



Adaptive Learning in Organizations: A System Dynamics-Based Exploration

Alessandro Lomi
Erik R. Larsen
University of Bologna
Ari Ginsberg
New York University

This paper employs a system dynamics-based framework to examine the limitations of experiential learning as a guide for decision-making in organizations. This framework departs from the more traditional approach to modelling experiential learning processes in organizations by emphasizing the systematic interaction between decision-making agents and their environments, rather than the effects of varying degrees of noise on performance. We present the results of a series of computer simulations that examined the consequences of adaptive learning in organizations by concentrating explicitly on the link between individual decisions and the system-level consequences generated by the interaction of individual choices. The results show that experience is a poor basis for learning primarily because the understanding of structural relations between individual actions and their aggregate consequences is confounded by nonlinear dynamics, time delays, and misperception of feedback.

Behavioral theories of the firm model organizations as target-oriented, routine-based systems that adapt incrementally to past experience (Cyert & March, 1963; March & Simon, 1958). In the tradition of behavioral research, learning is viewed as a continuous process of adjustment to changes in contingencies of reinforcement (Rachlin, 1991). As experience about individual contingencies accumulates, routines are developed that allow faster and faster response to recurrent stimuli, and the learning strategies encoded in these routines become the foundation for future rule-governed behavior. It is in this sense that routines are the basis for both organizational reliability and performance, and organizations are likely to behave in the future according to routines used in the past (Nelson & Winter, 1982; Hannan & Freeman, 1989).

Direct all correspondence to: Ari Ginsberg, Stern School of Business, New York University, 44 West 4th Street, New York, NY 10012.

In the behavioral decision theory literature, inference errors and judgment biases have traditionally been associated with uncertainty due to the probabilistic nature of the tasks (Kahneman, Slovic & Tversky, 1982). Following this experimental tradition, organizational decision theorists have typically modelled ambiguity as the aggregate result of random factors (noise) interfering with accurate assessment of an underlying "true" performance, or state of the environment (signal) (Lant & Mezas, 1990). The effectiveness of organizational learning, search routines, and the level of performance are then shown to vary as a function of the signal to noise ratio (Lant & Mezas, 1990, 1992). Specifically, a central focus of this work has been on outlining how conditions of ambiguity in the relationships among structure, action, and performance will affect the behavior of organizations modeled as experimental learning systems. Hence, behavioral models of incremental organizational learning are motivated on the ground that "if the world and preferences are stable, and experience prolonged enough, behavior will approach the behavior that would be chosen rationally on the basis of perfect knowledge." (March, 1989, p. 273).

In this paper we show that this coincidence of learning and planning, i.e., of heuristic problem solving and anticipatory action, will not necessarily be realized even in the absence of any random disturbance confounding the agents' perception of the environment. Indeed, the value of experiential learning, we argue, is problematic even in environments that are completely deterministic and almost completely stable. We claim this to be the case because organizational environments do not exist independently from the agents, but rather emerge from the network of interactions among them and from their perception of feedback.

To illustrate this point, we report the results of a series of computer simulations of processes of adaptive learning within a system dynamics framework, a modelling approach that is particularly helpful in examining how interactions between decision-making agents and their environments occur over time (Morecroft & Sterman, 1994). Our findings suggest that even in relatively simple nonstochastic environments, the interaction of multiple information feedbacks produces nonlinearities that result in complex—and, to the limit—chaotic, dynamic behavior. These results illustrate how the same decision-making processes that yield adaptive learning—behavioral re-enforcement based on past success—can also produce fundamental and irreversible failure.

We begin by discussing how problems associated with the basic components of organizational learning (limited memory, form of the search process, speed of learning, and the structure of information) affect the success and reliability of incremental trial and error learning. After briefly describing the methodology adopted in our study, we next report the results of a series of computer simulations based on a well known production-distribution management game. The final section derives the implications of the simulation results in the context of current theories of organizational learning.

Limitations of Experiential Learning in Organizations

According to the adaptive learning perspective of organizations, organizational choices stem from decision rules that adjust cumulatively on the basis of the

successes or failures associated with past trials and errors (Glynn, Lant & Milliken, 1994). Organizational change is seen as the result of a basic learning process resulting in the updating of routines based on the interpretation of experience (Lant & Mezias, 1992). Routines and beliefs are updated in response to experience through trial and error experimentation and incremental search. The likelihood that a given routine will be used is increased when it is associated with past success in meeting a target, and decreased when it is associated with failure (Cyert & March, 1963). Organizations gradually adopt routines that lead to favorable outcomes.

Adaptive models of organizational learning are typically based on the notion of bounded rationality, according to which individual agents have a limited ability to process information in uncertain and rapidly changing environments. As a result, "Most human decision-making, whether individual or organizational, is concerned with the discovery and selection of satisfactory alternatives; only in exceptional cases is it concerned with the discovery and selection of optimal alternatives" (March & Simon, 1958, pp. 140-141).

This search for satisfying solutions is *local* (i.e., new solutions emerge as marginal modification of pre-existing routines); *sequential* (i.e., alternatives are considered one at a time); and *driven by experience*, i.e., learning results in an updating of current routines on the basis of interpretation (or recollection) of past experience (Wall, 1993). Levinthal and March (1981) observed over a decade ago that search strategies may be problematic, i.e., they occur when performance is below target and emphasize relatively immediate refinements or they can be innovative, i.e., they occur when performance exceeds target and emphasize search activities that cannot be justified in terms of their expected return for the organization. This explains why the internal dynamics of organizational processes may sometimes produce outcomes of search processes that are not expected. In extending this line of reasoning to the subject of corporate renewal, Mezias and Glynn (1993) documented how organizational processes interact with the effects of environmental ambiguity to create the paradoxes of institution and revolution. In a more recent application of this idea to the management of technology, Mezias and Glynn (1995) found linkages between the degree of organizational inertia and subsequent levels of innovation and refinement that result from organizational processes.

Components of Adaptive Learning

The effectiveness of the process of experiential learning depends on the assumptions about: (1) *memory*, or the extent to which organizations are able to encode inference from history into routines that guide future behavior, (2) *form of the search process*, or how new information is obtained and in which direction the organization moves to obtain it, and (3) *speed of learning*, or the speed at which the organization can adjust its current state to the desired state. An additional element that is needed to understand the context of adaptive learning as an effective guide to strategic decision making is the structure of information or, more specifically, the terms on which information becomes available to organizational decision makers, i.e., the notion of feedback.

Memory. An organization with perfect memory would always be able to deduce from its current state everything that it once knew (Binmore, 1992). This game-theoretic assumption does not apply without qualifications to organizational learning since organizations and institutions, in general, are known to unlearn and "forget" previous practices (Hedberg, 1981). For example, organizational unlearning may be the result of continuous process related to membership turnover (Cohen, March & Olsen, 1972), or discrete events related to competence-destroying innovations that make previous organizational knowledge obsolete (Tushman & Anderson, 1986).

Determining how exactly limits on organizational memory affect organizational action and how institutions selectively retain information and direct attention is an important agenda for empirical research (Argote, Beckman & Epple, 1990; Carley, 1992; Walsh & Ungson, 1991). For the purpose of the current paper, it is sufficient to note that organizational decision makers cannot be assumed to be boundedly rational while, at the same time, having perfect recall. Since organizational learning is driven by a recursive relation between performance and goals (Levinthal & March, 1981; Lant, 1992), it is important to establish exactly what fraction of the total information derived from previous experience agents actually take into account in their decision-making activity.

Form of Search Process. The trial-and-error search for satisfying courses of action is usually modelled as a linear learning process similar to a hill-climbing search algorithm (Lounaama & March, 1987). When the decision maker is close to her goal (i.e., the distance between desired state and actual state is small) then a small step is taken. Conversely, if the perceived difference from the goal is large, the search will proceed in bigger steps in a direction indicated (or, in extreme cases, dictated) by past experience. This kind of search implies that—if after taking a given step in a given direction the goal appears to be farther away than before (i.e., the agent is worse off)—it is always possible to go back and restart the search in a different direction.

There are good theoretical, as well as practical, reasons to keep the search process simple and intuitive. Search is needed because decision makers cannot possibly know the optimal response and have to experiment to find one that is good enough given their current level of aspiration. Linear relations are used as realistic information processing mechanisms to capture causal connections between the organization and its environment.

Speed. A third issue related to learning by doing has to do with the speed of learning, i.e., the speed at which agents can modify organizational routines as the environment changes. Boundedly rational organizations cannot move from one equilibrium to the next as the environment changes without bearing some adjustment costs and without systematic implications for their life chances (Hannan & Freeman, 1989). In models of adaptive rationality, the speed of learning determines whether a stable (or sustainable) solution is obtained. If an organization closes the gap between its desired and actual state "too fast," then search may be prematurely terminated and the final state may be far from satisfactory or unsustainable (Levinthal & March, 1981; March, 1991). Similarly, if adaptation is slow, search is prolonged unnecessarily and the organization will have to bear the costs

of operating in a subsatisfactory situation for long periods of time. In competitive environments, both outcomes can result in an increased likelihood of organizational failure. The problem seems to be that decision makers often have no way of distinguishing *ex ante* between reacting rapidly to changes and overreacting to transient disturbances. As shown by Lounaama and March (1987, p. 122) in their computer simulations of the internal dynamics of learning organizations, "false lessons are learned as rapidly as true lessons."

Feedback. Feedback can be defined as the "process by which an environment returns to individuals a portion of the information in their response outcome necessary to compare their present strategy with a representation of an ideal strategy" (Balzer, Doherty & O'Connor, 1989, p. 412). The notion of feedback illustrates how the structure of information constrains learning from experience. By structure of information, we mean: (1) the temporal patterns of availability of information related to the delay interval between a given action and its observable outcome, and (2) the sources of information related to how, and from whom, information is obtained.

Many organizational decision-making situations are characterized by substantial time delays, i.e., response to action is not immediate. As the delay period increases, learning becomes more ambiguous due to the interferences of the simultaneous response (and learning) of others. Individuals in organizations cannot experience the complex system of organizational-environment relations directly but only through local information. In this situation, ambiguity is common since "the channels for diffusing observations and interpretations often obscure the events" (March & Olsen, 1976, p. 352). Thus, experience may not be a good teacher not only because of inaccuracy in the perception of performance feedback, but because the relation between actions of individuals and outcomes is confounded by the simultaneous learning and actions of other actors.

Using a computer simulation to explore the dynamics of mutual learning among different organizational members, Lounaama and March (1987) demonstrate that even when performance outcomes are clear, the simultaneous learning of group members and the coordinator can lead to mistakes that become larger over time. They suggest that the effectiveness of experiential learning may be significantly improved by using such heuristics as slowing the rate of adaptation, reducing the simultaneity of change, and scaling the size of change to be neither too small nor too big.

Exogenous Noise Versus Structural Uncertainty

To examine the limitations of experiential learning in organizations, the dominant approach to modelling learning processes has been to simulate the effects of varying degrees of noise on performance (see Lant, 1994, for a review). Thus, the "noise" interfering with accurate assessment of causal links between outcome and action is determined by exogenous sources.

In contrast, a system dynamics-based perspective of adaptive learning in organizations emphasizes the structured and time-dependent interaction between decision-making agents and their environments. Therefore, in a system dynamics-based simulation of organizational decision-making processes, "environments" do

not exist independently from the decision-making agents' understanding of their structure. Accordingly, the "uncertainty" interfering with accurate assessment of causal links between structure and action emerges as the *systematic* result of the misperception of feedback.

Rather than examining how search routines take decision makers from information to specific actions through the update of their beliefs, our model of adaptive organizational learning emphasizes the emergent structure of the decision-making landscape shaped by a stream of individual decisions. Therefore, we argue that search strategies that are typically in the form of linear hill-climbing algorithms are of no avail to decision makers in environments characterized by nonlinearities and rich feedback structures (Dorner, 1980), which is the case in many organizational decision-making situations. To explore whether the limitations of experiential learning in organizations may be even more severe than suggested by previous simulations (Lounaama & March, 1987; Mezias & Lant, 1994), we conducted a series of computer simulations of a production-distribution decision-making process. The next few sections describe the methodology and discuss our findings.

Methodology

Computer Simulation.

Since Cyert and March (1963) introduced the notion of organizational learning, computer simulation has been frequently used as a tool for developing and testing theories about how organizations and individuals learn or fail to learn (Lant, 1994; Lant & Mezias, 1990, 1992; Levinthal & March, 1981; Lounaama & March, 1987). In this modelling tradition, the central result is that when either the causal structure of the organizational environment or the link between action and performance is confounded by random components, learning from experience is not guaranteed to produce desirable results (Lant, 1994).

The computer simulations we conducted to examine the consequences of adaptive learning in organizations concentrate explicitly on the link between individual decisions and the system-level consequences generated by the interaction among individual choices. Using a system dynamics framework to model time-dependent behavior (Morecroft, 1988), we explore the performance outcomes of disequilibrium, complexity, and interdependence, which we claim to be prominent features in real-life economic and management systems.

The Beer Game

The series of computer simulations conducted were based on the *Beer Game*, a well-known experiential simulation that has been used in management education and development for more than 30 years (Forrester, 1961; Jarman, 1963; Sterman, 1989; Senge, 1990). In this experiential simulation, participants have to minimize cost by managing inventory levels in a production-distribution chain. Among the thousands of people from undergraduates to top managers who have played the game, most have discovered that what at the outset looks like a simple task is almost impossible to accomplish.

The *Beer Game* consists of 4 sectors: retailing, wholesaling, distributing, and production. There is an external customer whose demand is determined in advance, although the players do not know the demand pattern. Each sector has an initial small buffer inventory of 12 cases. All the retailer has to do is fill the orders received from the customer, and then decide how much to order from the wholesaler. The wholesaler has to fill the orders received from the retailer, and then decide how much to order from the distributor. The distributor similarly ships beers to the wholesaler and places orders at the factory. The factory delivers to the distributor, and then decides how much to produce. The production time for beer is two weeks.

The orders from the customer are represented as a stack of cards. Each week the retailer takes the top card that represents the customer's demand for that week. Customer demand is simple, a step increase in week 5 from 4 cases to 8 cases of beer and then kept constant at 8 cases for the rest of the game. As in real life, not everything can happen at the same time, so there are built-in mail and shipping delays: it takes two weeks to mail an order and two weeks to ship the requested amount of beer from one sector to the next. This delay structure exists between all sectors. Finally, it is not possible to cancel orders.

If a sector is unable to deliver the requested amount of beer to the sector downstream the remainder of the order will go into the backlog and will be delivered when the sector receives beer from its supplier, which might take one or more weeks. It is assumed that it is more costly to end up in a stock-out situation than in a situation with surplus of inventory. Inventory-carrying cost is \$0.5 per unit of beer per week. Stock-out associated with the possibility of losing customers provides an incentive to hold some inventory. In the game, the stock-out costs are \$1.0 per units of beer per week.

When the game is used in the role-playing mode, each sector is controlled by one or two people, who have to make the decision described above. In principle, the players of different sectors are not allowed to communicate with each other. This means that the retailer is the only one who knows the customer's demand. The wholesaler can try to estimate it from the orders he receives from the retailer, the distributor from the orders received from the wholesaler, and so forth. Effective inventory is defined as (inventory—backlog), i.e., if the effective inventory is negative there is a backlog.

Figure 1 shows the structure of the game in terms of the flow of orders and cases of beer. The system described is a typical cascaded production-distribution system, where each sector has its own small buffer stock, and it is not specific to beer production in any obvious way. Similar to other dynamic systems, a production-distribution system consists of two generic parts: the notions of a state (the essential information about a system) and a dynamic (a rule that describes how the state evolves with time). We next discuss the parameters of these basic components as they pertain to the development of a simulation model.

A Simulation Model of the Beer Game. To model the decision making of the participants in the Beer Game, Sterman (1989) used a heuristic known as anchoring and adjustment. The fundamental idea behind this heuristic is that the decision maker chooses a variable that is known to him (i.e., last week's order) and uses it

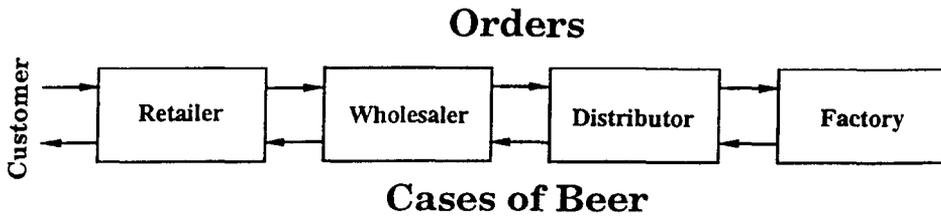


Figure 1. The Beer Game Decision Making Structure:
A Cascaded Production-Distribution System

as an indicator (or anchor), and then adjusts the order accordingly (i.e., the adjustment for inventory and supply line) (Tversky & Kahneman, 1974; Hogarth, 1980). Sterman (1989) confirmed the validity of this heuristic by conducting econometric tests of experimental evidence drawn from 48 trials (192 subjects consisting of undergraduate, MBA, Ph.D. students, and senior executives) collected over a period of four years. Appendix 1 presents the equations behind the decision-making model, which is based on the following three criteria:

1. As demand increases, orders increase as well and vice versa. There is a lag, however, in this response due to the time taken by decision makers to form a view of the extent and permanence of demand changes.
2. Each sector has a target inventory. It is possible to adjust the actual inventory to this target inventory (if there is less than desired inventory, ordering should be increased and vice versa).
3. Each sector keeps track of its supply line, i.e., what is ordered but not yet received. If the supply line grows larger than the desired supply line (given as the supply line target: desired supply line multiplied by expected orders), then orders are reduced to compensate, and vice versa.

Using the decision making heuristic validated by Sterman (1989), Mosekilde, Larsen, and Sterman (1991) developed a computer simulation model of the Beer Game that captures the physical flow as well as the decision-making process in the game, and is capable of producing a variety of different behaviors. As we will discuss below, we used this model to examine the boundaries between regions of stable and unstable behavior that can be found on a policy-making space.

Figures 2a and 2b show results from a short simulation of the model. Figure 2a shows the orders placed by the four sectors and Figure 2b the effective inventory (i.e. inventory—backlog). Shortly after the step increase in customer demand in week 5, the retailer realizes that his inventory is falling, and consequently starts to increase his orders to the wholesaler. As the demand from the retailer increases, the wholesaler's inventory will rapidly disappear. This means that the retailer cannot get the increasing amount of units already ordered. The retailer perceives that nothing is happening, and therefore places new larger orders. Nothing, however, gets delivered. Watching inventory fall, the retailer might even increase orders more, and may start to panic as backlog grows while still only a fraction of what was ordered is being

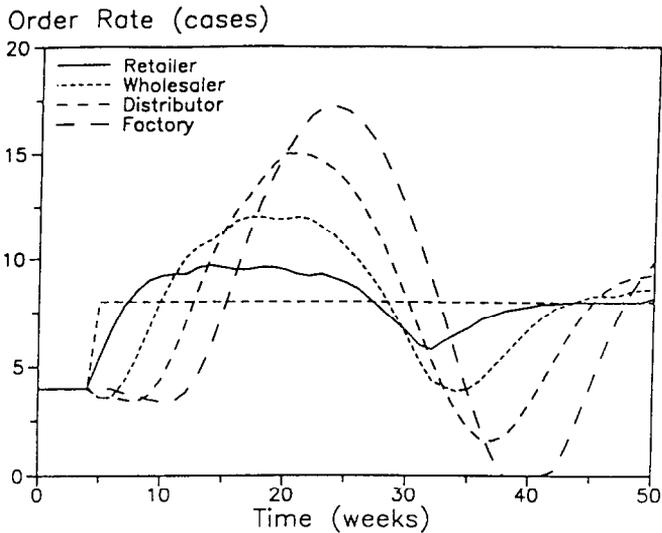


Figure 2a

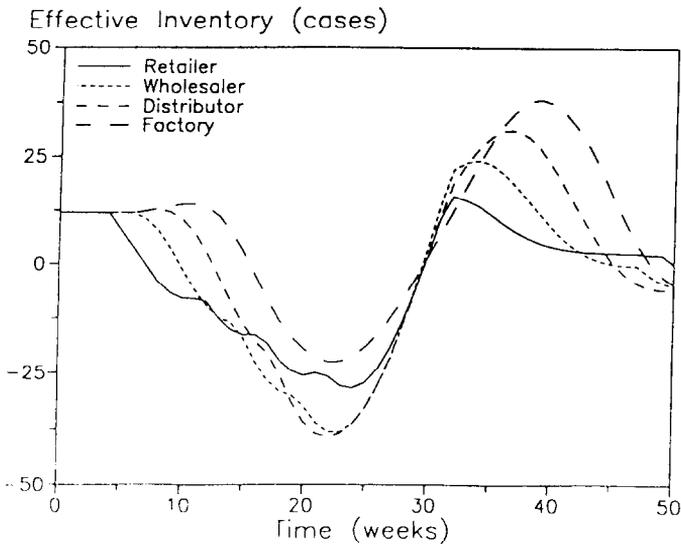


Figure 2b

Figure 2. Patterns of Policy Making Observed in the Beer Game Simulation

received. The wholesaler has the same experience, although it is likely to be slightly worse because of the retailer's panic, which causes more ordering than is necessary. The distributor is even worse off because of having to deal with the panic of both the retailer and wholesaler. As the factory finally discovers the explosive growing

demand and production is increased to cope, more than 10 weeks have already past. As the factory increases production all the sectors will now discover that they have been ordering too much and will stop ordering. This leaves all sectors with large inventories. Over time these inventories are finally eliminated, and a new cycle of backlog and excess inventory might start.

Basic Types of Behavior in the Model

The model described above is capable of producing a number of different types of behavior including stable, periodic, and complex (chaotic, and hyperchaotic) behavior (Devaney, 1988). Figures 3 to 5 illustrate some of these types by showing the variations of the distributors' effective inventory over the first 1000 weeks.

Stable Behavior. After some initial oscillations, the stable behavior model reaches a solution where all variables have a constant final value as shown in Figure 3. In Figure 3, it takes around 150 weeks to reach this stable state. This is by far the most common approach to modelling behavior, as most economic models have focused on equilibrium.

Periodic Behavior. After an initial transient (time to steady state, i.e., the time before the model reaches periodic behavior), the periodic behavior model repeats a certain pattern or sequence over and over again, as the model evolves. Each

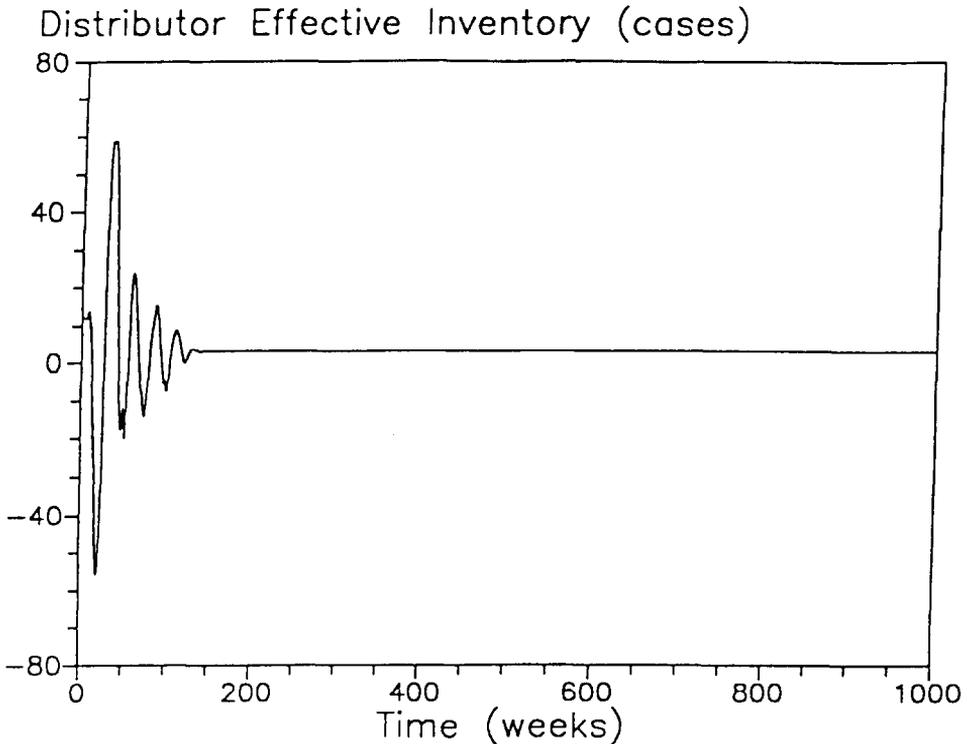


Figure 3. Stable Behavior Model of Effective Inventory

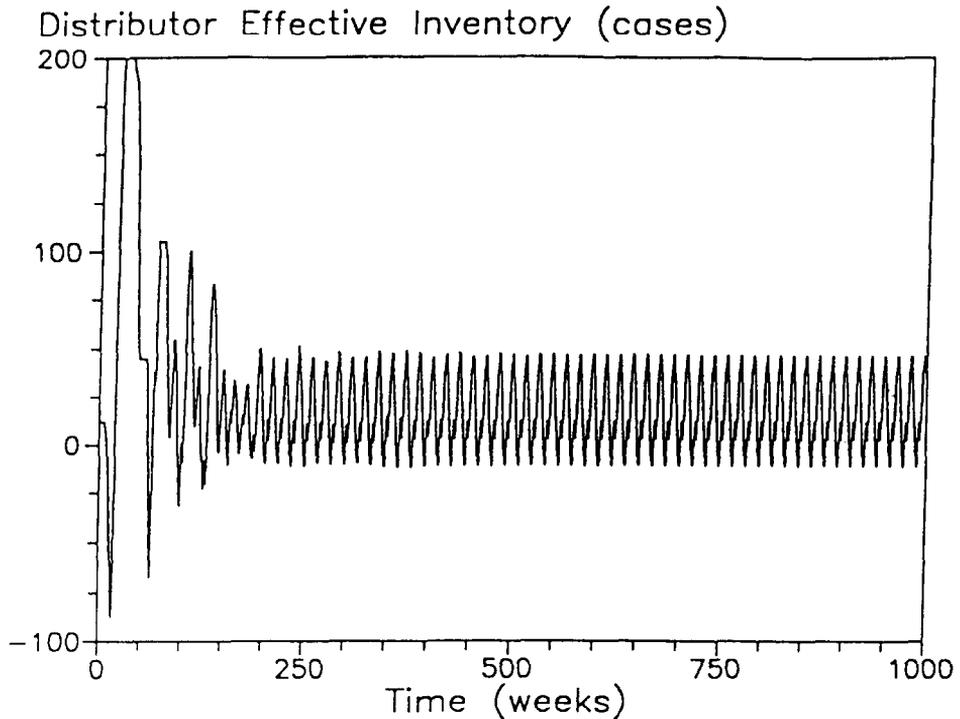


Figure 4. Periodic Behavior Model of Effective Inventory

sequence has fixed periodicity with an arbitrary length: A period can be one or less time unit or thousands or millions of time units long. Figure 4 shows an example of this. In this case, the period is relatively short—around 16 weeks. As shown in Figure 4, it takes some time before the system reaches this periodic mode.

Chaotic Behavior. For a model to be chaotic, the output must be characterized by the following three conditions.

1. The output from the model must be aperiodic, which means that there is no repeated pattern in the output. Figure 5 shows an example of this.
2. There must be sensitivity to initial conditions. Given an infinitesimal change to one of the initial conditions, a simulated trajectory of the modified model must diverge exponentially from the original model. An example is shown in Figure 6. Here, the model was simulated twice: first with all 4 inventories at 12.00 units, and then with the initial inventories as 11.999999999 units. The distance between the two trajectories was calculated, normalized and plotted. As shown in Figure 6, exponential growth of the distance between the two trajectories occurs until nonlinearities of the model prevent the distance from growing any further.

3. The excursion of the model must be bounded. Nonlinearities should prevent variables in the model from approaching plus or minus infinity.

Quasiperiodic behavior has some of the characteristics of periodic and chaotic behavior. Hyperchaos has all the characteristics of chaotic behavior in addition to other more complicated properties. Although the simulations that we conducted differentiate between stable, periodic, and complex behavior, they do not distinguish between different kinds of complexity (quasiperiodic, chaotic, and hyperchaotic). See Thomsen, Moskilde, and Sterman (1992) for a discussion of hyperchaos.

Results of the Simulation Experiments

Having described the dynamic structure of the policy-making system reflected in the *Beer Game* and the types of behavior it is capable of generating, we are now ready to simulate the model. Computer simulations were designed to investigate the range of qualitative behaviors the model produces as change occurs in the value of the decision parameters— α_s (how fast the inventory gets updated in case of a discrepancy between desired and actual inventory) and β

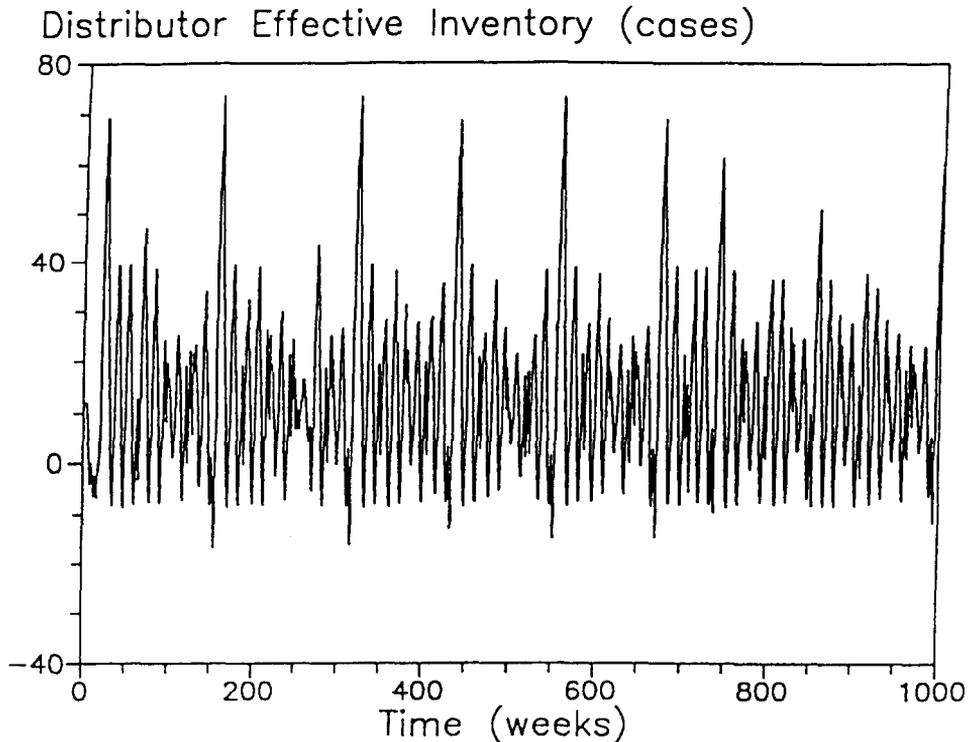


Figure 5. Chaotic Behavior Model of Effective Inventory

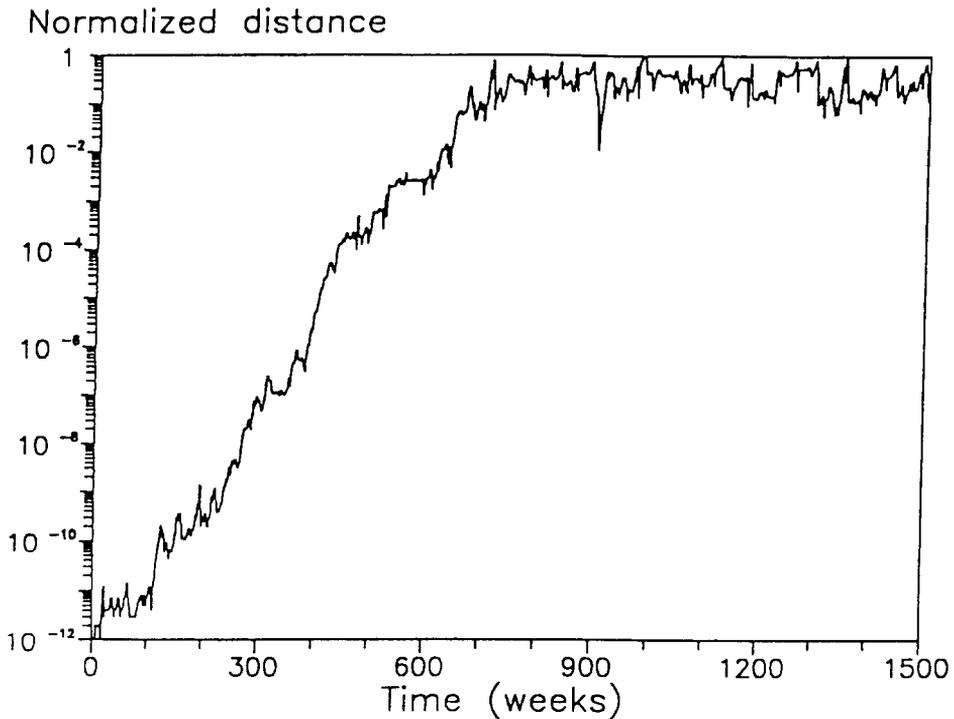


Figure 6. Sensitivity to Initial Conditions in a Chaotic Behavior Pattern

(what fraction of units ordered but not yet received has been taken into account when the new order is placed). The three main categories of behavior that can be found in the model—stable, periodic, and complex—are characterized in terms of the steady state behavior of the model, i.e., of the behavior after several thousands of simulated weeks.

Policy-Making Space. Figures 7a and 7b show the results of a series of simulations using the model described earlier. The range of qualitative behaviors that the model can produce as a function of the policy-making parameters— β and α_s —is represented by the policy-making space shown in Figures 7a and 7b, which was created by simulating the model (varying β and α_s) 201 * 201 times. The composition of the policy-making space is defined in terms of three kinds of qualitative behavior: Stable (represented in light grey), periodic (represented in black) and chaotic (represented in dark grey). The policy-making space can be thought of as a topological map. Where a topological map shows the altitude of the area for a given set of parameters, the policy-making space shows what kind of long-term qualitative behavior obtains for each possible pair of parameter values.

System Instability. Simulation results show that too fast an adjustment process will introduce instability in the system (Lyneis 1982; Morecroft 1983).

Mathematically, this is what happens when α_s is large. Figure 7a shows clearly that the aperiodic solutions are in the area where α_s is close to 1. This indicates that the adjustment between desired inventory and actual inventory is happening too fast. As soon as the adjustment process is slowed down the system becomes stable. When α_s is smaller than 0.3, only stable solutions appear, i.e., the adjustment process has to take at least 3 weeks.

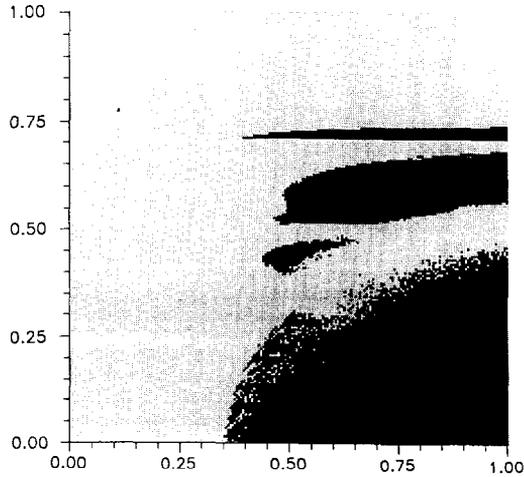


Figure 7a

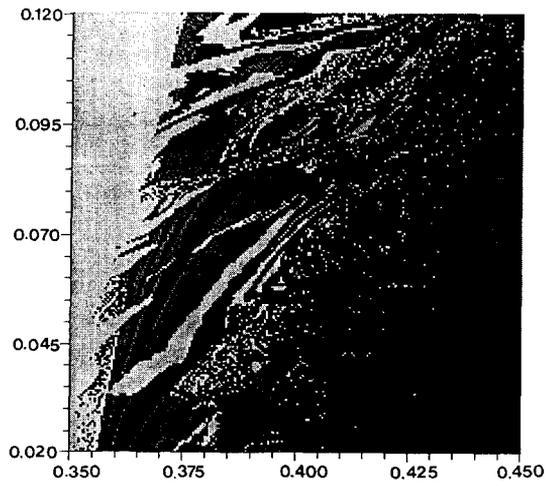


Figure 7b

Figure 7. Topological Maps of the Policy Making Space

When the other parameter— β —is small, i.e., close to 0, a large amount of already ordered goods is not taken into account when the next order is placed. As the units ordered that were not taken into account eventually arrive, however, excessive inventories will start to accumulate in the system. Figure 7a shows clearly that the aperiodic solutions are found only for relative small values of β . As β is increased, the decision maker takes more and more outstanding orders into account when making the new ordering decision, which, in turn, increases the stability of the system. It is important to note that there is no way of explaining the exact behavior found in Figure 7a. We can understand the “macro behavior,” i.e., where the various types of behavior can be found, but not the micro behavior, i.e., which kind of behavior a certain parameter combination of α and β might produce.

To get a better view of the complicated nature of the borderline between periodic and aperiodic behavior, we have magnified a part of the policy making space. The region $0.35 < \alpha_s < 0.45$ and $0.02 < \beta < 0.12$ is shown in Figure 7b. Here, too, 40,000 simulations were run in order to illustrate the behavior of the model over time. As we can see in Figure 7b, the structure is very complex, as “fingers” of stable behavior penetrate deeply into the region of aperiodic behavior. If one looks carefully at the figure, the mixture of solutions could not be more complicated. It is possible to find examples of stable behavior surrounded by periodic solutions, stable solutions surrounded by periodic behavior, aperiodic behavior surrounded by stable solutions, and so on.

Sensitivity to Initial Conditions. A further dimension of complexity is added to Figures 7a and 7b by looking at the initial conditions. Analyzing the behavior of a system by varying its initial conditions is a standard procedure in the study of nonlinear dynamics (Hilborn, 1994). To do this β and α_s are kept constant ($\beta = 0.0775$ and $\alpha_s = 0.3775$) and the number of initial units is left to vary in each simulation from 0 to 24 in Figure 8. In the previous set of simulations the initial inventory was fixed at 12 units. The result of each run in this case can be either a stable (white) or aperiodic (black) solution. A figure like this, which in this case resulted from 160,000 simulations of the model, is referred to as a *basin of attraction*.

Figure 8 illustrates one of the characteristics of complexity—sensitivity to initial conditions. There is a clear structure in Figure 8, although it is impossible to explain why this specific structure emerges in the basin of attraction. Notice that a small change in initial conditions will change the behavior dramatically thousands of weeks later. This reflects one of the key features of chaotic systems: *small differences get magnified and eventually influence macrobehavior*.

Discussion

By examining the policy-making landscapes shown in Figures 7a and 7b and the basin of attraction shown in Figure 8, we can derive some important implications of employing an adaptive learning approach to organizational decision making.

The main explicit objective in the game is to minimize cost. An implicit objective is to keep a stable stream of orders so as to avoid excessive inventories and backlogs and reduce oscillations. In a “real” production-distribution system,

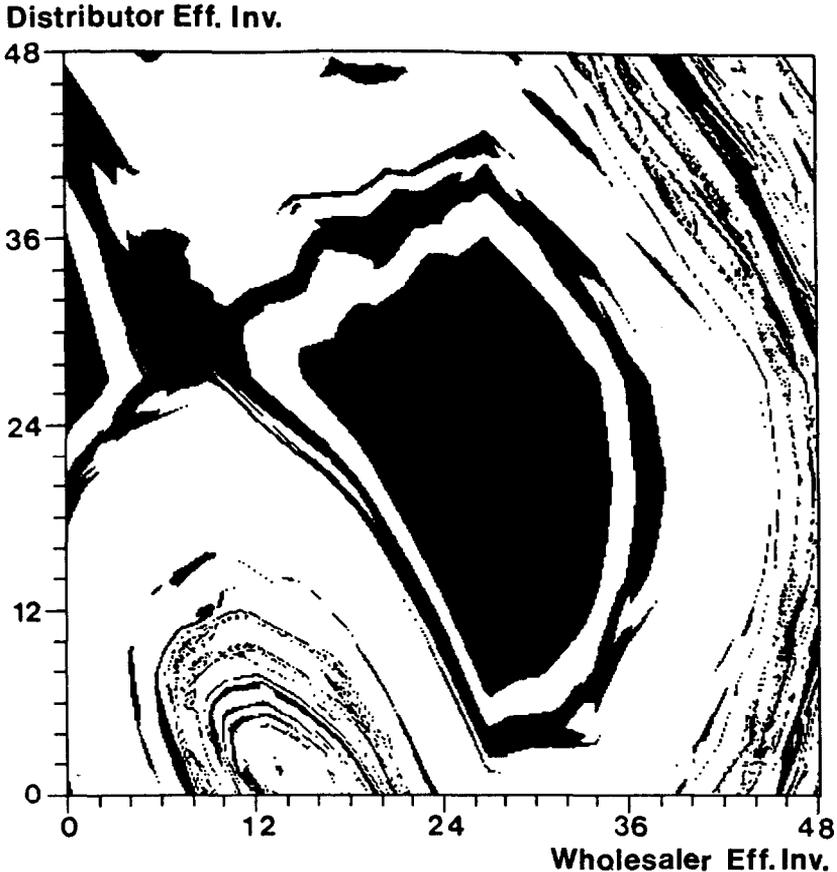


Figure 8. Topological Map of the 'Basin-of-Attraction' for the Beer Game

this is important since production capacity is limited and it has to be kept at a reasonable and steady level of utilization.

There is a close connection between behavior and cost (Mosekilde et al., 1991). For a given combination of the model parameters, the lowest cost is found in the stable area around a global minimum. A rational search strategy in the policy-making space will be, first, to locate the stable area and then to identify the lowest cost point. Once decision makers have located the stable area (i.e., the area where there is no mix of solutions and which is stable independent of the initial conditions), they will be able, in general, to reach the lowest cost using a hill-climbing trial-and-error algorithm.

But what if the decision maker alights in an area of the policy-making space other than the stable area described above? In general, there is no way of knowing if there is a global minimum or where it is likely to be found in case it exists. Even

if the players understood the connection between their behavior and cost, and were assigned the explicit objective of locating the stable area, still they could not rely on experience and past information to find their optimal course of action.

An adaptive learning strategy requires initialization in terms of specific values of α and β , i.e., an initial course of action. If the starting point turns out to be in the border region between stable and complex behavior (e.g., somewhere in the lower right corner of Figure 7b), the initial course of action chosen—the set of (α, β) —will generate a periodic solution. In the context of the policy-making space, learning by experience means selecting new values for α and/or β , which implies a series of small incremental adjustments, since—by definition—experience constrains the size of the area around which improvement is sought. Assuming the first step in a given direction makes the decision maker “better off,” i.e., takes the decision maker to a stable point, the improved condition will indicate that the next step will have to be taken in the same direction (a situation similar to following a course on a topological map determined with a compass). Because the policy-making space contains a mixture of different kinds of solutions, however, doing so could easily lead to a complex area, i.e., an undesirable state. In this situation, experience will not be a useful guide to action.

Things get even worse if we consider the situation described in Figure 8. The natural behavioral response to a move from a stable to a complex solution, i.e., from a relatively desirable to a relatively undesirable state, will be to move back to the original state and explore a different direction. In complex systems, this option may not be available to decision makers. In fact, as we try to return to the original state, we may end up in a totally different situation since the initial conditions changed. What the player may remember as a stable and relatively low-cost state may have turned into an area of high complexity and, consequently, high cost. This implies that even perfect organizational memory will not help adaptive learning, since there is no way of going back.

Communication Structure

Although the current research suggests that the model is capable of producing equally complex behavior even when the number of sectors is reduced to two, we should point out that the behavior of the model is strongly affected by the time delay in the receipt and shipment of orders, and by the structure of communication among the different sectors of our simplified value chain. The type of production-distribution system we simulated can be found in many industries, including automobile manufacturing, leather-shoe production, memory-chip production, and real estate (Senge, 1990). In all these cases, delays could be reduced, for example, by subsuming individual independent units under the hierarchical authority of a larger vertically integrated firm with organizational costs substituting market transaction costs. Among the documented advantages of vertical integration is the development of a coding system which increases communication efficiencies and provides stability in operations (Malmgren, 1961). Determining the circumstances under which vertical integration will actually improve the efficiency and effectiveness of the system as a whole, however, is a subject we leave to agency and transaction-cost theorists (see Mahoney, 1992, for a review).

The second main characteristic of the model that we mentioned above, namely the structure of information exchange and communication among sectors, also warrants some elaboration. It is not always realistic to assume that there is no other form of direct communication between sectors beyond the pure exchange of goods. This situation may represent a highly competitive market with no (or very low) supplier switching costs, where exchange is characterized by a series of spot market transactions among substitutable entities. The central point to this study, however, is not the realism of the model itself, but rather the fact that if complex dynamics can be shown to arise in simple (and simplified) situations, it will be even more likely to characterize systems with a more "realistic" structure of information exchange.

The consequences of alternative information structures can be explored by identifying and re-engineering the basic processes in the game. This includes the consequence of introducing advanced management and inventory information systems, a strategy that has been followed by a number of major retailers. System dynamics modelling and continuous time simulation can provide a useful framework for carrying the process of business process re-engineering all the way from mapping to successful redesign (Van Ackere, Larsen & Morecroft, 1993).

Conclusions

Our computer simulation results have broad methodological, theoretical and practical implications for organizational learning. In their review of the organization learning literature, Glynn et al. (1994) identify two dominant theories: the adaptive learning approach and the knowledge development approach. The first operationalizes organizational learning in terms of changes in organizational routines or search strategies, emphasizes the individual level of analysis, and generally assumes organizations to be unitary entities; the second operationalizes organizational learning as changes in knowledge structures, emphasizes the organizational level of analysis, and underscores the necessity of understanding the systems dynamics by which individual knowledge is transferred and retained within organizations. To advance the current state of understanding beyond what each of these perspectives has contributed, Glynn et al. (1994) call for an approach that helps bridge these two divergent perspectives. Although traditional theories of organizational learning have considered cross-level effects such as the relationship between individual level and organizational level learning, there is also a need to consider system interaction effects.

The approach we have taken to explore organizational learning processes is an *important step in this direction. From the methodological point of view, we build on a consolidated tradition of modelling organizational learning in the context of disequilibrium systems controlled by boundedly rational agents. The system dynamics-based simulation we presented shares with this tradition the fact that no equilibrium condition is assumed or imposed to the system. In fact, one of the main analytical insights that the simulation produced is that simple systems of decisions can have a very long transient (time to equilibrium) and—to the limit—generate periodic or chaotic motion. However, we have gone one step further by drawing the logical*

implications of abandoning equilibrium assumptions. For example, our results indicate that extreme sensitivity to initial conditions suggests limitations to the usefulness of traditional search methods that represent the core element in models of organizational and individual learning. Our models innovate on more conventional representations of learning processes by emphasizing systematic—rather than quasi-random—interaction between agents and their environments.

This methodological innovation has important implications for developing a theory that helps bridge the chasm between microlevel individual processes and macrolevel organizational effects. In our model, “environments” do not exist independently from the agents’ understanding of their structure, but rather the structure of the environment and the action of the decision makers are codetermined. Therefore, the structure of the policy-making landscape emerging from a stream of individual decisions becomes the main object of our modelling effort.

Our approach to theory building is consistent with the interpretive view of organizational life, which views the collective condition of human existence to be the source of meaning, and which emphasizes the centrality of “committed interpretation,” i.e., that people become bound to interacts rather than acts (Weick, 1993). It is also consonant with the work of Porac, Thomas, and Baden-Fuller (1989) on cognitive communities (1989) and of Brown and Duguid (1991) on communities-of-practice, studies that view individual learning as inseparable from collective learning. In the terminology of Drazin and Sandlands (1992), our methodology reflects an autogenetic, rather than an exogenetic or endogenetic, view of organizational behavior. As opposed to an exogenetic perspective, which emphasizes the importance of institutional forces outside the control of any individual actor, and an endogenetic perspective, which emphasizes the importance of intentional design choices, an autogenetic perspective emphasizes the importance of the self-organizing capacities of individuals interacting in a social field.

In examining a typical system of interacting decision makers, our analysis suggests that the problems of experiential learning in organizations may be even more fundamental than has been previously acknowledged. The situation represented by the simulated production/distribution system is defined in terms of 1) *a stable environment*—there is only one initial step change in customer orders, after which demand remains stable; 2) *stable objectives*—the players need only minimize total cost; and 3) *prolonged experience*—once within the policy-making space, decision makers have an arbitrarily long time to explore the “landscape” around themselves. Even in this quasi-ideal situation, decision makers seem to be unable to reconstruct the policy-making landscape accurately.

A disquieting conclusion that emerges from the present research is that stable preferences, lack of structural change, and lack of time constraints are not sufficient to ensure that behavior of economic agents will ever approach the behavior that “would be chosen rationally on the basis of perfect knowledge” (March, 1989: 273). Organizations and individuals learn from experience but experience requires interpretation and, as Lounaama and March put it, “simple learning is no more assured to be sensibly intelligent than anticipatory reason” (1987, p. 108).

Our analysis shows that path dependence—the main consequence of sensitivity to initial conditions—requires that decision makers reach a detailed under-

standing of the global implications of their interdependent actions, i.e., that they approach a learning situation with an articulated strategy. The results of the simulation experiments also suggest that lack of randomness and turbulence in the environment is not sufficient to ensure that optimal solutions to management problems will be found.

From the practical point of view, the fact that random components confounding the causal links between structure and action are not determined by exogenous sources implies that decision makers need to understand the implications of nonlinearities and time delays on the outcomes of their decisions. It is in this sense that we believe that a system dynamics view may provide new insights to managers and help them improve the quality of their decisions.

Acknowledgment: Portions of this article appear in a previous paper (Ginsberg, Larsen & Lomi, 1996) on the embeddedness of strategy. Support for this project was provided in part by the Danish Social Science Research Council and is gratefully acknowledged. The authors would like to thank two anonymous reviewers for their helpful comments.

Appendix

The following equations are based on Sterman (1989) and Mosekilde et al. (1991):

$$O_t = \text{Max}(0, IO_t) \quad (1)$$

Where O_t is the orders at time t . IO_t is the indicated orders at time t . The Max function ensures that orders are not negative.

$$IO_t = L_t^* + AS_t + ASL_t \quad (2)$$

IO is given as the sum of L_t^* the expected loss (i.e. shipments or expected demand). AS_t the adjustment for stock and ASL_t the adjustment for supply line all at time t .

$$L_t^* = \Theta L_{t-1} + (1 - \Theta) L_{t-1}^* \quad (3)$$

The expected loss L_t^* is formulated using adaptive expectations where L_{t-1}^* was the expected loss at time $t - 1$ and L_{t-1} the actual loss at time $t - 1$ - Θ is the relative weight of the expectations and the actual loss.

$$AS_t = \alpha_s(S^* - S_t) \quad (4)$$

Adjustment for stock AS is given as the difference between the desired stock S^* and the actual stock S . α_s is the constant that determinate the "speed" with which any discrepancy between the actual and desired inventory will be closed.

$$ASL_t = \alpha_{sl}(SL^* - SL_t) \quad (5)$$

Adjustment for supply line, ASL is formulated in the same way as AS. Where SL^* is the desired supply line and SL is the current supply line. The constant α_{sl} determines how much of the already ordered but not yet received orders are taken into account when the order decision is taken at any given time. β is defined as:

$$\beta = \alpha_{s1}/\alpha_s \quad (6)$$

and S' as

$$S' = S_t + \beta SL_t^* \quad (7)$$

If we merge equation (2) to (7) the decision rule can be described as follows:

$$IO = L + \alpha_s(S' - S_t^* - \beta SL_t) \quad (8)$$

[†]Note that in this formulation S' is kept as a constant although it might be more realistic to let the desired supply line SL_t^* be a variable proportional to the expected acquisition delay and the expected throughput (Sterman, 1989). However, this change will not alter the arguments in the paper, but make the computer simulations much more time consuming.

References

- Argote, L., Beckman, S. & Epple, D. (1990). The persistence and transfer of learning in industrial settings. *Management Science*, 36: 140-154.
- Balzer, W.K., Doherty, M.E. & O'Connor, R. (1989). Effects of cognitive feedback on performance. *Psychological Bulletin*, 106: 410-433.
- Binmore, K. (1992). *Fun and games*. Lexington, MA: D.C. Heath.
- Brown, J.S. & Duguid, P. (1991). Organizational learning and communities-of-practice: Toward a unified view of working, learning and innovation. *Organization Science*, 2: 40-57.
- Carley, K. (1992). Organizational learning and personnel turnover. *Organizational Science*, 3: 20-46.
- Cohen, M.D., March, J.G. & Olsen, P. (1972). A garbage can model of organizational choice. *Administrative Science Quarterly*, 17: 1-25.
- Cyert, R.M. & March, J.G. (1963). *A behavioral theory of the firm*. Englewood Cliffs, NJ: Prentice Hall.
- Devaney, R.L. (1988). *An introduction to chaotic dynamics*. New York: Benjamin.
- Dorner, D. (1980). On the problems people have in dealing with complexity. *Simulation and Games*, 11: 87-97.
- Drazin, R. & Sandlands, L. (1992). Autogenesis: A perspective of the process of organizing. *Organization Science*, 3: 230-249.
- Forrester, J.W. (1961). *Industrial dynamics*. Portland, OR: Productivity Press.
- Ginsberg, A., Larsen, E. & Lomi, A. (1996). Generating strategy from individual behavior: A dynamic model of structural embeddedness. Pp. 121-147 in J.A.C. Baum & J.E. Dutton (eds.), *Advances in strategic management*. Greenwich, CT: JAI Press.
- Glynn, M.A., Lant, T.K. & Milliken, F.J. (1994). Mapping learning processes in organizations: A multi-level framework linking learning and organizations. Pp. 43-84 in C. Stubbart, J.R. Meindl & J.F. Porac (eds.), *Advances in organizational cognition and organizational information processes*. Greenwich, CT: JAI Press.
- Hannan, M.T. & Freeman, J. (1989). *Organizational ecology*. Cambridge, MA: Harvard University Press.
- Hedberg, B. (1981). How organizations learn and unlearn. Pp. 3-27 in P.C. Nystrom & W.H. Starbuck (eds.), *Handbook of organizational design*, Vol. 1. Oxford: Oxford University Press.
- Hilborn, R.C. (1994). *Chaos and nonlinear dynamics*. New York: Oxford University Press

- Hogarth, R. (1980). *Judgement and choice: The psychology of decision*. New York: John Wiley & Sons.
- Jarmain, W.E. (1963). *Problems in industrial dynamics*. Cambridge, MA: MIT Press.
- Kahneman, D., Slovic, P. & Tversky, A. (1982). *Judgement under uncertainty: Heuristics and biases*. Cambridge, MA: Cambridge University Press.
- Lant, T.K. & Mezias, S.J. (1992). Aspiration level adaption: An empirical exploration. *Management Science*, 38: 623-644.
- . (1994). Computer simulations of organizations as experiential learning systems: Implications for organization theory. Pp. 195-215 in K.Carley & M. Prietula (eds.), *Computational Organization Theory*. Hillsdale, NJ: Erlbaum.
- Lant, T.K. & Mezias, S.J. (1990). Managing discontinuous change: A simulation study of organizational learning and entrepreneurial strategies. *Strategic Management Journal*, 11: 147-179.
- Lant, T.K. & Mezias, S.J. (1992). An organizational learning model of convergence and reorientation. *Organizational Science*, 3: 47-71.
- Levinthal, D. & March, J.G. (1981). A model of adaptive organizational search. *Journal of Economic Behavior and Organization*, 2: 307-333.
- Lounaama, P.H. & March, J.G. (1987). Adaptive coordination of a learning team. *Management Science*, 33: 107-123.
- Lyneis, J.M. (1982). *Corporate planning and policy design: A system dynamics approach*. Cambridge, MA: MIT Press.
- Mahoney, J.T. (1992). The choice of organizational form: Vertical financial ownership versus other methods of vertical integration. *Strategic Management Journal*, 13: 559-584.
- Malmgren, H.B. (1961). Information, expectations, and the theory of the firm. *Quarterly Journal of Economics*, 75: 399-421.
- March, J.G. (1989). *Decisions and organizations*. New York: Basil Blackwell.
- . (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2: 71-87.
- March, J.G. & Olsen, J.P. (1976). *Ambiguity and choice in organizations*. Bergen, Norway: Universitetsforlaget.
- March, J.G. & Simon, H.A. (1958). *Organizations*. New York: John Wiley.
- Mezias, S.J. & Glynn, M.A. (1993). The three faces of corporate renewal: Institution, revolution, and evolution. *Strategic Management Journal*, 14: 77-101.
- . (1995). Using computer simulations to understand the management of technology: Applications for theory development. *Technology Studies*, 2(2): 175-215.
- Mezias, S.J. & Lant, T.K. (1994). Mimetic learning and the evolution of organizational populations. Pp. 179-198 in J.A.C. Baum & J.V. Singh (eds.), *Evolutionary dynamics of organizations*. Oxford, UK: Oxford University Press.
- Morecroft, J.D.W. (1983). A systems perspective on material requirements planning. *Decision Science*, 14(1): 301-320.
- Morecroft, J.D.W. (1988). System dynamics and microworlds for policy makers. *European Journal of Operational Research*, 35(3): 301-320.
- Morecroft, J.D.W. & Sterman, J.D. (eds.). (1994). *Modelling for learning organizations*. Portland, OR: Productivity Press.
- Mosekilde, E., Larsen, E.R. & Sterman, J.D. (1991). Coping with complexity: Deterministic chaos in human decision making behavior. Pp. 199-229 in J. Casti & A. Karlquist (eds.), *Beyond belief: Randomness, prediction and explanation science*. Boca Raton, FL: CRC Press.
- Nelson, R. & Winter, S. (1982). *An evolutionary theory of economic change*. Cambridge, Mass: Belknap Press.
- Porac, J.F., Thomas, H. & Baden-Fuller, C. (1989). Competitive groups as cognitive communities: The case of Scottish knitwear manufacturers. *Journal of Management Studies*, 26: 397-416.
- Rachlin, H. (1991). *Introduction to modern behaviorism*. New York: W.H. Freeman.
- Senge, P.M. (1990). *The fifth discipline*. New York: Doubleday.
- Sterman, J.D. (1989). Modelling managerial behavior: Misperception of feedback in dynamic decision making. *Management Science*, 3: 321-339.
- Thomsen, J.S., Mosekilde, E. & Sterman, J.D. (1992). Hyperchaotic phenomena in dynamic decision making. *SAMS*, 9: 137-156.
- Tushman, M.L. & Anderson, P. (1986). Technological discontinuities and organizational environments. *Administrative Science Quarterly*, 3: 439-465.
- Tversky, A. & Kahneman, D. (1974). Judgement under uncertainty: Heuristics and biases. *Science*, 185: 1124-1131.
- Van Ackere, A., Larsen, E.R. & Morecroft, J.D.W. (1993). Systems thinking and business process redesign. *European Management Journal*, 11: 412-423.
- Wall, K. (1993). A model of decision making under bounded rationality. *Journal of Economic Behavior and Organization*, 21: 331-352.
- Walsh, J. & Ungson, G. (1991). Organizational memory. *Academy of Management Review*, 16: 57-91.
- Weick, K.E. (1993). Sensemaking in organizations: Small structures with large consequences. Pp. 10-37 in J.K. Murnighan (Ed.) *Social Psychology in Organizations*. Englewood Cliffs, NJ: Prentice Hall.