Learning in the Artificial Factory

Martin Natter, Markus Feurstein, Andreas Mild, Alfred Taudes, Michael Trcka
Vienna University of Economics & Business Administration
Pappenheimgasse 35, A-1200 Vienna, Austria
martin.natter@wu-wien.ac.at

Georg Dorffner, Christian Merz
Austrian Research Institute for Artificial Intelligence
georg@ai.univie.ac.at

Abstract

We study the effects of various incentive schemes on the learning behavior of teams in an artificial factory. Modeling the new product development process, we demonstrate, how production and marketing agents learn to coordinate their actions in order to produce the optimal product with respect to their incentive schemes. As a coordinating mechanism between marketing and production, we use the House of Quality framework of Hauser and Clausing [6]. The House of Quality methodology, which is used by real firms, contains important information from marketing and production. It is a procedure that facilitates the search for new promising (from market perspective) and feasible products (from a production/design perspective). We found that the House of Quality approach yields higher life cycle returns than the traditional search for new products - especially for a low number of search steps. This is an important finding recommending the application of the House of Quality since the number of search steps directly influences time to market. Thus, minimizing the number of steps could be an important competitive advantage in todays fast moving consumer markets.

1. Introduction

Traditionally, microeconomics has assumed that the form of the organizational network is hierarchical, where the result of a computation is only passed to several immediate subordinates or to one agent on the next higher level, respectively. It is assumed that management has a general model of the environment and the organization and on this basis derives the decision rules for the agents. However, economics agents have a limited capacity for computation and their knowledge is limited to their field of specialization.

Therefore, if the rationality of economic agents is bounded, the capacity of management to build models and derive sub-models must be bounded too, as models are built on communicated or observed data by estimating parameters and performing symbolic computations. Thus, the assumption that - also as a team - a management agent can make a total model of the firm and its environment looses its credibility in more complex situations. Furthermore, the environment must also be rather stable, so that the need for "reorganizing", i.e. determining new decision rules, does not arise too often so that the bounded rationality of management is exceeded (see e.g. [10]).

Both requirements are often violated today. Consider the following statement from Clark and Fujimoto [3] about the complexity of consumer behavior in today's car markets: "Car buyers don't choose between brands on prices, qualities, and functions alone anymore. In order for a brand to become appealing to them, certain "soft" variables such as "urban feeling" or "high-tech feeling" have to be added." (cf. [2]). Obviously, many variables enter into the functional relationship between the purchase probability and the technical specifications of a car in a nonlinear way in such a situation, where a number of variables have equivocal meanings and are not readily encodable. It thus seems highly unlikely that someone who is not constantly and directly exposed to customers can make a meaningful model.

When important knowledge is created through daily operations and "embedded" in the employees, one must think about how this tacit knowledge can be integrated into the *organizational knowledge base* so that via *organizational learning* good organizational decisions arise.

A number of empirical works indicate the importance of the knowledge integration view. Clark and Fujimoto [3], for instance, show that in the 80s Japanese car manufacturers that used multi-functional teams coordinated by a

0-7695-0493-0/00 $10.00 (c) 2000 IEEE

high-profile project manager outperformed the more loosely coupled, sequentially organized European and American competitors both in terms of development time and product quality [2]. A study by [11], which surveyed 788 new product developments in Japan, confirms this finding. Similarly, [5] find that the success of a new product increases with the intensity of communication between marketing and R & D. However, when R & D and marketing personnel become too close friends, the need for harmony prohibits an open discussion of conflicting arguments and consequently new product success decreases [5]. Lawrence and Lorsch compare different companies in different industries to formulate the ”contingency theory of organization”: the more complex the environment, the more differentiated the knowledge has to be and the stronger the need for high-band-width communication and non-hierarchical coordination [9].

The aim of this paper is to develop quantitative models for organizational learning in tactical planning. We focus on modeling the new product development process and demonstrate, how production and marketing agents learn to coordinate their actions in order to produce optimal products with respect to their incentive schemes. We will also show that methods of Total Quality Management (TQM) like the House of Quality [6] are coordinated search procedures for organizational learning and study under which conditions they work best (functional and institutional integration).

2. World Definition

In this section, we define the environment in which the agents live, i.e., their world. The environment consists of three major components:

- The market definition, which describes life-cycle effects and market share as a function of product attributes, price and promotion budgets.

- The production definition, which describes how production processes map into technical product features; i.e., the production and cost function of the firm.

- The interface definition, which describes how real technical product features are related to attributes as perceived by the consumers.

Figure 1 shows the interactions in this environment. Starting without any prior knowledge, the agents observe realizations of past actions and try to build their own model of the world. Observations of actions are realized when prototypes of a product are developed and used for a market study. Once a prototype is built, the real costs, technical characteristics and consumer preferences are known. These examples can then be used to learn their model of the world.

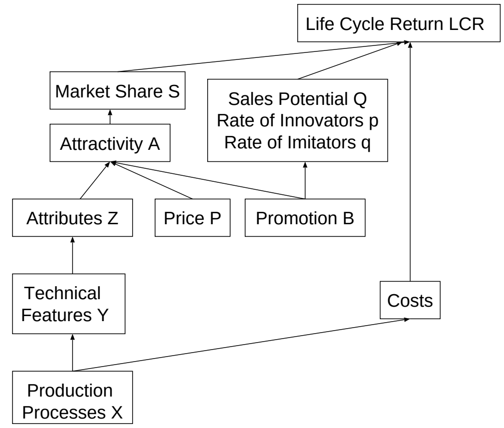

Figure 1. Structure of the environment

The learned relations can be used in the next step to decide which prototype to build next, etc.

Market Definition

The Life Cycle Return (LCR) of a product is given by the sum of profits over all periods, t:

$$LCR = \sum_{t=1}^{T} \pi_{(t)} \tag{1}$$

The profit of a period is determined by the following relation:

$$\pi_i(t) = [P_i(t) - \phi_i(Z_i(t))]S_i Q(t) - B(t) \tag{2}$$

where $\pi_i(t)$ represents the profit of firm i at time t, P the price, Z the attribute vector of the product, Q the market volume, and S the firms' market share. $\phi_i(Z)$ denotes the products' costs as a function of the product attributes.

Life Cycle (LC) effects are modeled by an extension of the classical Bass model [1] which finds strong empirical support [12]. With only three parameters (rate of innovators (p), rate of imitators (q), market potential (Q)) the sales quantity of each period is determined:

$$Q(t) = Q\left(\frac{p(t)(p(t)+q)^2 e^{-(p(t)+q)t}}{(p(t)+qe^{-(p(t)+q)t})^2}\right) \tag{3}$$

The coefficient of innovation, p, is modeled as function of the marketing budget, B, in period t.

$$p(t) = a + c\,ln[B(t)+1] \tag{4}$$

The market share of a product, S_i, is given by

$$S_i = \frac{A_i}{\sum_{j=1}^{J} A_j} \tag{5}$$

0-7695-0493-0/00 $10.00 (c) 2000 IEEE

where A denotes the attraction of the product with attribute vector z and price P.

$$A = f(z)g(P) \tag{6}$$

where f(z) is a function of the product position relative to the ideal point, z^*. Following Shocker and Srinivasan (1974), we model the distance of the product offering to the ideal point as a weighted Euclidean distance:

$$f_j(z_i) = 1 - \frac{[(z^* - z_i)'W(z^* - z_i)}{(z^{*\prime}Wz^*)} \tag{7}$$

with W representing a diagonal matrix whose diagonal elements w_i denote the weights consumers place on attribute i.

g(P) is a downward sloping function of price:

$$g_j(P_j) = (1 - bP_j/\rho) \tag{8}$$

with $P_j \leq \rho/b$ and ρ/b is the market's reservation price such that $0 \leq g_j(P_j) \leq 1$.

Production Environment

The relationship $X \to Y$ is captured by the following relationship (Cobb-Douglas)

$$Y_i = \prod_j X_j^{\beta_{i,j}} \tag{9}$$

with $0 \geq \beta_{i,j} \geq 1$.

The costs of a technical feature depend on a weighted sum of the chosen processes and materials X through which the feature Y is characterized and a penalty for deviations from the intended attributes.

$$\phi_i(Y_i) = \sum_{i,j=1}^{I,J} r_{ij}X_iX_j + d_i(Y_i - \hat{Y}_i) \tag{10}$$

Marketing-Production Interactions

Attributes Z_i of a product are a weighted function of technical features Y of that attribute:

$$Z_i = \sum_{j=1} \theta_{i,j}Y_j \tag{11}$$

with $\sum \theta_{i,j} = 1$

Costs of an attribute Z are given by

$$\phi_j(Z_j) = \sum_{i=1}^{I} \phi_i(Y_i) \tag{12}$$

Product attributes as perceived by the consumers, z, are a function of the real product attribute values, Z, last periods perceptions and advertising budget, B, which can be used to shift the perceptions from Z to z:

$$z(t) = \alpha_1 Z(t) + \alpha_2 z(t-1) + \alpha_3 ln[B(t)+1](z^* - z) \tag{13}$$

3. Learning and Decision Making

The marketing agent learns the expected market share of a product as a function of past realizations of product attributes, market shares, etc. The production agent learns the relationship between input factors X and technical features of the product Y on one hand and which input factors to use for a given set of target features Y, on the other hand, i.e., the inverse function. As described by equation 10, the production agent learns to manufacture a product in a way such that the costs and the distance to the target product are minimal. Agents' knowledge is modeled via artificial neural networks with one hidden layer using five hidden units (determined by trial and error). After having learned market reactions, the way products can be built and their production costs, the agents have to decide together which technical features and – as a result – which attributes their product should have. In contrast to individual learning, we define this process as organizational learning [7] or outer loop learning.

By using team decision methods or negotiations the agents try to develop an optimal product according to their incentives.

The process of coordination could be cooperation or negotiation. In the case of cooperation the two agents discuss different possible sets of technical features. For each set the production agent calculates the production cost and the marketing agent calculates the market share based on their approximations of the world. Together they are able to calculate the expected revenues and costs of producing the product - the Life Cycle Return. In a team they value different sets of technical features and choose the one that promises the highest Life Cycle Return. This method of coordination is possible if the agents can be motivated to find a solution that is optimal for the company (LCR) and if they have no different individual incentives. One way to reach such a status is to link the wage of the agents to the profit of the company.

In many cases such an incentive scheme is impossible to reach. Usual benchmarks for marketing agents are sales or market shares. Production agents are often rewarded for low costs. In our simulation the marketing agent tries to maximize the market share and the production agent minimizes production costs. In this setting the agents evaluate different peculiarities of the product features and accept a change only if it does not decrease their individual payoffs. So if they start with product features Y^{old} they will accept a change in design to features Y^{new} only if the new design increases market share and decreases costs. Such win-win situations are of high relevance in management practice when changes are implemented across several functional units.

For both ways of negotiation (cooperation, win-win negotiation) we modeled a situation where the agents search for new products randomly and one where they use the House of Quality to guide their search.

Search Models

In the process of coordination the agents have to find a target set of technical product features Y_t that maximizes expected Life Cycle Return (LCR). To model this search, we used Simulated Annealing, an optimization method that was first used by [8]. In our implementation the agents choose one product feature Y_i and change it according to the rule $Y_i{}' = (1-\xi) * Y_i + \xi * \beta$ where ξ is a parameter that allows to change the step width in each search (in our model, ξ is reduced during the search process) and β is uniformly distributed in the range $\beta \in [-1, 1]$.

The agents accept the change if the expected reward R' with features Y' is higher than the expected reward R for the original product. If the new reward is lower than the original one the change is accepted with a probability of

$$\frac{1}{1 + e^{\frac{R-R'}{\mathrm{Temp}}}} \tag{14}$$

where $Temp$ is a parameter which avoids local minima. $Temp$ is reduced during the search process - so that in the beginning of the search worse solutions are accepted and at the end only improvements are allowed.

For the two ways of negotiation, different measures of the return of a product consisting of a set of features are used.

- For cooperation the agents try to maximize LCR, i.e., $R = \mathrm{LCR}$.
- In the case of the win-win negotiation each agent has an individual reward. R_M of the marketing agent is the expected market share S and R_P are the production costs ϕ and estimated deviations from the target features.

House of Quality

The House of Quality aims at finding a favorable product/process specification. It is a ”kind of conceptual map that provides the means for inter-functional planning and communication” [6].

As its name indicates, these interfunctional relationships are graphically depicted in a house. Its body is a matrix that contains the size and 'strength' of interrelations between technical specifications (features) of a product plan and customer attributes of the product concept. The entries of the matrix indicate in what way (direction, strength) a change in Y affects Z. The original approach consists of 4 houses, linking product concept with product plan, product plan with parts design, parts design with process design and process design with quality control measures. The entries are made based on tacit knowledge enriched with explicit knowledge and experimental data. The roof of the house, contains correlations between the technical features Y.

Hauser and Clausing describe how one search step is conducted in practice, using the specifications of a car's door as example: ”Our doors are much more difficult to close from the outside than those of competitors' cars. We decide to look further because our marketing data says this customer attribute is important. From the central matrix, the body of the house, we identify the engineering characteristics that affect this customer attribute: energy to close door, peak closing force, and door seal resistance. Our engineers judge the energy to close the door and the peak closing force as good candidates for improvement together because they are strongly positively related to the customer's desire to close the door easily. They determine to consider all the engineering ramifications of door closing. Next, in the roof of the house, we identify other engineering characteristics that might be affected by changing the door closing energy. Door opening energy and peak closing force are positively related, but other engineering characteristics are bound to be changed in the process and are negatively related. It is not an easy decision. But with objective measures of competitors' doors, customer perceptions, and considering information on costs and technical difficulty, we decide that the benefits outweigh the costs. A new door closing target is set for our door - 7.5 foot-pounds of energy. This target, noted on the very bottom of the house directly below the relevant engineering characteristic, establishes the goal to have the door easiest to close.”

For the analyses of our problem, we used only the first House of Quality, where the marketing and the production agent meet. In the House of Quality, we represent the connection between different technical features Y_i (some features promote other features, some features restrict each other) - the roof matrix - and the connection between technical features Y_i and product attributes Z_j - the central matrix - using the correlations $c^r_{i,j} = \mathrm{Corr}(Y_i, Y_j)$ and $c^c_{i,j} = \mathrm{Corr}(Y_i, Z_j)$ calculated for 100 training samples. We also estimate the importance of the product attributes in sales I_i using the same samples by learning the relation

$$f(Z) = \sum_i I_i * Z_i + \epsilon \tag{15}$$

in equation (6) and the costs k_j of technical features by learning

$$\phi(Y) = \sum_i k_i * Y_i + \epsilon \tag{16}$$

in equation (10).

To use the House of Quality in the search process, we calculate a rating $W(Y_i)$ of each technical feature Y_i by

4

Table 1. Comparison between House of Quality and random search using the same incentive schemes

steps	mean HoQ	std HoQ	mean without HoQ	std without HoQ
10	0.157	0.038	0.144	0.030
50	0.287	0.059	0.265	0.075
100	0.384	0.024	0.359	0.047
250	0.414	0.018	0.410	0.018
500	0.459	0.019	0.446	0.024
1000	0.483	0.014	0.483	0.017

$$W(Y_i) = \sum_j I_j * c^c_{i,j} - k_i \tag{17}$$

This value represents the "isolated contribution" of (Y_i) to LCR. To represent inter-feature dependencies (changing one feature may result in the (unwanted) change of another one) a modified feature value (W_m) is calculated.

$$W_m(Y_i) = W(Y_i) + \gamma \sum_{i \neq j} c^r_{i,j} * W(Y_j) \tag{18}$$

If the agents use the House of Quality they use $W_m(Y_i)$ to decide which feature value should be changed next. Attributes with higher W_m are more likely to be changed.

4. Results

Tables 1 and 2 show average values and standard deviations of expected life cycle returns over 50 replications of our simulation. While Table 1 represents the results for the coordinated search, Table 2 reflects the outcomes of the negotiation based (win-win) situation. Each table contains the results for the House of Quality search and the random search for different numbers of product evaluation iterations.

Comparing the results of the simulation searching for new products indicates that the House of Quality approach always yields higher LCRs. However, the results from a random search approximate to the results of the House of Quality approach when the number of steps is high (see Figures 2 and 3).

Allowing for 10 search steps only, the estimated LCR using the House of Quality is 9.2% higher as compared to the random search for coordinated incentives and 7% for individual incentive schemes, respectively. After 1000 steps the results are almost identical. Furthermore, it can be noted that the standard deviation of the outcome decreases considerably when the number of search steps is higher than 100. The number of search steps should be regarded as the time available for discussing about new product possibilities to be introduced.

The importance of choosing the right incentive functions can be seen. Linking the agents (additional) payoff to the overall performance measure of the firm leads to a higher LCR in general (see Figure 4). Furthermore, product development is "guided" into the right direction, i.e., the agents will improve the life cycle return of a new product with increasing time to discuss several product opportunities.

Linking the individual payoffs to local performance measurement (such as production cost or market share) allows both agents (especially if there is enough time to search) to improve their individual payoff while lowering the firms profits.

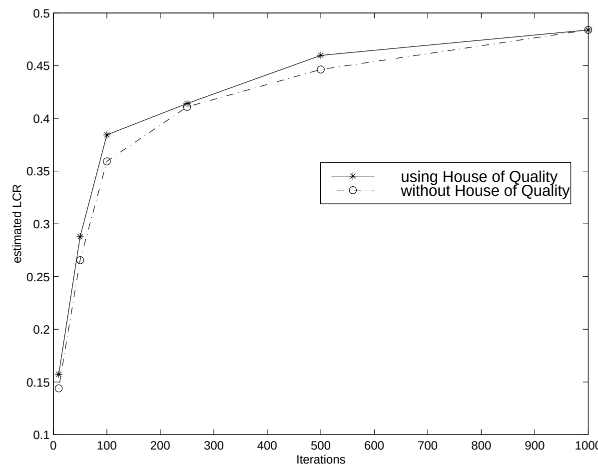

Figure 2. Results for the case of using the same incentive schemes

5. Conclusion

We have studied the product development process in an artificial firm using two different incentive schemes and two methods for product selection. We showed that coordinating incentive schemes - such as shares or options as part of the income - increase the performance (shorter time to market or higher product performance) of the firm [4]. As a coordinating mechanism between marketing and production, we used the House of Quality framework of Hauser and Clausing and compared it to a directed random search strategy. We found that the House of Quality approach yields higher life cycle returns than the traditional search for new

Table 2. Comparison between House of Quality and random search using conflicting incentive schemes

steps	mean HoQ	std HoQ	mean without HoQ	std without HoQ
10	0.137	0.029	0.128	0.031
50	0.253	0.049	0.213	0.049
100	0.343	0.041	0.283	0.054
250	0.371	0.006	0.367	0.007
500	0.367	0.004	0.363	0.002
1000	0.364	0.001	0.362	0.001

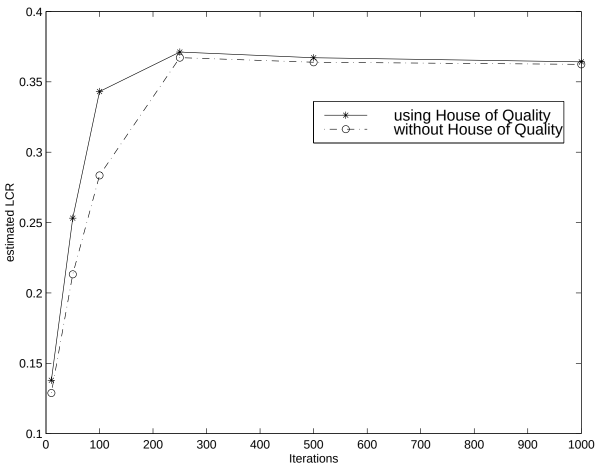

Figure 3. Results using individual incentive schemes

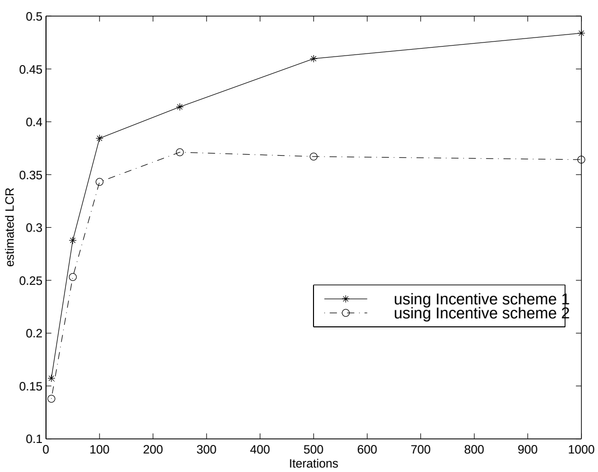

Figure 4. Comparing different incentive schemes

products - especially for a low number of search steps. This is an important finding recommending the application of the House of Quality since the number of search steps directly influences time to market. Thus, minimizing the number of steps could be an important competitive advantage in todays fast moving consumer markets.

References

[1] F. Bass. A new product growth model for consumer durables. *Management Science*, 15(4):216–227, 1969.

[2] H. Clark and C. Marshall. Definite reference and mutual knowledge. *Joshi A.K./B.L. Webber/I.A. Sag (eds.): Elements of Discourse Understanding, Cambridge University Press*, 1981.

[3] K. Clark and T. Fujimoto. The power of product integrity. *Harvard Business Review*, 68(6):107–118, 1990.

[4] M. Cohen, J. Eliashberg, and T. Ho. New product development: The performance and time to market tradeoff. *Management Science*, 42(2):173–186, 1996.

[5] A. D., R. Dahlstrohm, and S. Skinner. An exploratory investigation of organizational antecedents to new product success. *Journal of Marketing Research*, 34:107–116, 1997.

[6] H. J.R. and D. Clausing. The house of quality. *Harvard Business Review*, pages 63–73, 1988.

[7] D. H. Kim. The link between individual and organizational learning. *Sloan Management Review*, pages 37–50, 1993.

[8] S. Kirkpatrick, C. Gelatt, and M. Vecchi. Optimization by simulated annealing. *IBM Research Report RC*, 1982.

[9] P. Lawrence and J. Lorsch. Organization and environment - managing differentiation and integration. *Harvard Business School Press*, 1986.

[10] H. Mintzberg. Die mintzberg-struktur : Organisationen effektiver gestalten. *Moderne Industrie: Landsberg am Lech*, 1992.

[11] X. Song and M. Parry. The determinants of japanese new product success. *Journal of Marketing Research*, 34:64–76, 1997.

[12] F. Sultan, J. Farley, and D. Lehmann. A meta-analysis of applications of diffusion models. *Journal of Marketing Research*, 27(1):70–77, 1990.

Acknowledgement: This piece of research was supported by the Austrian Science Foundation (FWF) under grant SFB#010 ('Adaptive Information Systems and Modeling in Economics and Management Science').