Policy Evolution within an Organization

James H. Hines
Sloan School of Business, Massachusetts Institute of Technology, Cambridge, MA 02139

Jody Lee House
Department of Electrical and Computer Engineering, Oregon Graduate Institute, Beaverton, OR, 97006

Abstract: A plan of action is given for a newly funded research project on organizational evolution. In this study, our goals are to apply an evolutionary framework to organizational learning. The process includes a collaborative effort with partner companies to identify the working mechanisms behind the evolution of policies or decision rules. We also use computer simulations as a tool to examine our findings.

Introduction: This document describes a plan of action for a three-phase research project funded under the Innovation and Organizational Change program by the National Science Foundation in organizational learning using an evolutionary framework**.** The long-term objective of our ongoing research is to explicate principles that govern evolution within an organization. On a more practical side, we hope to provide managers with a set of rules or guidelines that will permit their companies to evolve more rapidly in desirable directions. These rules might deal with the appropriate number of business units, promotion policies, or team-based decision making, among other organizational characteristics [1].

Our first-phase goal is to perform *evolutionary audits*– that is, to collect data on the evolutionary potential of several partner companies. The evolutionary audits will aid in the development of a model of organizational evolution (phase II). The data we collect and the model that we build will form a foundation for in-depth explorations of the question: How can managers create an organization that will evolve quickly in the direction they desire? (Phase III) [2].

Partner companies will be involved in all three phases. Our partners include PriceWaterhouse-Coopers, Pugh Roberts Associates, Eastman Chemicals, Hewlett Packard, The Lincoln School and General Motors. We selected these particular companies for their interest in this work, their existing expertise in one of the underlying methodologies (system dynamics), and certain organizational characteristics (discussed below) which make them particularly good subjects for studying organizational evolution.

Justification: Two observations suggest this effort is worthwhile: First, we are addressing a significant problem, and, second, the problem is solvable.

The problem is that system-wide company improvement is difficult because companies are too complex to “solve” [3]. Creating a consistent set of beneficial management policies is difficult or impossible because the complexity of modern companies exceeds, by many orders of magnitude, our ability to understand. Managers today work on isolated issues that they can identify. Sometimes the issues are solved, but often this reactionary approach leads to unintended consequences, as intended solutions create problems in other areas of the organization [4].

The problem of improving organizations in the face of ignorance is solvable. In fact, it *has* been solved, just not by humans. Biological evolution has produced excellent natural organizations (i.e. organisms) even though the organizations themselves are completely ignorant of how they are put together or why they succeed. We are identifying analogs of natural evolution that will help companies to likewise excel.[1]

In this paper, the foundations of the study, the technical methodology, and our current status are presented. Our research plan follows this background material. A schedule follows the research plan.

Foundation: Central to our work is a particular analogy between biological and organizational evolution. Analogies can be dangerous when careless application leads to unwarranted transfer of conclusions from one domain to another. Properly employed, however, analogies are powerful mechanisms for using precious knowledge from one area to bootstrap understanding in another [5]. Today, evolution is increasingly seen as a general mechanism, not restricted to the biological realm [6-8][2]. Still, the evolutionary mechanisms we understand best are biological. Hence, in seeking to understand how an organization can evolve it is natural to seek fruitful and powerful analogies between the biological and the organizational.

A focal point of our research has been an analogy between organizational policies and biological genes. By *policy* we mean an explicit or implicit decision rule in the usual system dynamics sense [10]. For example, a manager might set prices by the implicit rule: Raise prices when inventories are low, and lower prices when inventories are high. The policy is *implicit* as long as it remains unspoken or unwritten. Articulating the policy, perhaps by recording it in a policy manual, could make it more explicit. Of course, people might change their approach to pricing even without updating the manual. In this case, the new approach would be a policy, while the old procedure recorded in the manual would no longer be a policy in our use of the term. A policy is a rule or procedure that people actually use to make a series of decisions. In this case, the policy gives rise to a continuing stream of particular decisions to raise or lower price.

A policy in an organization is comparable to a gene in a cell. A gene is a segment of DNA (or, in some organisms, RNA) that acts as a set of instructions for the ongoing production of a particular protein. The proteins then catalyze reactions in the cell. Indeed, no necessary chemical reaction occurs in a cell without a protein catalyst that is coded by a gene [11]. Genes produce a continuing stream of action in the cell, while policies produce a continuing stream of action in the company.

The creative mechanisms in evolution are mutation and recombination. In genetic mutation, part of a DNA molecule is physically changed, producing a new gene. In our analogy, genetic mutation corresponds to

[1] The difficulty of making organization-wide improvements clearly increases the value of an evolutionary approach. However, this does not mean that *evolutionary* management is mutually exclusive of other improvement efforts. Indeed, good evolutionary mechanisms will tend to *spread any beneficial change* whether the change is intentional or not. The same evolutionary mechanisms will *limit the spread of deleterious changes*, again whether or not the changes are intentional.

[2] Interestingly, Darwin himself was lead to his theory of evolution partly through an analogy between Adam Smith's political economy and biology. And, Adam Smith had been influenced by an analogy between Newtonian physics and political economy [9].

policy change, intentional or unintentional, and (like mutation) producing either favorable or unfavorable results [12]. The result of such a change, for better or worse, is a new policy. Genetic *recombination* occurs when two DNA molecules mix to form a new DNA molecule. In a company, genetic recombination corresponds to a particular kind of organizational learning: Inter-personal learning whereby a person combines a part of someone else's ideas with his or her own [5, 13].

Evolutionary management consists primarily of managing the environment in which policy change and learning (mutation and recombination) occur. One important task is to create mechanisms that will ensure the spread of more effective policies and the decay of less effective ones. That is, managers need to create mechanisms that correspond to nature's processes for the spread and selection of genes.

In higher animals, sexual reproduction encourages the dissemination and recombination of genetic material. Natural selection is the process by which beneficial recombinants (and mutations) are retained, while deleterious ones are discarded. Sex and selection in the natural world correspond in the corporate world to the various ways in which companies identify certain employees as exemplary and encourage other employees to learn from (or imitate) them [14-19]. We call these processes of identification and encouragement "pointing and pushing mechanisms". For example, pay and organizational position are two ways that a company can *point* to outstanding performers. Pay and position can serve a *pushing* function as well: Employees are motivated ("pushed") to learn via their desire to rise in the pay scale and in the hierarchy.[3] Other pointing and pushing mechanisms are also possible.

[3] *Pushing* successful people to share policies with others may also be necessary [20].

Methodology: We have developed theory and preliminary formal models of organizational evolution to deepen our understanding and to gain additional insight. Our approach to modeling has been to combine system dynamics modeling [10, 21], agent based modeling [22, 23] and genetic algorithms [24-26].

System dynamics models are built around non-linear feedback processes. The model formulations are similar to those found in feedback and control models of electronic, mechanical, and biological systems, but are used in system dynamics usually (though not exclusively) to investigate human-related systems. The system dynamics methodology has been applied to a broad diversity of topics in organizational dynamics including project management, inventory supply chain management, environmental systems management, and urban planning [10, 27-31]. Here, we use a system dynamics model to simulate the progress of projects within a company, given a set of policies.

Agent based models are built around simulated agents –or, in our case, *managers.* Each agent has a repertoire of behavior and can interact with other agents. The key behavior of our agents (managers) is making decisions based on policies and the key interaction is learning from another manager and thereby changing the policy.

Genetic algorithms are computer programs that solve problems by mimicking the biological process of evolution [24, 25]. A standard genetic algorithm would begin with a population of potential solutions to a problem and would evolve ever-fitter solutions through processes of mutation, recombination, and selection. Genetic algorithms have been applied to a wide range of optimization problems including microwave antenna design, circuit design and airline scheduling [32-35]. In most applications of this technique, the genetic algorithm is simply used as an optimization technique and is *not* intended to represent a real-world process operating within the problem domain. In our work, however, we interpret the genetic algorithm as representing a particular type of human learning – the process by which one person incorporates ideas from another into his or her own policies.

For example, perhaps a manager (agent) follows the previously mentioned policy of raising prices when inventories are low and lowering prices when inventories are high. In this way his price effectively responds to excess supply (high inventories) or demand (low inventories). Perhaps he comes into contact with another manager with a different policy, a manager who sets prices by taking costs and adding a fixed margin of two hundred percent. The second manager's policy ensures that prices never fall below costs. From this contact, our original manager might learn the idea of margin pricing. Henceforth he might price at a margin over cost, like the second manager, but in addition, he might vary the margin depending on inventory position, similar to his original policy. The new policy is actually a combination of two parent policies, combined by learning. In this example, the new policy might perform better than either original policy, because it responds to inventory position (supply) and at the same time ensures that price will never fall below cost[4].

Generally, we use system dynamics, agent-based modeling and genetic algorithms in the following way: A system dynamics model represents the underlying physics of the organization as well as all policies that are *not* evolving. *Evolving policies* are carried by agents (individual managers) who learn from one another via a process that is essentially a genetic algorithm.

The completed simulation environment will allow us to investigate the conditions and trade-offs that influence the rate at which organizations improve their policies. Conditions include number of teams, team size, frequency of mixing and evaluating, promotion (or other pointing and pushing) policies, number of managers, complexity of evaluation criteria, complexity of policies, number of policies, and many others. An example of a trade-off is between the ability to discriminate between good and bad policies on the one hand and, on the other hand, the speed by which a policy spreads. As the number of teams increase, the rate by which policy changes disseminate will slow, while the ability to discriminate between beneficial and deleterious changes will increase.

One important contribution of our methodological approach will be a tested method for building simulation models that will enable managers to investigate the efficiency of proposed or existing evolutionary mechanisms in their own companies.

Current status: We have initiated this study by developing a preliminary theory of organizational evolution and creating a proof-of-concept model to explore and deepen our theory of organizational evolution as well as to demonstrate the feasibility of combining system dynamics with agent based and genetic modeling. Throughout, we have sought examples from the business sector to support, clarify, or debunk our findings.

The proof-of-concept computer simulation environment coordinates two modules: a system dynamics model and an agent-based

[4] Of course the recombination can be detrimental as well [36].

module operating under a genetic algorithm. More specifically, we have created a simple system dynamics model of a company that simultaneously runs a number of projects (shown in Figure 1). We interpret the equations in terms of a software company or automobile manufacturer, which issues a stream of new releases for each of several product lines – for example, word processor, spread sheet, etc, or compact car, sedan, sport utility vehicle. In constructing this model we are able to build upon a rich tradition of project modeling in system dynamics [29, 30, 37-39].

Each manager (agent) is assigned to a team managing one of the projects (i.e. a release for one of the software products). Each management team determines certain policies for the underlying system dynamics model. In our current simulation environment, the success of each project is evaluated at the end of each project cycle. The managers receive promotions and demotions according to the relative success of the projects on which they worked. The managers are then mixed and reassigned to new teams, in which each manager has an opportunity to learn from a team member. Learning occurs via a process of policy recombination where the probability of learning *from a particular colleague* increases with that colleague's relative position in the management hierarchy (essentially, this is a genetic algorithm). The system dynamics model then simulates each project based on its managing team's policy. Performance of each team is evaluated; promotions and demotions are handed out; and the cycle begins anew.

This initial simulation environment has allowed us a preliminary examination of the impacts of team vs. individual evaluation, number of teams, and team size on evolutionary efficiency. For example: Early results suggest that, when done correctly, team-evaluation performs almost as well as individual-evaluation in terms of being able to "point" to people who perform well.[5] Further, as the number of teams increase, the likelihood of converging on an optimum policy increases, while the speed of convergence decreases [14, 17]. Finally, Team size appears to have at most only a small impact on organizational evolution.

Simulator specifics: Our proof-of-concept simulation environment has been developed as an object-oriented program [40-42]. A schematic of the class structure is shown in Figure 2.

The simulations are parameterized using a control panel as shown in Figure 3. The user can establish the number of teams, the learning profiles, and the range of policy values that can be established (chromLength). The number of teams impacts the number of generations required to see policy convergence, and thus the user can control the number and duration of the generations. Altering the start time and the integration time step (DT) controls the system dynamics simulation. Increasing the time step decreases the resolution of the model behavior. Finally, the user can choose the probability of learning (recombination) and the probability of innovations (mutations).

We are also able to change the learning process by running different types of company profiles. The individuals in the company can work in teams of varying sizes or can work alone. Performance in the company can also be rewarded by different promotion schemes. The company profiles are coded into the simulator.

[5] Teams must be randomly reformed periodically for this to work. Essentially, the time series of how an individual's teams have performed provides enough resolution to estimate that person's average contribution.

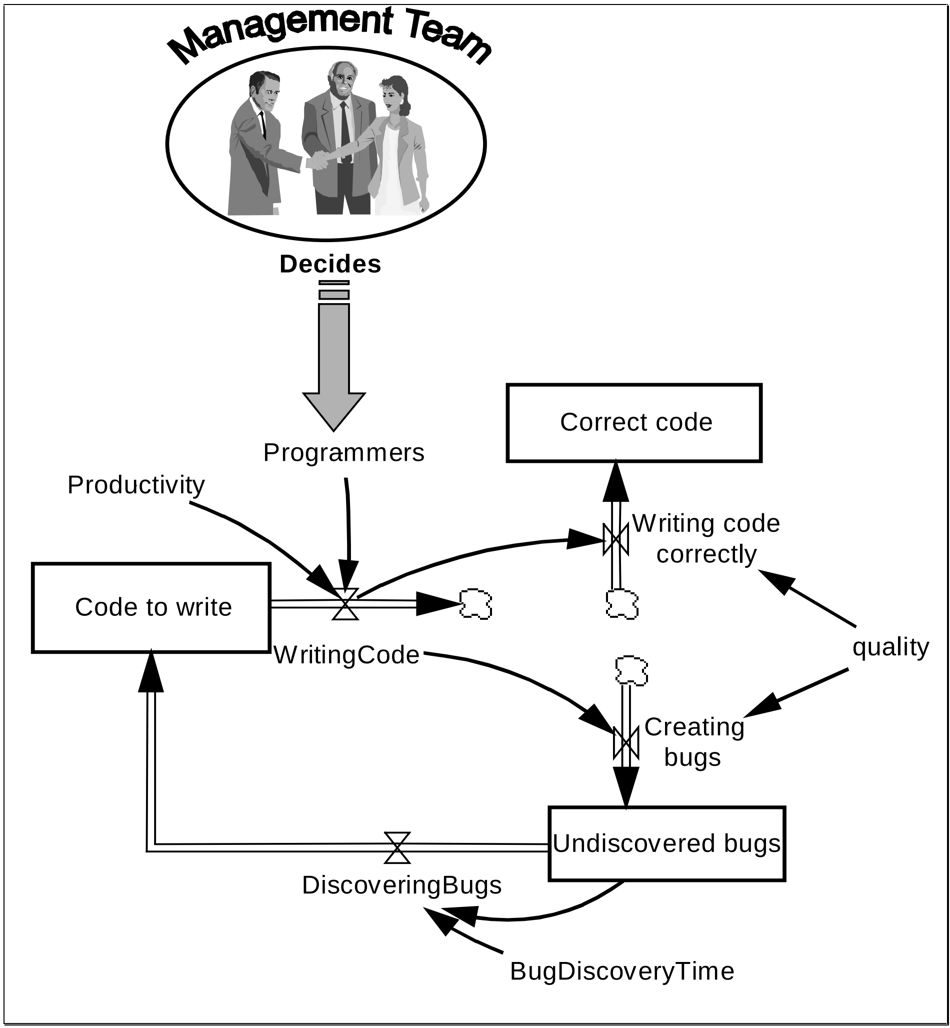

Figure 1: Schematic of project model under control of agents

Model output is either shown graphically or collected in tabular form. Figure 4 shows the model output relating to the policy values (in this case desired number of programmers) of the different managers on the Y-Axis plotted as a function of time on the X-Axis. The policy of each manager is determined at the end of every generation. As a result, the curves tend to show large jumps in values before the managers converge on a good policy [43-45].

Research Plan: The research plan is divided into three areas: evolutionary audits of partner companies, technical model development, and an in-depth evolution-centered examination of an important organizational issue, such as team learning.

[6] Teams must be randomly reformed periodically for this to work. Essentially, the time series of how an individual's teams have performed provides enough resolution to estimate that person's average contribution.

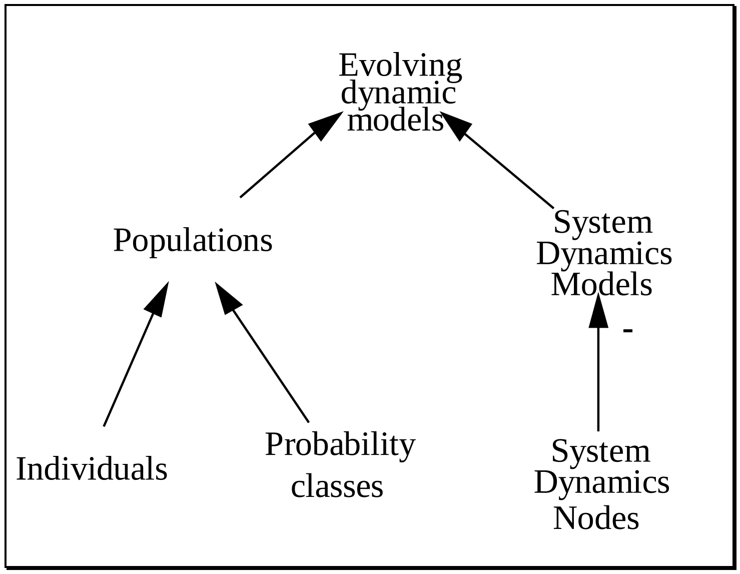

Figure 2: The simulation environment is an object-oriented program consisting of about 50 classes. Conceptually, an evolutionary model is composed of two parts: A system dynamics company model and a population of individuals who manage the company.

The results of our initial work have suggested that we need to work with members of industry to better identify and understand existing or potential evolutionary mechanisms within organizations, refine, extend and document our model, and then pursue further analysis and applications [46]. Accordingly, we have asked a number of companies and one school to work with us on this research. An advisory committee comprised of lead contacts from these organizations will be formed to provide advice, ideas, critiques, and direction. The organizations themselves will be research sites.

Partner Companies: Evolutionary theory suggests the characteristics of organizations that are likely to provide particularly fruitful research sites. First, evolution requires consistent "selection" criteria, suggesting that we work with companies where success is particularly easy to see. Second, evolution works on populations; so the easier it is to identify an organizational population, the easier it will be for us to work with a particular company.

An obviously prime class of companies to work with are those that have a population of projects whose success or failure is relatively public. Such companies include software firms, accounting firms, and consulting firms. Our partners from this class of firm are: PriceWaterhouse Coopers, the largest professional service firm in the world and Pugh Roberts Associates, a U.S. subsidiary of PA Consulting, Britain's largest consulting firm.

Schools with their "populations" of classrooms also fall into this organizational category. Consequently, our group of partners also includes the Lincoln School, an independent k-12 girls school in Providence, RI.

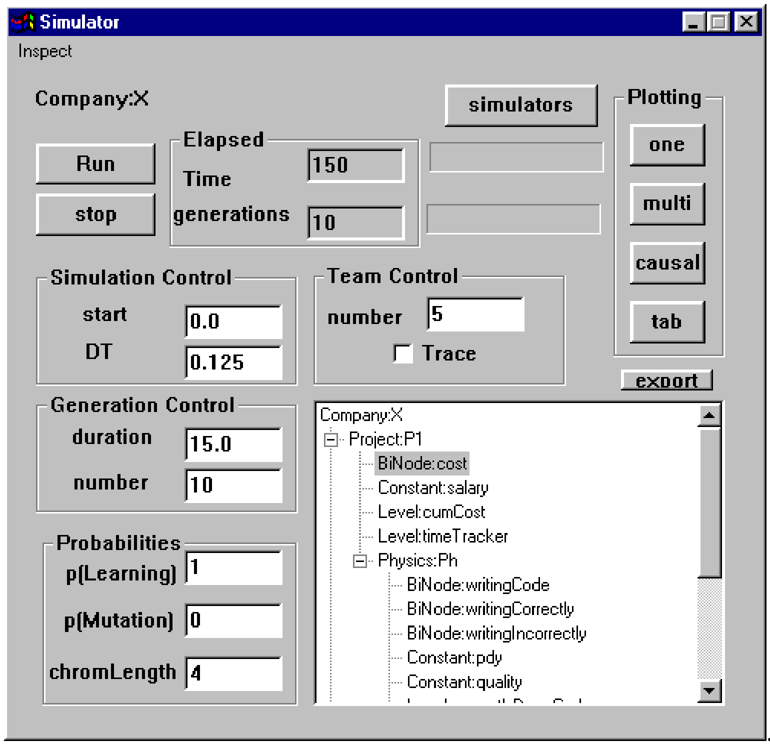

Figure 3: The control screen permits the user to choose a simulator, run a simulation, call for output, and to specify various parameters including number of generations, mutation rate, time step, and maximum generation duration.

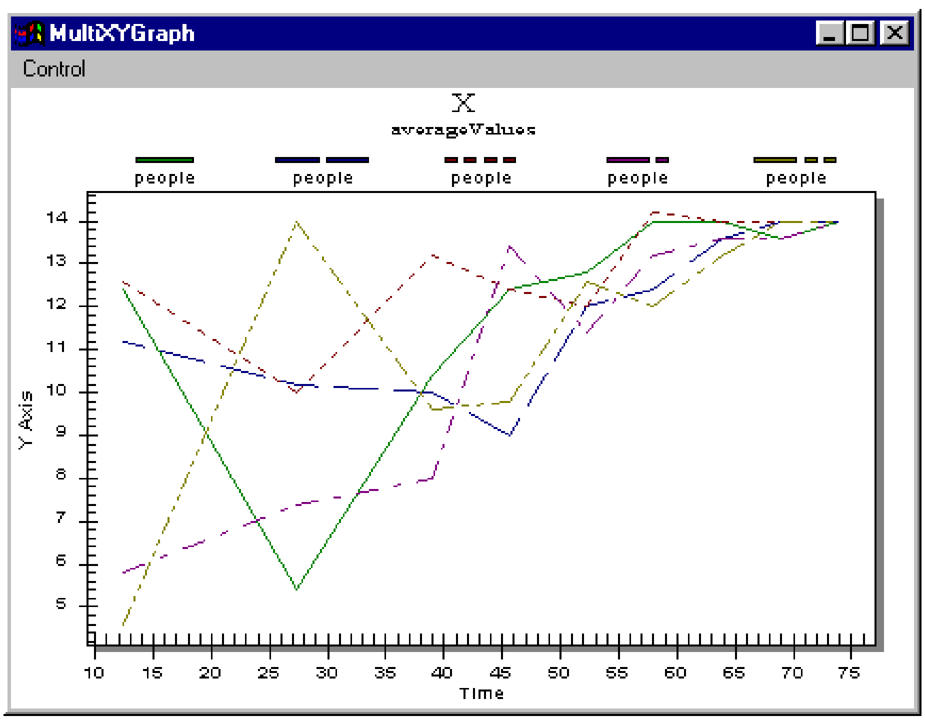

Figure 4: Graphical output from a simulation of a software company showing managers of different projects converging on a policy.

Although project-based firms such as these are important in their own right, we do not believe that the application of our findings will be limited to such companies. Here we apply the biological example -- where experiments with mice, flies and yeast yield insights that hold beyond these test organisms-- to organizational research [47]. To increase confidence that our findings have broad applicability we will also work with two additional companies whose operations, while including project based work, also have significant aspects that are not neatly described by project models (e.g. manufacturing and assembly). The three companies in this class are Eastman Chemical Company, a recent winner of the Baldrige Award, Hewlett Packard, and General Motors.

Phase 1. Evolutionary Audits: In the initial phase of the proposed research, we will audit the evolutionary potential of each of our partner companies. We will investigate the amount of policy-innovation, the amount of inter-personal policy learning, and the existence and effectiveness of selection mechanisms. In addition we will gather retrospective recollections about the evolution of certain company policies.
Data collection will focus on policy change and variability, policy recombination, and Pointing and Pushing mechanisms within the organization. In addition to collecting data on the current state of the subject companies, we will also attempt to chart the evolution of particular policies via recollections of managers.

Data collection will involve open-ended interviews with managers in our partner companies. There will be two classes of questions, (a) questions about specific evolutionary mechanisms and (b) questions about specific policies. Mechanism-specific questions will center on evolutionary processes (e.g. promotion criteria) and will aim at identifying actual examples. Our policy-specific questions reverse this process: We will consider one or more specific policies and probe managers about how the policy came into being, how it gives rise to decisions, and how it changes. We will map these specifics onto the variables of our theory.

Questions in the interviews will be guided by theory. Mechanism-specific questions will be guided by our developing theory on organizational evolution. Policy-specific questions will be guided by theory embedded in pre-existing system dynamics models. In the case of project-based partners, we will consider policies suggested by the project model, discussed earlier in this proposal. Policies treated in this model include:

- Hiring policies
- Completion policies (e.g. when do we "ship" a new product)
- Resource allocation policies
- Scheduling policies
- Work release policies
- Testing (QA) policies

In the case of our industrial and educational partners we will consider policies treated in models that the organizations themselves have built for specific issues. For example in the case of Eastman Chemicals we could look at:

- Pricing policies
- Capacity expansion policies
- Repatriation policies
- R&D policies

The specific policy we focus on for each company will be determined in consultation with our advisory committee.

We anticipate that during our study we will sharpen our techniques and protocols for knowledge elicitation, based on results of our efforts. The basic approach, however, will be one that we have used successfully for about a decade with similar subjects in a wide variety of companies. This approach entails repeated separate or combined interviews with two to three subjects. During the course of the interviews, we build a representation of a mental map, and reflect it back to the subjects for extension, modification, and ultimately confirmation [46].

By the end of this first phase we will have comparative information on the evolutionary mechanisms that are currently at work in organizations. Further, we will gain a practical understanding of constructs that we currently know primarily from theory. Finally, the data we collect will provide the foundation advancing our modeling work, which is the focus of the second phase of our research.

Phase II. Modeling extensions and modifications: In the second phase we will modify and extend our proof-of-concept model using the new information from our audits. Although the particular changes and additions we make will depend on the results of our audits, we anticipate changes in a number of areas.

The initial simulator portrays managers evolving a single policy – the number of programmers to put on a software release. Obviously, a typical manager employs a number of key policies, as our audits are virtually certain to show. These policies need to be coordinated. Closely related to the need for coordination of policies held by a single manager is the need for coordination of policies held by different managers, perhaps managers in different parts of the firm. Coordinating the policies of different managers is a potentially important aspect of organizational learning rather different from what we have considered heretofore (inter-personal learning). Evolution will require policy coordination at different scales within a company (i.e. within a person and between people)

Biological evolution has met similar challenges at microscopic scales (where, for example, one enzyme works hand in glove with another), at macroscopic scales (where for example upright posture coordinates with the peculiar human foot) and at “megascopic” (ecological) scales (where for example hawks and doves regulate each other [17] and where cleaner wrasses (Labroides dimidiatus) are adapted to pick parasites from barracuda who themselves are adapted *not* to eat the wrasses).

We believe that *organizational* evolution can provide similarly effective *policy* coordination. Indeed, policy coordination is of extreme importance to our work because a major source of the failure of organizational innovation stems from conflicts with existing policies or activities [4]. It is important that we deal with how different policies – whether held by a single person or by different people -- can evolve together so that they complement one another. Consequently, we will likely extend the model in two ways. First, we will alter the model to represent each manager as having policies governing multiple decision points in the underlying system dynamics model; say, permissible overtime, and willingness to slip the schedule in addition to the number of programmers. Second we will probably want to extend our model to allow team membership to evolve in order that complementary policies held by different managers may regularly be associated with one another [48].

Another extension to the model will be to introduce *multiple* criteria for evaluating and promoting managers. Currently, managers’ promotions (part of the pointing and pushing device) are based solely on how quickly they ship a release. Its clear, however, that real world managers are evaluated on a number of criteria, and we expect our audits to give us some sense of how many criteria a superior works with in evaluating his people. We will almost certainly extend the model to represent multiple criteria, say remaining bugs and cost as well as time to complete.

It is likely that our audits will reveal that a genetic algorithm is not the ideal way to represent an evolving policy. Genetic algorithms usually work with strings –such as a string of bits representing a number (e.g. the number of programmers). Our audits may reveal that we need a more flexible technology, most likely genetic programming [49]. In genetic programming computer programs or phrases in a mathematical equation evolve, rather than strings. Hence, by moving to genetic programming, we can simulate the evolution of any policy that can be represented in any way on a computer. Put another way, genetic programming will permit us to simulate changes not only of parameters, but changes of the structure of a policy as well.[7] While structural change can be represented as parameter changes, it will be easier and more faithful to the underlying reality to have structure evolve directly.

[7] A parameter is a constant, while structure relates to operations and inputs. For example say the pricing policy is:

$$price_t = cost_t * spread$$

where *price* and *cost* are variables and *spread* is a constant (say, 1.5). A genetic *algorithm* could be used to evolve a better spread (perhaps 2.8 rather than 1.5). A genetic *program* could also evolve a better spread, but in addition it could transform the policy into something different, for example:

$$price_t = \mathrm{MAX}\left[cost_t * \mathrm{f}\left(inventory_t\right), minimumPrice\right]$$ in which the *spread* is now a non-linear function of *inventory* and in which *price* is prevented from going below a minimum value.

Generations of learning are discrete in our simulator currently – occurring at the end of each release – even though the release itself is simulated in continuous time. We would like to permit learning to occur at any time that people are interacting, as we anticipate the audits will show is common.

Finally, we would like to enhance the user interface to a sufficient degree to permit our industry partners to play with the simulator. The company representatives are the sorts of people who will likely want hands-on involvement. Their involvement will be useful in a number of ways – it will make the experience richer for them, and it will give them a much more solid foundation for contributing to and critiquing the work.

Phase III. In-Depth Example in Evolutionary Context: The goal in the final phase of our proposed research program is to show that the evolutionary framework provides an approach for dealing with real, significant organizational problems. With our corporate and educational partners, we will take an in-depth look at a specific, significant organizational issue. We will wait to choose a topic until we are closer to the third phase. However, the topic will flow out of the first phases of the investigation, and will be coupled to the interests of the advisory committee. Based on our work to date, we can give several examples of the kind of issue that we might consider in the third phase of our project.

1. Team learning and promotion: We have explored simulations of organizational evolution where employees are evaluated individually and where only their teams are evaluated. Surprisingly team evaluation worked almost as well as individual evaluation. Currently, teams and team learning are issues of great interest among both

managers and researchers [50]. We may be in a position to help companies get even more value from these teams than they anticipated, for example by *evolving* good team-practices [51, 52].

2. Company size: Evolution works on populations. We simulate populations of business units and managers. The number of business units and the number of managers affect evolution's speed and ability to discriminate between degrees of fitness. New traits can spread rapidly through small populations, however large populations are better at discriminating between good and better policies. We are currently in a period where downsizing and "sticking to the knitting" seems to be giving way to mergers. We may want to consider what the evolutionary pros and cons are of size and diversity [53].

3. Selection mechanisms: Evolution requires a selection mechanism. One selection mechanism we have represented is managerial promotions. We anticipate that we will learn of several other potential mechanisms in the first phase, and may be able to model them in the second phase. Much has recently been written about the desirability of a flat organization. One possible area of investigation in the third phase will be an assessment of the evolutionary costs of flattening the hierarchy and a concrete understanding of other pointing and pushing mechanisms that could replace hierarchy.

Summary: We hope to develop a better understanding of organizational learning from an evolutionary framework during the course of the three phase study that has been outlined in this paper. The long-term objective of our ongoing research is to explicate principles that govern evolution within an organization. On a more practical side, we hope to provide managers with a set of rules or guidelines that will permit their companies to evolve more rapidly in desirable directions.

Acknowledgements: This research is supported by the National Science Foundation, under the Innovation and Organizational Change Program - Award # 9975942, and our business partners; Eastman Chemicals, PriceWaterhouse Coopers, Hewlett Packard, Pugh Roberts Associates, The Lincoln School, and General Motors.

References

[1.] Lounamaa, P. and J.G. March, *Adaptive Coordination of a Learning Team.* Management Science, 1987. **33**(1): p. 107-123.

[2.] Hines, J.H., *Evolutionary Management: Five Rules for Organizational Improvement.* Systems Thinker, 1998. **9**(6): p. 1-3.

[3.] Simon, H.A., *Adminstrative Behavior*. 1976, New York: The Free Press.

[4.] Kofman, F., J.D. Sterman, and N.P. Repenning, *Unanticipated Side Effects of Successful Quality Programs: Exploring a Paradox of Organizational Improvement.* Management Science, 1997. **43**(4): p. 503-521.

[5.] Holland, J.H., *et al.*, *Induction: Processes of Inference, Learning, and Discovery*. 1986, Cambridge, MA: MIT Press.

[6.] Collins, J.C. and J.I. Porras, *Build to Last*. 1994, New York: Harper Collins.

[7.] Dawkins, R., *The Selfish Gene*. 1990, Oxford: University of Oxford Press.

[8.] Holland, J.H., *Hidden Order: How Adaptation Builds Complexity*. 1995, Reading, MA: Addison-Wesley.

[9.] Depew, D.J. and B.H. Weber, *Darwinism Evolving*. 1995, Cambridge, MA: MIT Press.

[10.] Forrester, J.W., *Industrial Dynamics*. 1961, Portland, OR: Productivity Press.

[11.] Steele, E.J., R.A. Lindley, and R.V. Blanden, *Lamarck's Signature: How Retrogenes are Changing Darwin's Natural Selection Paradigm*. 1998, Reading, MA: Perseus Books.

[12.] Meszias, S.J. and M.A. Glynn, *The three faces of corporate renewal: Institution, revolution, and evolution.* Strategic Management Journal, 1993. **14**(2): p. 77-101.

[13.] Miner, A.S. and S.J. Mezias, *Ugly Duckling No More: Pasts and Futures of Organizational Learning Research.* Organization Science, 1996. **7**(1): p. 88-99.

[14.] Nei, M., *Molecular Evolutionary Genetics*. 1987, New York: Columbia University Press.

[15.] Gillespie, J.H., *The Causes of Molecular Evolution*. 1991, New York: Oxford University Press.

[16.] Leach, D.R.F., *Genetic Recombination*. 1996, Cambridge, MA: Blackwell Science, Ltd.

[17.] Smith, J.M., *Evolutionary Genetics*. 1989, New York: Oxford University Press.

[18.] Brown, T.A., *Genetics: A Molecular Approach*. 1992, New York: Chapman & Hall.

[19.] Brown, T.A., *Gene Cloning: An Introduction*. 3 ed. 1995, New York: Chapman & Hall.

[20.] Constant, D., L. Sproull, and S. Kiesler, *The kindness of strangers: The usefulness of electronic weak ties for technical advice.* Organizational Science, 1996. **7**(2): p. 119-135.

[21.] Richardson, G. and A. Pugh, *Introduction to System Dynamics with Dynamo.* 1981, Portland, OR: Productivity Press.

[22.] Resnick, M., *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds*. 1994, Cambridge, MA: MIT Press.

[23.] Epstein, J.M. and R. Axtell, *Growing Artificial Societies: Social Science from the Bottom Up*. 1996, Cambridge, MA: MIT Press.

[24.] Goldberg, D.E., *Genetic Algorithms in Search, Optimization and Machine Learning.* 1989, Reading, MA: Addison Wesley Publishing Company, Inc.

[25.] Holland, J.H., *Adaptation in Natural and Artificial Systems*. 2 ed. 1992, Cambridge, MA: MIT Press.

[26.] Herriott, S.R., D. Levinthal, and J.G. March, *Learning from experience in organizations.* American Economic Review, 1985. **75**(2): p. 298-302.

[27.] Forrester, J., *Urban Dynamics*. 1969, Portland, OR: Productivity Press.
[28.] Forrester, J., *World Dynamics*. 1971, Portland, OR: Productivity Press.
[29.] Cooper, K., *Naval Ship Production: a Claim Settled and a Framework Built.* Interface, 1980. **10**(6).
30.] Cooper, K.G., *The Rework Cycle: Vital Insights Into Managing Projects.* IEE Engineering Management Review, 1993. **21**(3): p. 4-12.
[31.] Meadows, D. and D. Meadows, *Toward Global Equilibrium*. 1973, Portland, OR: Productivity Press.
[32.] Majhi, A.K., L.M. Patnaik, and S. Raman, *A genetic algorithm-based circuit partitioner for MCMs.* Microprocessing & Microprogramming, 1995. **41**(1): p. 83-96.
[33.] Eclercy, D., A. Reineix, and B. Jecko, *FDTD genetic algorithm for antenna optimization.* Microwave and optical technology, 1997. **16**(2): p. 72.
[34.] Langerman, J.J. and E.M. Ehlers, *Agent-based airline scheduling.* Computers & Industrial Engineering, 1997. **v33n3,4**: p. p. 849-852.
[35.] Koonce, D.A., C.-H. Fang, and S.-C. Tsai, *A data mining tool for learning from manufacturing systems.* Computers & Industrial Engineering, 1997. **v33n1,2**: p. p. 27-30.
[36.] Champaigne, A.B., L.E. Klopfer, and J.H. Anderson, *Factors influencing the learning of classical mechanics.* American Journal of Physics, 1980. **48**: p. 1074-1079.
[37.] Abdel-Hamid, T. and S. Madnick, *Software Project Dynamics: An Integrated Approach*. 1991, Englewood Cliffs, NJ: Prentice Hall.
[38.] Ford, D.N. and J.D. Sterman, *Expert knowledge elicitation to improve mental and formal models.* System Dynamics Review, 1997. **in print**.
[39.] Ford, D.N. and J.D. Sterman, *Dynamic modeling of product development processes.* System Dynamics Review, 1998. **14**(4): p. 309-340.
[40.] Booch, G., *Object Oriented Design with Applications*. 1991, New York: Benjamin/Cumming Publishing Company, Inc.
[41.] Wirfs-Brock, R., B. Wilkerson, and L. Wiener, *Designing Object Oriented Software*. 1990, Englewood Cliffs, NJ: Prentice Hall.
[42.] Pinson, L.J. and R.S. Wiener, *An Introduction to Object-Oriented Programming and Smalltalk*. 1988, Reading, MA: Addison-Wesley Publishing Company.
[43.] Hines, J. *Organizational Evolution*. in *International System Dynamics Conference*. 1996. Cambridge, MA.
[44.] Hines, J. *Organizational Evolution*. in *New England Complex Systems International Conference*. 1997. Nashua, NH.
[45.] Hines, J. and J. House. *Team promotion to foster policy evolution*. in *New England International Complex Systems Conference*. 1998. Nashua, NH.
[46.] Carley, K.M., *Extracting team mental models through textual analysis.* Journal of Organizational Behavior, 1997. **18**: p. 533-558.
[47.] Fine, C.H., *Clock speed: Winning industry control in the age of temporary advantage*. 1998, Reading, MA: Perseus Books.
[48.] Eisenhardt, K.M., J.L. Kahwajy, and L.G. Bourgeois, *Conflice and strategic choice: How top management teams disagree.* California Management Review, 1997. **39**(2): p. 42-62.
[49.] Koza, J.R., *Genetic Programming*. 1992, Cambridge, MA: MIT Press.

[50.] Lant, T.K. and A.B. Eisner, *Pharmaceutical R&D in an era of managed healthcare: Using integrative teams to produce enduring competitive advantage.* International Journal of Technology Management, 1998. **15**(3,4,5): p. 299-321.
[51.] Eisenhardt, K.M., J.L. Kahwajy, and L.J. Bourgeois, *How management teams can have a good fight.* Harvard Business Review, 1997. **75**(4): p. 77-85.
[52.] Ancona, D.G. and D.E. Caldwell, *Cross-functional teams: blessing or curse for new product development?*, in *Transforming organizations*, T.A. Kochan and M. Useem, Editors. 1992, Oxford University Press: New York. p. 154-166.
[53.] Miller, K., *et al.*, *The evolution of strategic simpliciy: Exploring two models of organizational adaptation.* Journal of Management, 1996. **22**(6): p. 863-887.