuing on studying more complex systems. But do larger numbers always increase the complexity of the processes? To answer this question we have to take a look at the most important characteristics of the allocation processes. These are task rotation and using skills in a declining order with respect to expertise. These characteristics, as well as their effects, will not be different in larger groups nor do we expect this to be conditional for other phenomena to emerge. Perhaps studies on coordination costs would require larger groups, but then again, that is something that is beyond the scope of this study.

Fourth, 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 more actions instead of just 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. [9]. All in all, 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.

Concerning the empirical value of the parameter values regarding the expertise and motivational processes, we state that we simply selected a parameter space that would result in interesting processes: For instance, a higher forget speed would result in a group of agents that will forget parts of their skill repertoire which makes performance impossible. A lower forget speed on the other hand, would not result in noticeable expertise loss, which is also not realistic. Further, identical boredom and recovery rates are also questionable. However, these parameters are difficult to adjust to real life situation because empirical studies that indicate such parameter values are yet to be done. Therefore, this study may serve as a start to formalise group dynamical processes concerning expertise, motivational and performance development, relating its parameter values to specific types of tasks.

With regards to the platform that we used, WORKMATE has been programmed in DELHPI6. In the future it would be useful to formalise WORKMATE on other platforms such as SWARM or NETLOGO, in order to test the robustness of the effects in a model-to-model context.

Finally, in analogy to the work of Dellarocas and Klein [7], we used a simple model that is more plausible than traditional models with rational agents or mechanisms that facilitate an optimum. With this model we described interactions between higher levels, i.e. societies or work groups, and lower levels of description, i.e. the individual agents. As their work, ours should be considered as a first step of exploration. Furthermore, ours needs to be tested much more before we may draw definite conclusions regarding the topic that we studied. We will present two future research lines:

First, in the present study we compared two groups of homogenous agents. This study poses the interesting question what would happen in a mixed group. Is there an optimal fit between task characteristics and the distribution of specialists and generalists? How does a mixed group perform in relation to the groups that we described in this paper? A second future line of research is the influence of newcomers on task allocation processes of existing work teams [31]. How do new specialists and generalists fit into an existing team? By manipulating task and team structure we could design different settings in which extra members are needed.

To end up, the key contribution of the current study concerns our systematic analysis of traditional problems concerning having specialists and generalists in teams by using theories on groups and task performance.

References

1. Arrow, H., & Crosson, S. (2003). Musical chairs: Membership dynamics in self-organized group formation. *Small Group Research, 34*(35), 523–556.
2. Arrow, H., McGrath, J. E., & Berdahl, J. L. (2000). *Small groups as complex systems: Formation, coordination, development, and adaptation*. London: SAGE Publications.
3. Ashby, W. R. (1956). Self-regulation and requisite variety. In *Introduction to cybernetics* (Chap. 11, pp. 202–218). London: Chapman & Hall.
4. Axelrod, R., & Cohen, M. D. (2000). *Harnessing complexity: Organizational implications of a scientific frontier*. New York: Basic Books.
5. Bandura, A. (1977). *Social learning theory*. New York: General Learning Press.
6. Dalenoort, G. J. (1989). The paradigm of self-organization: Studies of autonomous systems. In G. J. Dalenoort (Ed.), *The paradigm of self-organization* (pp. 1–22). London: Gordon and Breach.
7. Dellarocas, C., & Klein, M. (1999). Civil agent societies: Tools for inventing open agent mediated electronic marketplaces. In *Proc. workshop in agent-mediated electronic commerce (co-located with IJCAI'99)*. July, 1999.
8. Gilbert, N., & Troitzsch, K. G. (1999). *Simulation for the social scientist*. Buckingham: Open University Press.
9. Hackman, R. J., & Oldham, G. R. (1980). *Work redesign*. Reading: Addison-Wesley Publishing Company.
10. Heylighen, F. (1997). Self-organization, emergence and the architecture of complexity. In F. Heylighen, C. Joslyn, & V. Turchin (Eds.), *Principia cybernetica web (Principia Cybernetica, Brussels)*, URL:ftp://ftp.vub.ac.be/pub/projects/Principia_Cybernetica/Papers_Heylighen/Self-Organization_Complexity.txt

11. Hunt, R. G. (1976). On the work itself: Observations concerning relations between tasks and organizational processes. In E. J. Miller (Ed.), *Task and organization*. London: Tavistock Institute of Human Relations, Wiley & Sons.
12. Kitts, J. A., Macy, M. W., & Flache, A. (1999). Structural learning: Attraction and conformity in task-oriented groups. *Computational & Mathematical Organization Theory, 5*(2), 129–145.
13. Levine, J. M., Moreland, R. L., Argote, L., & Carley, K. M. (2005). Personnel turnover and team performance. United States Army Research Institute for the Behavioral and Social Sciences, Technical Report 1157.
14. Molleman, E., & Slomp, J. (1999). Functional flexibility and team performance. *International Journal of Production Research, 37*, 1837–1858.
15. Nembhard, D. A. (2000). The effects of task complexity and experience on learning and forgetting: A field study. *Human Factors, 42*(2), 272–286.
16. Newell, A., Rosenbloom, P. S., & Laird, J. E. (1989). Symbolic architectures for cognition. In M.I. Posner (Ed.), *Foundations of cognitive science*. MIT, Cambridge.
17. Ortega, J. (2001). Job rotation as a learning mechanism. *Management Science, 47*(10), 1361–1370
18. Panait, L., & Luke, S. (2005). Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems, 11*, 387–434.
19. Pynadath, D. V., & Tambe, M. (2002). The communicative multiagent team decision problem: Analyzing teamwork theories and models. *Journal of Artificial Intelligence Research, 16*, 389–423.
20. Pynadath, D. V., & Tambe, M. (2002). Team coordination among distributed agents: Analyzing key teamwork theories and models. AAAI Spring Symposium on Intelligent Distributed and Embedded Systems.
21. Steiner, I. D. (1972). *Group process and productivity*. New York: Academic Press Inc.
22. Tschan, F., & von Cranach, M. (1996). Group task structure, processes and outcome. In M. A. West (Ed.), *Handbook of work group psychology*. New York: Wiley.
23. Van den Beukel, A. L. (2003). Task flexibility in team-based work: A multilevel investigation into outcomes and practices. Doctoral Thesis, University of Groningen, Faculty of Management and Organization, Groningen.
24. Weick, K. (1979). *The social psychology of organizing*, 2nd ed. New York: McGraw-Hill.
25. Wilke, H. A. M., & Meertens, R. W. (1994). *Group performance*. London: Routledge.
26. Wood, R. E. (1986). Task complexity: Definition of the construct. *Organizational behavior and human decision processes, 37*, 60–82.
27. Yeatts, D. E., & Hyten, C. (1998). *High performing self-managed teams, a comparison of theory to practice*. London: Sage Publications.
28. Zoethout, K. (2006). Self-organising processes of task allocation: A multi-agent study. Doctoral Thesis Groningen: University of Groningen, Faculty of Management and Organization, Labyrinth Publication, Ridderkerk.
29. Zoethout, K., Jager, W., & Molleman, E. (2006). Formalizing self-organizing processes of task allocation. *Simulation Modelling Theory and Practice, special issue on simulating organisational processes, 14*, 342–359.
30. Zoethout, K., Jager, W., & Molleman, E. (2006). Simulating the emergence of task rotation. *Journal of Artificial Societies and Social Simulation*, 9–1, URL: http://jasss.soc.surrey.ac.uk/9/1/5.html.
31. Zoethout, K., Jager, W., & Molleman, E. Newcomers in self-organising task groups: A pilot study. In *Advancing social simulation: The first world congress*. Tokyo: Springer Verlag (forthcoming).