

MULTI-AGENT SIMULATION OF PURCHASING ACTIVITIES IN ORGANIZATIONS

Mark J.R. Ebben
Luitzen de Boer
Corina E. Pop Sitar

Department of Operations Management and Logistics
University of Twente
P.O. Box 217
7500 AE Enschede, THE NETHERLANDS

ABSTRACT

In this paper we present a Multi-Agent simulation model to investigate purchasing activities in an organizational environment. Starting point is the observation that the majority of purchasing activities in organizations is usually performed without any involvement of the organization's purchasing department. The purpose of the experiments is to investigate if and how certain factors determine the degree to which Purchasing professionals become involved in the purchasing of Non-Product Related (NPR) items and services. Among the factors investigated are: corporate purchasing policies, available information, and the nature of the various purchasing activities. Preliminary results show that the behavior of the multi-agent simulation model is an acceptable representation of reality. Furthermore, the results show the limits of a Purchasing department's added value and the important role of organizational learning in that respect. The paper provides directions for further research.

1 INTRODUCTION

A number of studies (U.S. Bureau of the Census, 1991; Fearon and Bales, 1995) indicate that the role of the Purchasing Department is much smaller in the purchase of nontraditional (non-product related) goods and services than commonly believed. These 'nontraditional' (non-product related) areas, such as insurance, consultancy, and advertising, are likely to be purchased without the involvement of professional buying practice and skills. In a study in 1993 the Center for Advanced Purchasing Studies (CAPS, Fearon and Bales, 1993) found that in a surprisingly large percentage of firms, the Purchasing Department had no input to the purchase of particular nontraditional goods and services. Given the often huge sums of money involved in NPR-purchasing and the supposed saving potential by bringing in professional Purchasing expertise, it

seems worthwhile to find out which factors apparently hinder this Purchasing involvement.

The literature about this subject is very limited. Most researchers who studied this subject used surveys or questionnaires as a methodology for their study. De Boer and Pop-Sitar (2001) contemplated on some theoretical explanations for the limited and often problematic involvement of the Purchasing department in NPR purchasing. They conclude that the problem is very dynamic and complex. That is why we want to investigate the usefulness of multi-agent simulation for studying the problem. What we want to study is:

- In what way and to which extent do certain factors influence the involvement of the purchasing department in the NPR purchasing activities? Examples of such factors are: corporate purchasing policies, different types of purchases, and the skills, goals and perceptions of the agents in the organization.
- How does the involvement of the purchasing department in the NPR purchasing change over time?
- How does Purchasing involvement correlate with the attainment of organizational goals?

Simulation could be a valuable tool in addition to the existing approaches because of the possibility to model and analyze dynamic and complex systems and control the assumptions and the conditions of the simulation model, such as corporate policies, and investigate the effects of these changes in a controlled environment. In the simulation model each department can be naturally represented by an agent, which executes jobs that arise within the department.

The paper is organized as follows. Section 2 describes the problem setting and basic assumptions used in the simulation model. In section 3 we discuss the different

types of agents in the multi-agent simulation model. Section 4 and 5 present the simulation model and the results of preliminary experiments. After discussion of the results in section 6, we draw conclusions in section 7.

2 PROBLEM SETTING

Before we describe the problem in more detail we give two definitions of the term “agent”. According to Wooldridge and Jennings (1995) an agent is a computer system that has the following properties:

- autonomy
- social ability
- reactivity
- pro-activeness

Franklin and Graesser (1996) state that “an autonomous agent is a system situated within and part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to effect what it senses in the future”. The agents in our model can make their own decisions, based on their knowledge and information about the environment, and interact with other agents in order to realize their goals.

We introduce three types of agents: one Company Agent (CA), which represents company management, one Purchasing Agent (PA), which represents the purchasing department or purchasing officer, and multiple Internal Customer Agents (ICA), which represent the remaining departments or business units in the organization. There is a relation between an ICA and the PA. There is no relation between ICAs, i.e. we assume they do not communicate about purchasing issues. All ICA and the PA have a relation with the CA.

Every ICA is assigned the following properties: (1) a specific level of knowledge about the function and technical properties of a class of purchased items (content skills) and knowledge about the process-side of purchasing (process skills) (2) a specific target level for his own goal for such a class (3) a perception of the required skills for performing a job and (4) a perception of the skills of the purchasing agent. These aspects will be discussed in more detail in the next section.

There are three distinctive ways for the ICA to complete a job: perform the complete job himself; outsource it to the PA; or cooperate with the PA. The agents make local decisions based on the information and knowledge gathered during the process. As a foundation for the decision process in the simulation model we used Cyert and March (1963). Their work provides a set of clear and empirically grounded mechanisms for modeling decision-making:

1. Quasi-resolution of conflict by sequential attention to various organizational goals

2. Uncertainty avoidance by using gradually changing routines for decision-making
3. Problem-driven search for decision-alternatives: trial-and-error like search only after existing routines fail to perform satisfactory
4. Learning: (a) insight in past performance guides the agent in adapting his routines and (b) aspiration levels for goal attainment rise as the agent becomes better at his work through learning.

The idea of the model is that the ICA receive purchasing jobs and decide whether or not they are able, given their skills and capacity, or allowed, given the corporate policies, to execute the job. The agent’s skills increase by executing jobs; the marginal learning effect is decreasing. From feedback on the realized quality of a job, the ICA can determine whether the decision rule achieved the goal. If not, the ICA reconsiders this rule.

We introduce five so-called application areas in our model: office expenses, professional services, human resource expenses, automation and communication, and accommodation. We emphasize that these labels are merely illustrative. For defining the type of purchase situation we used Faris et al.’s (1967) buyclass framework which distinguishes three types of jobs: Straight Rebuy (SR), Modified Rebuy (MR) and New Task (NT) jobs. Specific content (c) and process skills (p) are required for performing a job. Thus, a job is classified according to job type and application area, and requires content and process skills. For sake of readability we mostly omit indices of job type and application area.

Within the department of an ICA purchasing jobs arise according to a Poisson process. Depending on whether such a job has already been performed in the past, the job can be classified as an SR, MR or NT job. The skills required for a job (c, p) depend on the job type (SR, MR, NT) and application area. A job requires different types of skills. We assume that only two skills are required: process skills and content skills. The model can easily be extended to include more skill types. The required skills are drawn from a triangular distribution as shown in Table 1 (see e.g. Law and Kelton, 2000) with mean μ_c (μ_p), minimum m_c (m_p) and maximum M_c (M_p). The agents do not know this distribution exactly.




Figure 1: Skill level distribution

3 THREE AGENT TYPES

3.1 Company Agent

The CA represents company management, sets target quality levels, and can use corporate policies to influence the purchasing process. There is a target quality level (\bar{Q}) for each combination of job type and application area. Our definition of quality is an abstract notion in which we also include cost aspects, i.e. the time and money spent on the job. The target level for a particular job type in an application area is based on the past performance of all agents on that job type (ICA and PA), where only the last n jobs are taken into account. So all agents report to the CA. This means that target levels increase when the ICA or PA improve their performance. Target levels are updated after each time period; we work with discrete time.

The company agent can influence the decision process of an ICA by enforcing corporate policies. A policy states that a job of a particular type and application area should always be executed in the same execution mode. All NT jobs should be done in cooperation, e.g., or all jobs related to Office Expenses should be done by the PA. An ICA has to comply with these policies. By setting corporate policies the CA might be able to influence the purchasing process and increase the quality achieved.

3.2 Internal Customer Agent

The internal customer agent represents a department within the company, but not the purchasing department. The decision process of the ICA closely follows Cyert and March (1963) and is shown in Figure 2.

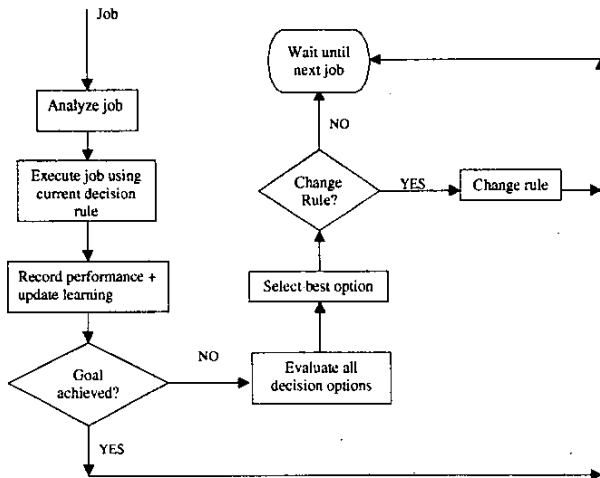


Figure 2: The decision process of an Internal Customer Agent

The ICA has to handle purchasing jobs that arise within the department. Jobs arrive according to a Poisson process, with an expected number of arrivals in an application area in one time interval of λ_{aa} . The ICA analyzes the job and determines whether it is a SR, MR or NT job and observes the application area. Given these attributes the ICA estimates the required skills (\hat{c}, \hat{p}) , drawing from a triangular distribution as shown in Figure 1, with parameters mean $\mu_{\hat{c}}, (\mu_{\hat{p}})$, minimum $m_{\hat{c}}, (m_{\hat{p}})$ and maximum $M_{\hat{c}}, (M_{\hat{p}})$. The ICA does not know the exact distribution of skills required. The ICA works with a distribution with a larger variance, but this variance can decrease because of learning effects. The ICA estimates both the required content and process skills. The goal of the ICA is to reach for all purchasing jobs the target quality level (\bar{Q}) that is set by the Company Agent.

The performance of the ICA is determined by a quality function that relates the available and required skills to the realized quality. This quality function is of the form:

$$f(c, p) = \beta_{0r} + \beta_{1r} \cdot c^2 + \beta_{2r} \cdot p^2, \quad \beta_{0r} > 0, \beta_{1r}, \beta_{2r} < 0, \quad (1)$$

where c denotes the required content skills and p the required process skills. The values of the β 's represent the available skills of the agent. The ICA knows the functional form of this quality function.

We use a quadratic form to express that the more difficult a job, the faster the quality decreases. We also implicitly included things like time spent on a job in our aggregate quality measure, and then the more difficult the job, the more time might be required. In Figure 3 we show an example of such a quality function in which $\beta_{0r}=5, \beta_{1r}= -0.1$ and $\beta_{2r}= -0.1$.

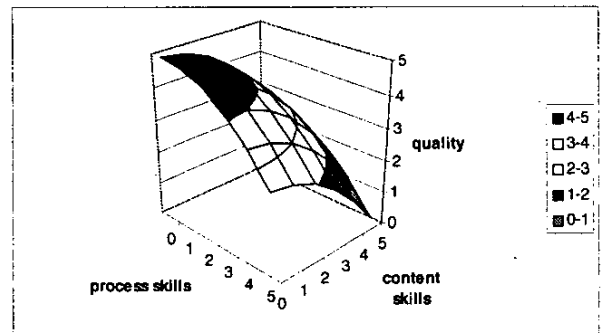


Figure 3: Example of a quality function; the quality given the required process and content skills

The ICA has specific skills available and the larger the gap between the required skills and the available skills, the lower the quality. Furthermore, the higher the required

skills, the more effort (time) is required to execute the job, and the lower the realized quality will be. Note that for difficult jobs the target quality level will also be lower because the target quality level is based on the past performance of all agents. The parameters, β_{0t} , β_{1t} , and β_{2t} , can be different for each agent and the index t indicates that they can change over time.

For a specific agent, β_{1t} can be different for each application area, representing that the ICA does not possess the same amount of content knowledge in all application areas. The parameter β_{2t} is related to the process skills, and is assumed to be the same for each application area. β_{0t} is the same for all application areas and represents the maximum achievable quality. We distinguish the following quality functions: $f_i^j(\cdot)$ is the quality function of ICA j at time t , $h_i(\cdot)$ is the quality function of the PA at time t , and $fh_i^j(\cdot)$ is the quality function in case ICA j cooperates with the PA at time t . The realized quality of job i , executed at time t by ICA j , is then equal to $Q_i = f_i^j(c_i, p_i)$.

In the decision process the ICA uses decision rules to determine the execution mode. Possible execution modes are to do the job self, to let the PA do the job, or to cooperate. A decision rule indicates, given a job type and application area, which execution mode should be chosen for a particular job. Stated more formal: (job type, application area) \Rightarrow action, with action \in {self, PA, cooperation}. Decision rules can change over time, but can also be fixed by a corporate policy, enforced by the CA. The ICA communicates with the PA whenever necessary, i.e. when the PA is involved in executing the job.

After the job has been executed the result of the action is evaluated. The fulfillment of job i results in feedback about the realized quality (Q_i), which is equal to $f_i(c, p)$, $h_i(c, p)$ or $fh_i(c, p)$, depending on execution mode. As long as the target quality level is being met ($Q_i > \tilde{Q}$), the decision rule apparently suffices, and nothing has to be changed. Otherwise, the rule has to be reconsidered.

Based on the expected skill level requirements (\hat{c}, \hat{p}) for the last job and given the knowledge of the PA's quality function, the ICA determines the expected quality of the possible actions: $f_i(\hat{c}, \hat{p})$, $h_i(\hat{c}, \hat{p})$, $fh_i(\hat{c}, \hat{p})$. The action that yields the highest expected quality is selected, and if necessary the decision rule is changed. Note that it is possible that none of the actions is expected to achieve the target quality level.

If the PA was involved in executing the job, the ICA receives feedback about the performance of the PA. The ICA thinks that $Q_i = h_i(\hat{c}, \hat{p})$, or in case of cooperation $Q_i = fh_i(\hat{c}, \hat{p})$, based on the skill requirements as estimated during job analysis. Q_i is an observation of the function h or fh . Based on several observations the ICA can estimate these quality functions by using a multiple linear regression method (see e.g. Freund and Walpole, 1987). To solve

the set of linear equations in the multiple linear regression the ICA uses Gauss-Jordan elimination, because it is a straightforward and understandable method (see e.g. Press et al. 1994). The ICA knows the functional form of the quality function and uses the last X observations to estimate the parameter values of this function. He does not use all observations because the performance of the PA can change over time (quality increases). The ICA can use the estimated functions \hat{h} and \hat{fh} in reconsidering and comparing the decision options. Given that the ICA does not know (c, p) , but only has estimates (\hat{c}, \hat{p}) , the ICA can only approximate the quality function of the PA.

Besides the learning process through the feedback from the PA, i.e. the ICA learns about the performance of the purchasing department, there are two learning aspects that are included as learning by doing. First, the ICA learns to estimate the required skills for a particular job; m_c converges to m_c and M_c to M_c , which reduces the uncertainty. The same holds for the process skills.

$$\begin{aligned} m_c &:= m_c + \delta m_c \quad \text{as long as } m_c < m_c \\ M_c &:= M_c - \delta M_c \quad \text{as long as } M_c > M_c \end{aligned} \quad (2)$$

The uncertainty is reduced after a particular number of jobs of one type and application area have been executed. The ICA stops learning when the real distribution is known (see Figure 1), which means that there will remain uncertainty about the required skills.

Second, the skills of the ICA, represented by the β 's in the quality function, increase by performing purchasing jobs. These learning effects are linked to the number of jobs executed for a job type in an application area. The marginal learning effect is decreasing over time, i.e. the more jobs have been done of a specific job type or application area, the smaller the additional knowledge, as is represented in equation (3).

$$\beta_{1t} := \beta_{1t} + (\delta\beta_1) / x \quad (3)$$

where x is the number of times this particular parameter has been increased. The skills are increased after a number of jobs of one type have been executed in one specific application area. The learning effect might be larger for NT jobs than for MR jobs. There are also possible differences between content and process skills, and differences in learning from different application areas.

3.3 Purchasing Agent

The purchasing agent represents the purchasing department, assuming that such a dedicated and specialized unit or officer exists in the organization. In this model the PA is passive and modeled in a simplified way. In further re-

search the PA could be modeled in a way similar to the ICA, where the PA tries to realize his own goals. Now the PA waits until he receives a request from an ICA to execute a purchasing job. We assume that the PA has infinite capacity. After execution the ICA is informed about the realized quality of that job, which is based on the quality function of the purchasing agent, $Q_i = h_i(c_i, p_i)$. This quality function has the same functional form as the quality function of the ICA, but with different parameter values (β 's). The outcome Q_i is feedback for the ICA, who can base his future decisions on the performance of the PA. Note that, although the ICA perceives the realized quality, he can only approximate the quality function of the PA, because he does not know the exact skill levels required for the job.

4 THE SIMULATION MODEL

Using the three agent types presented in the previous section, we constructed a simulation model, programmed in eM-Plant (Tecnomatix, 2000). In the simulation model there is 1 CA, 1 PA and there are several ICA. In this section we discuss two additional aspects of the simulation model: the experimental factors and the output of the model. Numerous experiments are possible with the multi-agent simulation model. Based on a theoretical and empirical study of the topic (De Boer and Pop-Sitar 2001), we have decided to concentrate on the following factors.

1. *Corporate strategies.* The CA can use different corporate policies. A corporate policy can be related to job type, application area, and also to a specific agent. If there are no fixed decision rules, we assume that the ICA has autonomy and decides what is best in a given situation. In practice, organizations often declare very strict formal policies stating that all purchasing should be handled by the purchasing department. However, the actual compliance to such policies is often quite low.
2. *Information.* What is the amount and quality of information the ICA possesses? One extreme assumption is that the ICA knows everything, i.e. he has complete and perfect information about what is required for a job and who is best equipped to do that job. To be more specific, the ICA knows the required skills c and p , and the quality functions h and fh . This is not the most realistic situation, normally there will be uncertainty in the required skills and in the knowledge about the performance of the PA. There is a whole region of scenarios with limited information and different levels of uncertainty.
3. *Time.* It can be interesting to see the development of the organization over time. By running the model for a different number of time periods we can investigate the short and long term effects. One interesting aspect is the involvement of the

PA over time. Another question is whether the model goes to a steady state in the long run.

4. *Learning.* Several experiments are possible with respect to the learning behavior and knowledge of the agents. The ICA can be more or less knowledgeable about purchasing and be fast or slow learners. The initial skills for all ICA might be the same ($\beta_{0,0}$, $\beta_{1,0}$, $\beta_{2,0}$), but there may also be differences between them.
5. *Job mix.* The job mix is yet another possible variable in the experiments. The number of application areas can be varied, just as the frequencies of the number of jobs of one type (SR, MR, NT) in an application area. There is also a choice whether all ICAs receive the same job mix, or that the ICA only receive jobs in one application area (generalization versus specialization).

One of the features of a simulation model is that we can run the model over a period of time. The simulation model is run for a specified number of time periods. Each time period the ICA receives a mix of purchasing jobs and decides who should execute a particular purchasing job. The output of the simulation model is a list of all jobs with their attributes (application area, job type), the skill requirements and expected quality as estimated by the ICA, the target quality level, the execution mode, and the realized quality. Furthermore, a table with aggregate output data is presented, in which the average performance of each agent for each job type and application area is indicated. Also the involvement of the purchasing department is computed, i.e. the percentage of jobs that is executed by the purchasing department, as well as the percentage of jobs that was done in cooperation.

5 SIMULATION EXPERIMENTS AND RESULTS

In this section we highlight interesting results for some of the experimental factors. We cannot present all details here, but we give some general results. We start with a basic scenario, see Table 1, which is the starting point for the other experiments. In each of the following experiments, one of the parameters in the basic scenario is varied.

Table 1: Basic scenario

Experimental area	Basic value
Corporate strategy	ICA free to decide
Information	$m = \mu - 10\%$, $M = \mu + 10\%$ $\hat{m} = \mu - 20\%$, $\hat{M} = \mu + 20\%$
Time	40
Job type frequency	(SR,MR,NT)=(1/3,1/3,1/3)
Number of jobs per application area	Several jobs in 1 area and a few in remaining areas

There are five ICA and there is one PA in our model. In the basic scenario all job types have equal probability of occurrence (1/3). The required skill levels are drawn from a triangular distribution as shown in Figure 1 with $m = \mu - 10\%$ and $M = \mu + 10\%$, where μ is the average required skill level on a scale of 1 to 5, depending on job type and application area. We assume that all ICA have infinite capacity, autonomous control, identical initial quality functions and similar learning processes. Their expectations about skill requirements are $\hat{m} = \mu - 20\%$ and $\hat{M} = \mu + 20\%$. In each time period the ICA receive several jobs in one application area and a few jobs in the remaining four areas, such that each ICA specializes in one area. The PA has infinite capacity and has less content skills, but more process skills than the ICA. With respect to content skills, a lot is learned from NT jobs, little from MR jobs and nothing from SR jobs. The same holds for the process skills.

We assume that the PA possesses more process skills than the ICA, i.e. for the PA $\beta_{0,0} = 6$, $\beta_{1,0} = -0.15$ and $\beta_{2,0} = -0.1$, and for the ICA $\beta_{0,0} = 6$, $\beta_{1,0} = -0.1$ and $\beta_{2,0} = -0.15$. Furthermore, the maximum achievable process skills of the PA after learning is also higher than the level that can be obtained by the ICA. This does not mean that the PA always has higher process skills than an ICA because when the PA is not involved in the purchasing process the ICA learn and can surpass the PA with respect to process skills. For the content skills the opposite holds. We think that this is a reasonable assumption, because every department has its own area of expertise. Given this assumption, it is clear that the best results can be achieved by always cooperating. In this case they both learn and use each other's experience. However, the ICA does not know that this is the best option. The question is whether in the model the ICA learn to cooperate with the PA on difficult jobs. Furthermore, the results for different corporate strategies can be compared with the best case, i.e. all jobs by cooperation.

Experiments were done within all areas described in Table 1. When discussing the experiments we will focus on general results. Rather than looking into too much detail, we are more interested in the differences between scenarios and developments over time.

In the basic scenario we see that the involvement of the purchasing department depends on the relation between the required process and content skills. When the required process skills are low, the ICA can execute the job himself, there is no reason to cooperate. Also when the required content skills are higher than the required process skills, the content skills are more important for the final outcome than the process skills and the ICA has a lower incentive to involve the PA. Only when the required process skills are high the PA is involved. In Figure 4 we see the relation between the required skill levels and the involvement of the PA. The different areas in Figure 4 indicate which agent executes the majority of the purchasing jobs for the given skill requirements. The ICA performs most jobs that re-

quire low process skills, but it does not mean that the PA is never involved in the execution of these jobs.




Figure 4: Relation between required skill levels and who executes the majority of jobs

Based on the experiments we conclude that corporate strategies can have a serious influence on the results. If agents cannot cooperate, the results are worse than when agents are allowed to cooperate. In Figure 5 the lowest three lines represent strategies which do not include cooperation. Cooperation speeds up the learning process; both agents learn at the same time. As mentioned before, a situation in which all jobs are done in cooperation leads to the best result. But cooperation is not always required to reach the goals, therefore the ICA do a lot of the jobs themselves when they have autonomy. When no cooperation is allowed, then jobs which require high process skills are mostly outsourced to the PA, the ICA does not have the required process skills. Jobs which require higher content skills than process skills are done by the ICA. So the ICA learns that he possesses more content skills and less process skills than the PA.




Figure 5: Performance over time for the different corporate strategies

When there is a lot of uncertainty with respect to the required skills, the decision rules change more often. In case of perfect information the ICA will find out, after a

learning period, that cooperation is the best alternative and after some time most jobs will be done in cooperation (cf. Figure 5). However, changes in the involvement of the PA due to uncertainty do not necessarily lead to a worse final outcome. When the ICA performs more jobs himself, he will learn about the process, increase his process skills and in this way decrease the negative effects of – at the time – non-optimal decisions. Thus, there is a trade-off between a good decision now, without a serious learning effect, or a worse decision that, due to the learning effects, leads to improved solutions in the future.

Over time the ICA learn more about the environment, the PA, and increase their skills as a result of learning by doing. In the first time periods the cooperation between the ICA and PA increases, because the ICA notices that the PA has more process skills. After a while, the ICA has learned more about the process and has increased his own process skills. From that moment on, the necessity to cooperate with the PA fades. After a while a steady state is reached, where only the jobs which require high process skills are done by the PA or by cooperation (see Figure 6).




Figure 6: Involvement of the PA over time

Because of the learning effects the overall quality increases over time (see Figure 5). The more the ICA learns about the purchasing process, the lower the need to involve the PA. Therefore, the lines that outline the border between the decision options in Figure 4 will move slowly to the right to indicate that the PA is only involved in jobs that require high process skills.

6 DISCUSSION

Let us now look at the questions we set out to answer. Obviously, the degree of Purchasing involvement is first of all a direct result of the reigning corporate policy. As far as policy allows Purchasing to become involved, the degree of involvement is logically related to the required process and content skills: if a PA possesses complementary skills and the ICA is aware of this, the former will be asked to help out. The awareness is however not given and the ICA must learn what the PA can do and cannot do.

The number of jobs per application area seems to determine the degree of involvement through the way target quality levels are defined. If target levels are linked throughout the firm, one or few departments may cause these levels to rise strongly, thereby 'forcing' other departments to do better and bring Purchasing in. The involvement therefore seems to hinge on the degree of fragmentation of jobs per application area and the mechanism for setting target levels. However, this needs to be confirmed in more elaborate experiments.

In the hypothetical case of perfect information, the involvement of the PA converges to 100%, while this percentage drops with increasing uncertainty. This underlines the importance of creating sufficient mutual awareness of each other's added value.

There is clearly a systematic pattern of Purchasing involvement over time. After a brief 'training period', the ICA learn to do things themselves and the PA's added value shrinks. Given the initial distribution of skills between the ICA and the PA, the overall superiority of cooperation in terms of goal attainment is hardly surprising. Still, it is clear that a more 'liberal' corporate policy leads to higher levels of goal attainment than the more strict 'either/or' or 'Purchasing only' policies. Even the 'ideal' mode of cooperation is not always necessary as long as the target level for the goal is achieved. This explains why 'perfect information' not immediately leads to switching to cooperation.

In terms of immediate recommendations, it appears that it is not a good idea to impose strict corporate policies because they might prevent good results. Purchasing should not be done in isolation, but the knowledge of other departments should be exploited as well. It is therefore important to increase the awareness of the departments of the skills of the purchasing department, but also the other way around. Also the quality of information is very important, it should be clear what skills are required for a job and who possesses these skills.

7 CONCLUSIONS AND FURTHER RESEARCH

From the simulation experiments it turns out that we can model a realistic organizational setting by using a multi-agent simulation model. We see that over time the agents learn that cooperating is useful, although they do not know at the beginning that it is the best option. They find ways to achieve their goals by trial and error. The knowledge of the agents increases over time and leads to better decisions and increased quality. In addition, the model evolves into a steady state.

Obviously, the simulation model is a simplified representation of reality. We did not explicitly include costs, which are very important in practice. Cooperation might be the best option from a knowledge perspective, it might not be from a cost perspective. To make the model more realis-

tic, we will try to include secondary goals, such as cost, which will lead to more complex decision rules for the ICA. In addition, the PA will be upgraded to the level of sophistication of the ICA in the sense of learning ability, having the freedom to take the initiative and/or to refuse jobs and working with limited capacity. The PA can also have a completely different goal than the ICA.

So far the choice for experimental factors was based on a theoretical and empirical study. There are a lot more factors that can be varied in the simulation experiments. We want to use an experimental design (see e.g. Law and Kelton, 2000) to select the factors that have a significant effect on the involvement of the purchasing agent. Furthermore, when interaction effects are found it can lead to a better understanding of the behavior observed in the model.

As to using Multi-Agent Simulation in Purchasing and Supply Management research: the potential seems huge but requires many choices about assumptions in order to get started. However, this is where the existing body of literature in Purchasing and Supply provides useful guidance. A continuous challenge is to not give in too easily to the idea that the results will be obvious, because it is exactly by resisting these intuitive stimuli that new insights will emerge when the system becomes to complex and dynamic for us to oversee.

REFERENCES

- Cyert, R.M. and J.G. March. 1963. *A behavioral theory of the firm*. Prentice Hall, Englewood Cliffs.
- De Boer, L. and C.E. Pop-Sitar. 2001. Managing purchasing of non-product related (NPR) goods and services. On horizontal frictions and vertical ignorance. In *Proceedings 10th IPSERA Conference*, Jonkoping, Sweden, pp. 131-140.
- Faris, C.W., P.J. Robinson and Y. Wind. 1967. *Industrial buying and creative marketing*. Boston: Allyn and Bacon.
- Fearon, H.E. and A.W. Bales. 1995. Purchasing of nontraditional goods and services. CAPS Report, NAPM.
- Fearon, H.E. and A.W. Bales. 1993. CEOs'/Presidents' perception and expectations of the purchasing function. CAPS Report, NAPM.
- Franklin, S. and A. Graesser. 1996. Is it an Agent, or just a Program?: A Taxonomy for Autonomous Agents. In *Proceedings of the Third International Workshop on Agent Theories, Architectures and Languages*. Springer-Verlag.
- Freund, J.E. and E. Walpole. 1987. *Mathematical Statistics*. Prentice Hall, New Jersey.
- Law, A.M. and W.D. Kelton. 2000. *Simulation modeling and analysis*. McGraw-Hill, Singapore.
- Press, W.H., B.P. Flannery, S.A. Teukolsky and W.T. Vetterling. 1994. *Numerical recipes in Pascal*. Cambridge University Press, Cambridge.
- Tecnomatix. 2000. *eM-Plant Reference Manual version 5.5*. Tecnomatix Technologies Ltd.
- U.S. Bureau of the Census. 1991. *Annual Survey of Manufactures*(Washington, D.C.). Statistics for Industry Groups and Industries.
- Wooldridge, M. and N.R. Jennings. 1995. Intelligent Agents: Theory and Practice. *Knowledge Engineering Review* 10(2), 115-152.

AUTHOR BIOGRAPHIES

MARK EBBEN is a post doctoral researcher at the Operations Management and Logistics department at the University of Twente. He received his Ph.D. from the University of Twente in 2001. His research interests include simulation, scheduling, logistics and transportation. His email address is <m.j.r.ebben@sms.utwente.nl>.

LUITZEN DE BOER is an assistant-professor at the Operations Management and Logistics department at the University of Twente. He received his Ph.D. from the University of Twente in 1998. His research interests include purchasing management, decision-making, and organization science. His email address is <L.deboer@sms.utwente.nl>

CORINA POP SITAR is a doctoral candidate at the Operations Management and Logistics department at the University of Twente. Her research is sponsored by the Dutch Association for Purchasing Management (NEVI) and deals with managing purchasing processes outside the Purchasing department. Her email address is <c.e.pop-sitar@sms.utwente.nl>.