Mathematical Organization Theory Workshop

April 25 and 26, 1993

Chicago, IL

Organized by

Kathleen Carley, Associate Professor
Department of Social and Decision Sciences
Carnegie Mellon University, Pittsburgh, PA
and
Michael Prietula, Associate Professor
Graduate School of Industrial Administration
Carnegie Mellon University, Pittsburgh, PA

Table of Contents

Agenda

Saturday, April 25

8:30- 9:00	Registration, Coffee and Rolls
9:00- 9:30	Welcome and overview Kathleen Carley (CMU)
9:30-10:00	
10:00-10:30	
10:30-11:00	Break
11:00-11:30	
11:30-12:00	
12:00-12:30	Lunch
12:30- 1:00	Lunch
1:00- 1:30	Lunch
1:30- 2:00	Lunch
2:00- 2:30	
2:30- 3:00	
3:00- 3:30	
3:30- 4:00	Break
4:00- 4:30	Al Wallace (Rensselaer Polytechnic Institute)
4:30- 5:00	
5:00- 5:30	
5:30- 6:00	Discussion: Michael Prietula and Kathleen Carley (CMU)

Sunday, April 26

8:30- 9:00	Coffee and Rolls
9:00- 9:30	
9:30-10:00	
10:00-10:30	
10:30-11:00	Break
11:00-11:15	
11:15-11:45	Richard Burton (Duke University)
11:45-12:15	
12:15-12:45	
12:45- 1:00	Wrapup

Speakers

Linda Argote	(412)268-2290	LAOM+@ANDREW.CMU.EDU
Richard Burton	(919)660-7700	BURTON%DUKESSB.BITNET
Kathleen Carley	(412)268-3252	Kathleen.Carley@centro.soar.cs.cmu.edu
Kevin Crowston	(313)763-2373	Kevin_Crowston@um.cc.unich.edu
Les Gasser	(213)740-4046	gasser@usc.edu
Raimo Hämäläinen		raimo@hut.fi
Bernardo Huberman	(415)812-4147	Huberman@parc.zerox.com
David Kleinman	(203)486-3066	
Benn Konsynski	(617)495-6436	Bkonsynski@HBS.HBS. Harvard.Edu
Panagiotis Kouvelis	(919)660-7833	kouvelis@dukefsb
David Krackhardt	(412)268-6122	krack+@andrew.cmu.edu
Theresa Lant	(212)285-6058	tlant@rnd.stern.nyu.edu
Arie Lewin	(919)660-7832	
Michael Masuch		michael@ccsom.nl
Michael Prietula	(412)268-5067	mp2j@andrew.cmu.edu
Eduardo Salas	(407)380-4651	
Shyam Sunder	(412)268-2103	SS8A+@Andrew.CMU.EDU
William Wallace	41-1256-2361(Switzerland)	
Massimo Warglien	39-41-5208657(Venezia Italy)	

Mathematical Modeling of Parkinson's Law

by
Panagiotis Kouvelis

A widely accepted behavioral law, Parkinson's Law, manifests that "work expands so as to fill the time available for its completion". We provide a modeling framework designed for project management activities that incorporates worker and/or subcontractor behavioral issues as depicted in Parkinson's Law. We develop a descriptive model of the expected completion time of an activity. This model expresses completion time as a function of the time allocated for that activity, i.e., the deadline set by the project manager, and the actual amount of work required by the specific project activity. The developed model is of an analytic nature, and is used to enhance our understanding of activity management.

The Acquisition and Depreciation of Knowledge in Manufacturing: Turnover and the Learning Curve

by
Linda Argote, Dennis Epple, Rukmini Devadas and Kenneth Murphy

Large increases in productivity are typically realized as organizations gain experience in production. As organizations produce more of a product, the unit cost of production typically decreases at a decreasing rate and quality may improve at a decreasing rate. These "learning curves" have been found in many organizations, in both the manufacturing and service sector (Yelle, 1979). There is considerable variation, however, in the rate at which organizations learn (Argote & Epple, 1990; Dutton & Thomas, 1984). Some organizations show remarkable productivity growth while others evidence little or no learning.

The presumption is often made in the organizational learning curve literature that learning is cumulative -- that it persists through time and does not evidence depreciation. There is new evidence, however, that learning may depreciate. Argote, Beckman and Epple (1990) found that recent output was a more important predictor of current production than cumulative output.

The current study investigates the acquisition and depreciation of knowledge in a manufacturing plant. We examine: (1) whether there is evidence of learning at the plant; (2) whether knowledge acquired through learning by doing depreciates, and (3) whether turnover affects the rate of learning and depreciation.

Our data are from a North American truck plant that began production in the late 1980s. We have two years of data since the start of production at the plant. We have collected weekly data on the number of units produced, the labor hours worked, the number of days of operation, and the amount of personnel movement into and out of the plant (turnover). The turnover data are broken out by the reason employees leave the plant. Our analysis approach involves estimating production functions which allow for the possibility that knowledge depreciates.

Preliminary analysis indicate that there is evidence of learning in this production environment: with each doubling of cumulative output, unit cost decline to approximately 83% of their previous value. The relationship between movement of employees into the plant is nonmonotonic. Initially, their movement into the plant has a positive effect on productivity but the effect then becomes disruptive and hurts productivity. The movement out of the plant of highly skilled employees who were selected for apprenticeship programs hurt productivity. None of the other categories of turnover had a significant effect on productivity, although the coefficient of the variable representing those discharged for poor performance was positive. Results are interpreted and theoretical and practical applications are developed.

References

Argote, L. , S.L. Beckman, and D. Epple (1990). The persistence and transfer of learning in industrial settings, *Management Science*, *36*, 140-154.

Argote, L., and D. Epple (1990). Learning curves in manufacturing, *Science 247*, 920-924.

Dutton, J.M., and A. Thomas (1984). Treating progress functions as a managerial opportunity, *Academy of Management Review 9*, 235-247.

Yelle, L.E. (1979). The learning curve: Historical review and comprehensive survey, *Decision Sciences, 10*, 302-328.

A Logic for Bounded Rationality

by
Zhishen Huang, Michael Masuch and Laszlo Polos

Simon's Bounded Rationality serves as the point of departure for modern organization theory. Unfortunately, the original notion of bounded rationality is somewhat ambiguous, and subsequent attempts at clarification have not been very successful. We argue that modern formal logic can be brought to bear on this issue. We propose a logic for bounded rationality by combining notions from V. Wright's preference logic, Pratt's dynamic logic, and Stalnaker's conditional logic. Although presently restricted to the one-agent case, the logic can be extended to multi-agent settings.

The Unicorn Factory

by
Massimo Warglien

Virtually all organizations must cope with unrelenting novelty. Problems and situations are never identical. But they may be similar enough to allow an organization to work as if they were the same. It is commonly assumed that there is often enough regularity to justify describing and modelling organizational choice and action as routine-driven processes. While this approach covers much of the spectrum of situations, organizations must occasionally face extremely novel problems and undertake substantially new activities.

Highly novel situations pose peculiar challenges to both organizational choice and adaptation. Rational analysis and assessment of the situation may find insuperable obstacles, while feedback-driven adaptive behavior is impaired by the absence of prior direct experience concerning the new situation. Under such conditions, organizations are condemned to transfer their stock of existing skills and routines by constructing some mapping of past problems and solutions, already experienced, onto new ones. Thus they use the past to define both the structure of the problem to be solved and the solution plans.

A driving force of the process of transfering past experience onto new problems is the recognition of some similarity between a situation experienced in the past and the current context. The process is, in many respects, the organizational equivalent of individual analogy making, and it functions similarly, as a substitute for direct experience. In typical organizations no one person possesses all the expertise necessary to assess the new situation and define a new plan or project, and skills and experience are distributed among organizational members. A kind of collective analogy has to be built from smaller pieces of knowledge held by different individuals, who may have different preferences, capabilities, and beliefs as well as differentiated access to the organizational decision making process.

The paper draws on a case study of a biomedical innovation. It describes how, in a new product development project, a network of beliefs is assembled through a process which tightly weaves cognitive judgments with political, coalition-building processes. Filling the gaps of experience and clustering beliefs are processes that intimately blend what actors believe and what they like.

A computer model, inspired by the case study, explores how a set of actors with different beliefs, different preferences and differentiated access to the design and decision making processes can engage in constructive activities by mapping knowledge of past situations on new problems. The model allows to access the dynamics of cooperation and conflict, and the consequences of different ways of structuring the decision making process.

TOWARDS THE DEVELOPMENT OF A COGNITIVELY MOTIVATED THEORY OF ORGANIZATIONS

by
Kathleen Carley and Michael Prietula

The proposed research investigates how artificial intelligence (AI) can be applied to the study of organizations of intelligent individuals. We propose to construct a full AI model of a small organization of humans communicating and cooperating to accomplish a task within particular organizational structures. We will use a collection of computers where on each computer there is running an artificial intelligence model of an intelligent agent. Each artificial agent will be represented on its own computer using an extremely sophisticated software architecture (called Soar) that is capable of serving as a basis for general intelligence and can learn from experience. Furthermore, each agent model is capable of communicating with other agents (as permitted) and can learn through experience. At the same time we will be conducting human studies involving the same number of agents (i.e., subjects), the same task, and the same constraints on specific parameters of interest (training, information sharing, communication, organizational structure). In order to orient this research we will focus in on the question: what coordination schemes are most effective under various conditions? The specific conditions we will examine are: the level of training, the type and amount of information shared, and the complexity of the task. A prototype version of the system, called Plural-Soar, has demonstrated the feasibility of this approach. Plural-Soar uses the warehouse task.

Our research is based on a functioning software architecture, called Soar, which is capable of learning through experience (Laird, Newell and Rosenbloom, 1987) and therefore of being trained. Within that architecture, we developed a prototype system and have demonstrated that a set of networked computers, each running a Soar model of an intelligent agent, can effectively perform a small, but realistic task of retrieving items from a warehouse. Over the next two years we intend to further develop the system to represent several versions of the task reflecting concurrent research into groups of human subjects in similar task environments. This research will,

1)extend and redefine the Plural-Soar model using a new task, the radar tracking task;
2)realize the model as multiple soar agents that can communicate;
3)develop a c-based communication protocol for these agents;
4)help to develop a computationally based theory of micro-organizational behavior.

Team Decision-Making in Complex Environments: Implications for Training

by
Eduardo Salas, Janis A. Cannon-Bowers and Catherine V. Baker

Critical performance in many complex systems depends on the coordinated activity of a team of individuals. Aircrews, fire fighting teams and command and control teams are all examples of teams who operate in situations where ineffective performance can have disastrous consequences. Despite a considerable amount of research into the area of team training, however, relatively little is known about how to train teams or the appropriate training content for teams. This is particularly true in the area of team decision-making where teams must gather, process and integrate information in support of a decision. In recent years, several researchers have suggested that team performance can be understood in terms of shared mental models of the task and team among team members.

The notion of team mental models and how these relate to team effectiveness has several implications for the understanding of team performance and training. As an explanatory mechanism, the team mental model construct is useful in understanding how teams are able to coordinate behavior and select task strategies in absence of explicit coordination activities. Under conditions of high workload, time pressure and other kinds of stress, such implicit coordination appears to be critical (Kleinman & Serfaty, 1989).

With respect to training, the shared mental model idea suggests that training strategies designed to foster development of shared mental models has the potential to improve team performance. Recent research regarding the success of efforts to train mental models for system operation offers preliminary evidence that such training may be possible (Rouse & Morris, 1986). For example, research suggesting that particular knowledge structures (i.e., mental models) can be trained provides support for the notion that common expectations for the task and team can be developed through training.

In summary, this presentation will address three related points: 1) team performance in complex systems requires application of new methods to capture performance (e.g., petri nets) as well as new instructional strategies for training; 2) it can be reasonably hypothesized that team performance effectiveness will be a function of the extent to which team members hold shared or common mental models; and 3) specific training strategies (e.g., cross-training) can be advanced that may be useful to foster shared mental models.

The Dynamics of Cooperation

by
N. Glance and B. A. Huberman

We present and study a dynamical model of ongoing collective action among intentional agents whose choices depend not only on the past but also on their expectations as to how their actions will affect those of others. In this model agents act on the basis of information that can be uncertain at times. We show that under these conditions the onset of overall cooperation can take place in a sudden and unexpected way. Likewise, defection can appear out of nowhere in very large, previously cooperating groups. These outbreaks mark the end of long transient states in which defection or cooperation persists in groups that cannot sustain it indefinitely.

Collective action problems are characterized by the impossibility of exclusion; that is, no member of a group engaged in collective action can be excluded from enjoying the benefits of the group's efforts. On the other hand, collective goods can have varying amounts of jointness of supply, which is the degree to which one agent's consumption of the good does not reduce the amount available to any other. The amount of jointness determines in part the dependence of cooperative outcomes on the size of the group. Thus, for example, some studies have shown that overall cooperation is undermined as the group increases in size, others that, to the contrary, that it is more likely for larger groups. The latter result was obtained for public goods, a special subset exhibiting perfect jointness of supply, while the former holds for divisible goods with limited jointness of supply. In this paper, we consider divisible collective goods.

Regardless of whether or not a group of a given size can exhibit ongoing cooperation, there remains the issue of how is it that such a state is reached, if ever. For one can imagine situations whereby a group initially exhibits cooperative behavior in spite of its being too large to do so, to only gradually evolve into collective defection. Conversely, a large non-cooperative group could undergo a drastic reduction in size and the relevant question would then be how long does it take for it to switch to global cooperation.

Following a long tradition, we study the free rider problem as a repeated-person prisoner's dilemma, wherein the benefit obtained by an individual from cooperating in producing the good is outweighed by the cost of cooperation for the one-shot game. We recast the interaction as an asynchronous dynamic game in which each individual reconsiders its decision at an average reevaluation rate, using delayed information of the level of provision of the good. The iterated game is of finite duration, and individuals decide to cooperate or defect by determining which choice maximizes their expected share of the good for the remainder of the game. In addition, uncertainty is introduced into the relation between individual effort and group performance to model imperfect information and bounded rationality.

The type of expectations that we consider consist of two components: each individual believes (1) that future aggregate collective behavior is directly influenced by the individual's choices in inverse proportion to the group size; and (2) that the interaction is of finite duration characterized by a horizon length. Thus, individuals believe that in the long run their actions encourage similar actions on the part of others in the group. There are various mechanisms by which this might occur: imitation and establishment of conventions or norms are a few. In addition, individuals believe that the ability of their actions to encourage like actions increases with the number of individuals who have chosen to contribute to the collective good (as opposed to free riding). In some sense, this reflects a belief that contributing individuals form a core group whose reactions are much more sensitive to fluctuations in the amount of the good produced.

We show that an individually rational strategy of conditional cooperation emerges from the individuals' expectations and beliefs. Individuals cooperate if they perceive the fraction cooperating to be greater than some critical amount and defect otherwise. This strategy of conditional cooperation is reminiscent of the successful tit-for-tat cooperates if and only if its opponent cooperated in the previous turn and defects otherwise.

We then analyze the dynamics of fluctuations away from Nash equilibria using a thermodynamis-like formalism. Besides confirming that there exists a critical group size beyond which cooperation is not sustainable, we find additional, intriguing effects. Our results reveal several different dynamical regimes that should be observable as the size of a group changes. In one regime, although cooperation may persist for very long times even for groups exceeding a critical size, group behavior eventually decays to overall defection. In another situation a system can be stuck in a non-cooperative state even though its size is well below that guaranteeing long term cooperation. These effects are shown to depend strongly on the degree of uncertainty pervading the system, as well as the length of the individual's horizons.

We also include the results of computer simulations which confirm our analytical predictions and present an example of a phenomenon in which the average behavior is not typical. By this we mean group behavior that flip-flops between the strategies of mutual cooperation and mutual defection, without ever settling into one or the other.

The Effects of Team Reward Structure on Team Coordination

by
P. Shi, P. B. Luh, and D. L. Kleinman

Our previous research on human team decision-making focused on dynamic multi-task environments in which the players always shared the common goal, e.e., each individual within the team was rewarded on the basis of the team's performance. There was no study of how humans coordinate their actions and resources when they are rewarded based solely on their individual performances.

This research, which has joint empirical and analytical components, examines the mechanisms by which different (team vs. individual) reward structures affect decision-making and coordination strategies within a simple two-person parallel team. A problem is formulated in which two decision-makers (DMs), each owning a limited amount of resources, are to process multiple types of tasks, some of which require coordinated actions by both DMs. A team-in-the loop experiment is designed and operationalized across three independent variables: reward structure (team vs. individual), task coordination requirement (sequential vs. parallel task processing actions), and resource availability (low vs. high).

The main empirical findings show that under different reward structures DMs behave differently in selecting tasks, coordinating and scheduling their resources. Less explicit coordination activities (via communication) and better performance are observed under the team reward structure. Our interpretation is that a team reward structure reduces a DM's uncertainty of his teammate's likely actions, hence enabling more implicit coordination on task processing and a reduced need for explicit communication. The experimental results also show that (1) a team reward structure is most beneficial within a more interdependent (parallel) task processing environment; and (2) a variation in resource availability has a more pronounced effect on teams operating under an individual reward structure.

A mathematical model is presented for describing the team decision-making and coordination processes. The model uses a decision tree for the team reward case and a game tree for the individual reward case, but instead of constructing a complete decision/game tree at each decision epoch the model only considers two stages. The first stage evaluates the current choices available to both DMs, while the second stage measures the impacts of the next decision by combining the current reward with an estimate of future outcome at each leaf node of the decision/game tree. The model seeks the maximum team reward for the team case, and a Nash equilibrium solution for the individual case.

A Decision Logic for Crisis Management

by
Giampiero E.G. Beroggi and William A. Wallace

Events that initiate a crisis situation occur very infrequently, so infrequent that we can only consider their possibility not their probability. In addition, the consequences of these events represent an opportunity for significant gains or losses to the organization - so significant that they threaten the very existence of the organization in its present form. Therefore, crises initiate active response from the organization as opposed to the use of guidelines or directives based upon formal plans. These events cannot be "planned for" but must be "responded to". And the response must be made in a timely manner, where the time available to make decisions and take action is determined by the event itself.

Research on crisis management has proposed and analyzed various organizational forms using, for example, cases based on post-crisis field investigations, field studies of organizations engaged routinely in activities with the potential for high impact failures, and computer simulation. In addition, decision aids have been developed and assessed for their contribution to crisis response. Literature on decision making under stress has also provided insights into the crisis situation. This research does conclude that organizations that attempt to plan for crises by prespecifying an organizational form and training their managers in procedures and directives based upon this form will be faced with even graver problems in response when the event is at all different in its characteristics or potential consequences. Our work focuses on providing decision support to the crisis manager in these situations, when the vent is at all different than anticipated or the planned for actions cannot be taken or have unanticipated consequences.

The objective of the research reported on in this paper is to propose and test a decision logic for crisis management. The logic is based upon the normative view of rationality, modified by the limitations placed upon human cognition by the crisis situation. It assumes that managing a crisis consists of continuous monitoring the assessment of (1) the environment in which crises take place, (2) the planned set of activities for response, and (3) the anticipated results of those activities. Decisions to change the planned activities to a new course of action are based on the crisis manager's experience and intuition - which are limited in a crisis situation.

We propose a new approach for crisis management that capitalizes upon advances in information and communications technology and that considers the fact that humans will always be involved in the process consisting of monitoring, assessment, and decision making. The real-time environment in which crises take place is defined using a terminology analogous to classical control theory. A control unit capable of sensing and reasoning supports the on-site managers in determining the optimal course of action. Real-time events (RTE) occur unexpectedly and are assessed by the control unit. Suggestions to alter the course of action are transferred to the on-site managers. The decision logic is based on a preference algebra. We address both the case of assessment and decision making by a single human as well as that of small groups.

Each activity in crisis management is proceeded by a decision to select this activity and leads to a new decision on which activity should be taken next. Therefore, decision D_j has been made to select activity A_{ij} which leads itself to decision D_j. A course of action CA^k consists of k+1 decisions and k activities: $CA^k = \{A_{12},...,A_{ij},...,A_{rk}\}$. The last decision (k+1) does not lead to a new activity but to the end of the crisis situation. This notation can be used for a representation of the relations between the possible decisions in a graph structure. Let $CM_t(D,R)$ be the topological structure of the crisis management graph which consists of a set of decisions D (nodes, vertices) and a relation R which is irreflexive. The adjacency matrix A has as its a_{ij} entry the value 1 if decision D_i can lead to decision D_j (i.e., if activity A_{ij} can be taken) and O otherwise. Let P be the set of preference classes. the entry p_{ij} is the preference of taking activity A_{ij}. A course of action (CA) is a (directed)I walk on the topological graph structure. It consists of a finite, ordered set of activities A_{ij}, such that for each pair of adjacent activities the following property holds: $CA=\{...,A_{ij},A_{jk},....\}$; i.e., the end-node of any activity is identical with the start-node of the succeeding activity. A walk of k activities is called a k-walk; i.e., a k-walk is a course of actions involving k activities. The preference for taking a k-walk is the intersection of the preferences of all activities. It is computed according to the preference algebra. A graph $CM_t(D,R)$ where the activities have assigned preferences at time t is called the preference graph.

According to this terminology, crisis management activities are represented in a implicit way. The most preferred course of action for a particular situation is computed with a graph search algorithm which is based on the preference algebra. Due to the nature of the initiating event, the operator must reassess the benefits, risks, and other attributes of all the activities that are affected-in real-time. After this reassessment, new preferred courses of action are generated.

When more than one manager is involved in the assessment and choice process, individual assessments might have to be reconciliated because of possible differences.

Assuming that the individuals assess independently of each other, and that they are "equal", i.e., they have the same background, experience, and available information, statistical tests are used to determine if an assessment is feasible. The assessment is feasible if it is consistent and efficient where (1) consistency holds if the rank-orders pass the rank-test, and (2) efficiency holds if the assessments fulfill certain concordance and discordance criteria. Feasibility refers to both the assessment and the choice problem. Since the courses of actions (alternatives) are represented implicitly to the decision makers, no classical strategies for choice can be applied (e.g., elimination by aspect or by alternative). Therefore, feasibility tests can be applied both for assessment and for choice. Several strategies for composing the different assessments and generated alternatives are investigated. One approach is to check the assessments for feasibility and to combine them into a single assessment. The computed most preferred course of action is therefore unique, unless several courses of action have the same overall preferences. Another approach is not to test the assessments for feasibility but only the different courses of action. If feasibility does not hold, reassessment is required. A third approach is to test both assessment and choice for feasibility.

Research is ongoing to test the preference algebra for single individuals in simulated crisis situations. Based upon that work, we will propose and illustrate experiments to test the group preference process using the decision logic.

Simulating Organizations as Experiential, Rule-Based Learning Systems: Implications for Organization Theory

by
Theresa K. Lant

The field of organization theory has begun to recognize that our theories of organizing need to become more longitudinal and evolutionary in nature. However, there is a lack of sufficient theory development in this area because of the difficulty of theorizing about complex phenomena over time. It is difficult to explicate how organizational outcomes. The unfolding of these processes can be observed, however, in a computer simulation. A computer simulation can take a complex set of assumptions, simulate a set of organizational processes, and are present the implications of these processes for organizational outcomes. Behavioral computer simulations, which model the procedures, processes, and decision rules used by a system (Crecine, 1969), have played an important, if limited, role in organizational theory development (Cohen and Cyert, 1965).

In this presentation, I will discuss one approach to using computer simulations to develop longitudinal theories of organizing. This approach models organizations as rule-based learning systems with adaptive aspiration levels. The simulation that was developed is built on the assumption that organizations are governed by an experiential learning process. I will discuss the use of this simulation to help build longitudinal theory in three areas of organization theory: 1. entrepreneurial strategies in established organizations (Lant and Mezias, 1990); 2. the punctuated equilibrium perspective of organizational change (Lant and Mezias, 1992); and 3. the role of organizational change in population dynamics (Mezias and Lant, 1992).

The learning model used in the simulation depicts stylized organizations that follow rules of behavior as specified in the literature on organizational learning. First, organizations have a target level of performance or aspiration level to which they compare their actual performance. Second, performance below or above aspiration level affects the likelihood of organizational change because it defines the organization's perceptions of success and failure. Change in behavior is more likely when performance is below aspiration level than when it is above aspiration level. Third, aspiration levels adapt over time in response to performance. Fourth, organizations are boundedly rational; the acquisition and processing in information is costly. In addition, organizations in the simulation are characterized by four core dimensions; one each dimensions firms have two alternatives. Organizational change is measure by observing changes in these dimensions. The simulation has been used to examine the implications of ambiguity, environmental change, and different search rules (adaptive, imitative, fixed, and garbage can) for organizational change and effectiveness.

The Lant and Mezias (1990) paper examines the effectiveness of three search rules (adaptive, imitative, and fixed) in established organizations that are faced with a fundamental restructuring of their environment. They examine the effect of search strategy, level of search, and their interaction on longitudinal patterns of organizational performance, resources, and death in ambiguous and unambiguous environments during stable periods and following an environmental restructuring. The Lant and Mezias (1992) paper examines how patterns of organizational change are affected by environmental conditions, levels of ambiguity, organizational size, search rules (adaptive, imitative, and garbage can), and organizational performance. They demonstrate that a learning model can account, to a large degree, for a punctuated equilibrium pattern of organizational change. The Mezias and Lant (forthcoming) paper explores the role of organizational level change in population dynamics. Specifically, they examine the conditions under which a significant proportion of firms using an imitation strategy will survive in populations of organizations consisting of imitating firms, which incur costs of search and change, and fixed firms, which do not incur costs of search and change.

The discussion of these papers will illustrate how computer simulations can be a flexible tool for developing theories of complex organizational phenomena. Simulations of organizational processes can yield compelling research implications, raise important research questions, and aid in theory development; these are issues of concern to all organizational theorists. I hope to demonstrate that organizational research based on computer simulations can be made accessible and relevant to a wide audience of organizational theories.

References

Cohen, K.J. and R.M. Cyert (1965), Simulation of organizational behavior, in J.G. March (ed.) *Handbook of Organizations*, 305-354, Chicago: Rand McNally.

Crecine, J.P. (1969), *A Computer Simulation of Municipal Budgeting*, Chicago: Rand McNally.

Lant, T.K. and S.J. Mezias, Managing Discontinuous change: A simulation study of organizational learning and entrepreneurial strategies, *Strategic Management Journal* (1990) 11: 147-179.

Lant, T.K. and S.J. Mezias (1992) An organizational learning model of convergence and reorientation. *Organization Science*, 3, 47-71.

Mezias, S.J. and T.K. Lant (forthcoming). Mimetic learning and the evolution of organizational populations. In J.A.C. Baum and J.V. Singh (eds.) *Evolutionary Dynamics of Organizations*. Oxford University Press.

Characteristics of Markets as Organizations comprising Nonrational Individuals

by
Dhananjay K. Gode and Shyam Sunder

We study simple double auction markets in which artificial nonrational individuals trade. These artificial traders are simple computer programs that randomly pick their actions from a large set prespecified by the rules of the market. Since these traders have no power to observe, remember, seek profits, or to maximize, we label them "zero intelligence" (ZI) traders. Rules of the market, applied to the actions picked by individual traders, determine the market outcomes.

It is know from past work that when human traders, given their private information alone, and having little idea of the market equilibrium price, trade in such markets they converge to the close neighborhood of the theoretical equilibrium predicted by the supply and demand functions as measured by price, quantity allocations, and efficiency of these markets. However, these market supply and demand functions can be constructed only by someone who has access to the private information of all individuals. Since there is no such omniscient individual in these markets, the empirical result has remained a source of some surprise to many observers; given the rules of a double auction, it has not been possible to prove that the market will converge to equilibrium.

When we replace human traders by the ZI traders, without placing any additional constraints on their behavior, markets do not converge to theoretical equilibrium. However, imposition of a budget constraint that restricts the action choices of the ZI traders to the set in which they can live within their means, is found to be sufficient to cause these markets to converge to equilibrium. While the approach to equilibrium price is noisier than in human markets, it occurs nonetheless with budget constrained ZI traders. Furthermore, efficiency of the market (as measured by the aggregate profits of all traders) is virtually the same as the efficiency of markets with human traders--which is close to hundred percent.

These markets can be seen as an example of organizations in which non-rational individual behavior yields rational aggregate behavior. The empirical described above have led to mathematical analysis that shows that the convergence to equilibrium and high levels of efficiency is a characteristic that attaches to the design of the organization (in this case the rules of double auction market), it is not a characteristic that derives from the behavior of individuals in these markets. There are, however, other characteristics of the markets outcomes (such as the across-trader dispersion of individual profits) that are determined by individual behavior as well.

Organization is often viewed as an arena in which individuals play. It has not always been possible to isolate those performance characteristics of organizations that derive from their design from those that derive from the behavior of individuals in them. The zero intelligence concept may be generalizable beyond markets to other forms of organizations, and holds some promise to serve as a means to isolate these two effects.

Evolving Novel Oganizational Structures

by
Kevin Crowston

Organizations take many structural forms. A key problem in organization theory is to explain the distribution of these forms. Population ecologists explain them by analogy to the diversity of biological species: organizations compete against each other and organizations with structures more appropriate for the environment are more successful and therefore take over. Organization with less effective structures are unable to compete and eventually disappear. Thus we observe a match between the structure of the organization and the environment in which it operates.

A second goal of organization theory is to suggest new organizational structure. I am particularly interested in organizational design because information technology seems to be changing some of the fundamental constrainsts on organizations, e.g., by making communication and data storage much cheaper and allowing interaction across time or space. I would like to look for new organizational structures that might be more appropriate to the changing conditions. However, as stated population ecology theories only say that such a new structure will eventual replace the existing structures; they provide no insight into what the new structure might be.

I therefore propose to explore the space of possible organizational structures by simulating the evolution of organizations on a computer. This technique (called genetic algorithms) has been used to evolve many kinds of objects, e.g., strategies for bargining games. Most of these studies involved fairly simply objects, but the technique has recently been applied to objects with complex structures, e.g., lisp programs, suggesting that it might be applicable to organizations as well. In this paper, I will describe the key elements of such an evolutionary strategy and outline options for implementing such a strategy.

Only Diamonds are forever: Caching Knowledge in Episotic Classification

by
Benn R. Konsynski

The talk will deal with formal models for representing the transient activities that are prevalent in organizations. Issues in delegationtechnologies, learning, intelligent agents and cognitive reapportionment will be discussed. the central theme of representing the dynamics of organizational entities will be explained in each of these dimensions.

Graph Theoretical Measures of Informal Organizations

by
David Krackhardt

Four measures of the degree of hierarchical structure are developed using graph theory: Graph Connectivity, Graph Hierarcy, Graph Efficiency, and Least-Upper-Boundedness (LUB). Each measure, ranging in continuous values from 0 to 1, corresponds to a necessary and sufficient degree to which one of the four necessary and sufficient conditions for the existence of a "pure hierearchy" (known as an outtree in graph theory). In addition, each measure captures a different social aspect of informal organizations. Random graphs are used to demonstrate the sensitivity of each measure to graph size and density.

HITOP-A: Reasoning About Organization Design

by
Les Gasser

The introduction of new production and enterprise integration technologies will change organizational arrangements within and across firms. The kinds of organizational changes entailed go far beyond hardware, network, and interoperability concerns, and involve new job designs, new coordination structures, new performance management systems, new skill and knowledge requirements, and others. Integrating knowledge from these multiple disciplines in the design and implementation of organizations with such embedded technical systems can be difficult to achieve for a variety of reasons. Moreover, organizational arrangements and production or integration technologies interact; the need for new organizational arrangements such as new coordination structures can change the requirements for and loads on production and enterprise integration technologies themselves. The basic question underlying our research is how to design computer-based production and enterprise-integration technologies as open systems where the people, organization, and hardware/software components are designed simultaneously and early on. Academically, this question translates into an attempt to use computer-based analysis and modeling techniques to create a practical theory of sociotechnical systems design which involves identifying how technology, people, organization, and environment can be matched for a more effective total system. Practically, this question translates into the development of a knowledge-based system---HITOP-A---for supporting the design, planning, and implementation of social and organizational infrastructures simultaneously with technical (e.g., product and process) design. The HITOP-A prototype is running, and we report on our experiences using it to generate theory and to design organizations.

Validating the Organizational Consultant on the Fly

by
Helmy H. Baligh, Richard M. Burton and Borge Obel

The Organizational Consultant is a knowledge base expert system to help in the design of organizations. New knowledge is generated continually through experiments, case studies and applications. A strategy of learning by doing is developed. The goal is an improving expert system which is validated on the fly. The interplay between system improvement and basic theory are complementary. The Turing test is argued to be generally (ir) relevant for a managerial decision support. For management the goal is to help so that their knowledge for choice is enhanced.

Multi-Agent Interaction Processes: From Oligopoly Theory to Decentralized Artificial Intelligence

by
Raimo Hamalainen

This paper examines the group processes studied in oligopoly theory and in decentralized artificial intelligence. We develop a unifying perspective for the research on the behavior of autonomous interacting agents. Among the many questions of interest in these disciplines are the ways of creating and reaching cooperation by a group of self-interested independent decentralized agents. In this respect, the models and results of oligopoly theory can also be used both in decentralized artificial intelligence and in many other areas of research, such as group decision making, negotiation support, and organizational theory. Especially, the important idea of reshaping goals with strategic information sharing and transmission--incentive communication--has received little attention outside the field of economics. On the other hand, oligopoly theory and experimental economics, in particular, can benefit from the computational methods and tools of artificial intelligence and modern decision support technology. To demonstrate this we have built a prototype of an experimental market analysis environment. Its potential in the analysis of group processes is illustrated with examples.

Keywords: artificial intelligence, game theory, oligopoly theory, group decision making, negotiation support, organization theory

Prof. Linda Argote
Graduate School of Industrial Administration
Carnegie Mellon University
Pittsburgh, Pa 15213

Prof. Richard Burton
Fuqua School of Business
Duke University
Durham, NC 27706

Prof. Kathleen Carley
Carnegie Mellon University
Department of Social & Decision Sciences
Pittsburgh, PA 15213

Prof. Kevin Crowston
Graduate School of Business Administration
University of Michigan
Ann Arbor, MI 48109-1234

Prof. Les Gasser
Computational Organization Design
Lab Institute for Safety & Systems Management.
University of Southern California
Los Angeles, CA 90089

Prof. David Krackhardt
School of Urban and Public Affairs
Carenegie Mellon University
Pittsburgh, Pa 15213

Prof. Raimo Hämäläinen
Systems Analysis Laboratory
Helsinki University of Technology
Otakaari 1 M, SF-02150 Espoo
FINLAND

Prof. Bernardo Huberman
Xerox Palo Alto Research Center
333 Coyote Hill Road
Palo Alto, CA 94304

Prof. David Kleinman
The School of Engineering
Dept. of Elec. & Sys. Engr.
University of CT
Room 312, UK-157
260 Glenbrook Road
Storrs, CT 06268

Prof. Benn Konsynski
Management Information Systems
Harvard Business School
Boston, MA 02163

Prof. Panagiotis Kouvelis
Fuqua School of Business
Duke University
Durham, NC 27706

Prof. Eduardo Salas
Naval Training Systems Center
12350 Research Parkway
Orlando, FL 32728-3224

Prof. Theresa Lant
Dept. of Mgmt/Organ. Behavior
Stern School of Business
90 Trinity Place
New York, NY 10006

Prof. Arie Lewin
Fuqua School of Business
Duke University
Durham, NC 27706

Prof. Michael Masuch
Universitati van Amsterdam
Oude Tursmarkt 151
1012 GC Amsterdam
HOLLAND

Prof. Michael J. Prietula
Carnegie Mellon University
Graduate School of Industrial Admin.
Pittsburgh, PA 15213

Prof. Shyam Sunder
Graduate School of Industrial
Administration
Carnegie Mellon University
Pittsburgh, PA 15213

Prof. William A. Wallace
Center for Industrial Innovation
RPI
Troy, NY 12180-3590
or
Polyproject "Risk & Safety of
Technological Systems"
Swiss Federal Institute of Technology
ETH-Center
8092 Zurich
SWITZERLAND

Prof. Massimo Warglien
Universita Degli
Studi di Venezia
ITALY