TOWARDS A THEORETICAL MODEL FOR ADAPTIVE ENTREPRENEURIAL ORGANISATIONS USING GENETIC ALGORITHMS

by

D. F. MUZYKA*
and
A. J. DE KONING**

96/85/ENT

* IAF Professor of Entrepreneurship and Associate Dean for the MBA programme, at INSEAD, Boulevard de Constance, Fontainebleau 77305 Cedex, France.

** PhD Candidate, at INSEAD, Boulevard de Constance, Fontainebleau 77305 Cedex, France.

Printed at INSEAD, Fontainebleau, France.

TOWARDS A THEORETICAL MODEL FOR ADAPTIVE ENTREPRENEURIAL ORGANISATIONS USING GENETIC ALGORITHMS

Daniel F. Muzyka
Alice J. de Koning

ABSTRACT

This paper proposes the use of genetic algorithm modelling techniques to develop a rigorous theory for adaptive entrepreneurial organisations. In particular, we focus on the need for more rigorous theory of adaptation to sort out the comprehensive action lists that have been generated by years of field research by many people. After describing how a basic genetic algorithm model works, and what assumptions drive the model, we develop a series of propositions which specifically link the GA models to earlier findings on entrepreneurial organisations. The paper is a beginning in the process of building a theoretical model.

INTRODUCTION

Despite intensive research in many disciplines, it is often difficult to separate what is consequential or associated behaviour from the key elements of adaptation in organisations. Research the field of entrepreneurship, and in other fields looking at how we build more robust economic organisations, have been helpful in raising our awareness of some of the elements that may be appropriate to an organisation "built to last" (to the extent that this is a title of a recent book on the subject) . We have described the need for entrepreneurship; the need to focus the organisation on opportunity and its capture. We have described mechanisms to permit more independent entrepreneurs to co-exist within the "normal" structures and processes inherent in economic organisations. We have highlighted the importance of various elements in making the organisation more opportunity focused and adaptable, including: learning and knowledge systems, different modes of resource allocation, the ability to manage change, and the strategic importance of building skills and competencies. In the end, however, we often have comprehensive lists. What we still lack is a cohesive, internally consistent model to describe the overall adaptation process in organisations in terms of structure, systems, and process. In particular, what we still lack is a short-list of key causal elements for adaptive behaviour.

Our research objective is to attempt to better understand the relevant and key elements of adaptation in an organisational system as suggested by the mapping of genetic algorithm models developed in artificial intelligence. We are attempting to open the discussion as to whether these models can be of use to us in better describing the essential elements required for organisational adaptation. Even though these models do not provide deterministic solutions to the problem of adaptation, they can be valuable in providing us a more rigorous, integrated structure for describing both the adaptation process and the essential organisational elements. They may be also be useful to us in describing the

organisational and managerial behaviours required to profit from our current environment, in the same way that the models of decision sciences were in the 1950s and 60s were useful to the Carnegie School researchers in describing managerial decision-making in their operating environment.

The basic thesis we are asserting is that these models are valuable and provide a valuable analog for what is required to ensure organisational adaptation. Even more than merely an analog, detailed modelling of adaptive organisations as genetic algorithm models will allow us to test the logic and rigour of our theories of organisational adaptation. We will begin by testing the appropriateness of genetic algorithm models in this article by applying evidence from long-run field research dedicated to better understanding Adaptive Entrepreneurial Organisations (AEOs). In our earlier papers with Churchill (e.g. 1994, 1995), we described Entrepreneurial Organisations as those that through the regular pursuit, evaluation and capture of new opportunity are economically viable over the long run. To be specific, however, we are not here attempting to describe the opportunity seeking and capture processes through application and discussion of the AI models . Rather, we are hoping to develop more insight into other major elements - the managerial levers - in creating self-renewing organisations: how do we get the organisation to efficiently and effectively form around the new opportunity and therefore redeploy the assets and focus on the organisation?

We do not discuss the need for adaptive organisations in this paper: discussions of the increasing chaos and unpredictability of our world abound, and along with arguments for the need for organisations to respond efficiently to changes and even to create changes to gain competitive advantage. The accelerated technological change, which continues to challenge the ability of organisations to respond (e.g., the average product life-cycle of computer products has dropped from 3 years in the 1980s to less than one year), is well documented (e.g. Tushman, 1986). Story (1993) has highlighted broader reasons for many environmental perturbations: one, a very high percentage of the individuals who have ever lived are alive today and two, political and economic relationships and barriers have been massively realigned in the past five years. The "bottom line" for business is that the economic life cycle of products is decreasing, competitive positions are less stable and that established competitors are often realizing defeat at the hands of newcomers.

We have chosen to focus on our own approach to adaptive organisations, as developed in our work on EOs. Other models of entrepreneurship in organisations exist, as we recognise, each one concerned with how management ensures that the organisation realizes more opportunity and becomes more adaptive to environmental and technological change. Specifically with entrepreneurship research, research streams in corporate venturing (e.g., Block and MacMillan, 1993; Burgelman, 1983), the building and support of internal entrepreneurs (e.g., Pinchot, 1985), and those working on the interface between product development and entrepreneurship in the area of innovation (e.g., Van de Ven et al., 1989). Our interest in pervasive organisation-wide systems leads us to continue focusing on EOs, which we have renamed Adaptive Entrepreneurial Organisations (AEOs) to highlight our specific interest in corporations which are self-renewing, re-organising around specific new opportunities as the business evolves.

USING ARTIFICIAL INTELLIGENCE TO MODEL KNOWLEDGE SYSTEMS

Attempts to replicate human and organisational thinking on a machine, particularly the inferential processes of humans, have produced some interesting insights and models which we believe may be useful in studying how to build more adaptive organisations. At the heart of the research in entrepreneurship, learning, change, etc., at the heart of opportunity and adaptation, is the ability to build new inferences or possibilities from existing knowledge (de Koning & Muzyka, 1996). Renewing the organisation, as all research streams cited earlier will validate, is more a knowledge game than a resource or structural game. The question in both self-renewing organisations and artificial intelligence research is how to get systems to improve themselves. This is echoed in Holland, et al. (1986):

> Researchers in artificial intelligence are able to build sophisticated game players and powerful expert systems, but have been largely stymied in attempts to program these systems to improve themselves. Indeed, alteration of these systems by simple insertion of new knowledge typically produces chaos, which can be undone only by extensive reprogramming.

Similarly, in management and economic research, the performance of the modern corporation in the last 20 years would suggest that we built "sophisticated" and "powerful" systems, but that these systems are not good at processing and realising the benefits of new knowledge without "extensive reprogramming" or transformation and reengineering. We would argue that we should at least attempt to learn from the models and work in artificial intelligence, particularly in the area of building self-renewing, or adaptive models.

Within the field of artificial intelligence, the primary stream of research into models which can not only deal with new knowledge but also create knowledge is related to a class of models referred to as genetic algorithms (see Holland, 1975) for an early definition of such models). The work in this field, which has led to some definition of "frameworks" rather than a complete theory of induction (see Holland et al., 1986). However, the artificial life models developed from this research have had great success in modelling natural evolution (see Koza, 1992). We propose to explore the linkage between these frameworks and organisational adaptation.

One could argue that genetic algorithms are just one class of adaptive and machine learning paradigms. Therefore, why not explore the others? The other models include neural networks, simple hill climbing, induction of decision trees, and simulated annealing. Each of these models promises to be adaptive (i.e., involves changing of some structure so that it performs better in its environment). However, the evidence suggests that the genetic algorithms models and the evolutionary development they produce are more robust and efficient models that the others, applicable in a wider variety of circumstances. This is a difficult determination to make in a field (i.e., AI) which is in some ways still in its infancy or growth phase, but the weight of evidence (see the journals in the field such as Complex Systems and Adaptive Behaviour) suggests that genetic algorithms are most generally applicable, especially in an environment that can be truly described as "chaotic".

What indications do we have that these models may be appropriate to today's unpredictable and fast-changing environment? The answer is that we have several. These models can deal with multiple and evolving goals and objectives. There is no assumption or requirement for an optimal solution or linearity. All one has to be able to indicate is whether the outcome is relatively "good" or "bad", at the moment. The model does not require that we know everything that is happening: it permits real inference. It makes no assumptions about environmental stability or the nature of environmental evolution: discrete steps and/or changes in relationships among those things observed is tolerated. In summary, genetic algorithms, as a class of models, would appear to ideally model the relationship between an organisation and the current environment.

RESEARCH METHODOLOGY

Research Questions

To fulfill these research needs, we specified several research questions:

Q1. What are the assumptions, structure, and processes that make up genetic algorithms?

Q2. Are these models a useful analog to organisations? Do they yield clear and appropriate implications or propositions which may be validated by empirical research in entrepreneurial organisations?

Q3. What are the implications of genetic algorithms for organisational systems which would like to be better able to adapt around evolving opportunities?

As part of the research project, in this paper we will be working with genetic algorithm models as currently described in the AI literature, and attempting to map the models into the context of economic organisations. We will develop a set of propositions which we will explore based on data we already have of successful AEOs. Should the propositions not find some grounding in our understanding of successful AEOs, one could then certainly challenge the propositions and the applicability of the genetic algorithms in this context.

Database

For purposes of this work, two sets of data we used are most important. The first is the current research database on genetic algorithms. This is available from a variety of researchers in the field of artificial intelligence, particularly Holland et al. (1986) and Koza (1992). We have conducted a relatively thorough investigation of their work as part of building our propositions in this paper.

The second is an understanding of entrepreneurial corporations and those attempting to become EOs. As part of our larger research program designed to better understand the management mechanisms in companies that have displayed a long-term entrepreneurial orientation (see Muzyka, De Koning and Churchill, 1995) and to understand how to make

existing, more traditional corporations more entrepreneurial, we have developed a large database of information which we can draw upon. We have carried out research which is described in cases, interview data and programmes developed in both larger and smaller corporations. The larger organisations in which research has been conducted include 3M, Digital Equipment Corporation, Johnson and Johnson, ABB, Rhone-Poulenc, and Siemens-Nixdorf. The "middle" companies or wholly owned divisions of larger organisations that have been investigated include Trinkhaus and Burkhardt (a private bank), Mettler-Toledo (weigh scale company, part of Ciba-Geigy), Festo Tool (power tool company), MHZ (window treatments), and Kolb and Schule (textiles and clothing). Some compilations of the data have been published (Muzyka et al., 1995; Churchill et al., 1996), and other data has been published in its reduced form (e.g., 3M Europe case in EMJ).

Caveats

In the end, this paper produces an argument that specific models developed from another discipline can be applied to help us understand the phenomenon of organisational adaptation. This paper represents a first attempt at building this argument and we will be continuing to develop the argument in future papers. We are, in this paper, linking observations from field research with insights into the nature of an evolving class of models to determine whether a cohesive argument can be made regarding the insights this class of models may provide us. In particular, we are in the course of building a computer model of a multi-business corporation which we hope will operationalise this class of models in an effort to build more adaptive organisations. We do not purport to demonstrate a conclusive research finding, but are attempting to make the aforementioned link as part of the early work on the subject. We are therefore open to suggestions and criticism of all sorts.

INTRODUCTION TO GENETIC ALGORITHMS

A genetic algorithm model is essentially a population of "genes" interacting with an environment, trying to achieve reasonably good performance vis à vis that environment. The models consist of a tournament of possible rules which can act on an environment and in being selected and performing well are rewarded so that they are more likely to be called upon again. Good performance is the effectiveness of the cluster of genes to match specified, and possibly changing goals. The environment is also set, and possibly changed, by the programmer. Within the context of specific hypotheses, researchers typically focus on the "behaviour" of the cluster of genes, looking at both the speed at which the cluster converges on local optimal performance, and how quickly the cluster reconverges after a shift in the environmental conditions or goals. Other researchers use genetic algorithm models to find solutions to complex problems more efficiently, and thus look for the best performing "gene" of the population of all the generations.

The Rule or Genetic String

The genetic algorithm model works on a cluster of rules (or genetic strings) which interact with each other and the environment. The rules have three segments, condition (i.e., what triggers the rule), message (for the future), and action, plus a strength metre which 'remembers' the use and usefulness of the rule in the system in previous periods. Each segment consists of a number of positions, each of which is filled by a 1 (yes), 0 (no) or # (doesn't matter or not clear). Messages from the environment are structured in a similar way.

The Three Basic Operational Cycles

The primary cycle of the model, which we will call the operations cycle, is the interaction between the cluster of rules and the environment, and is usually repeated thousands of times. In addition, at regular intervals the selection cycles are run. The selection cycles are the attribution cycle, which resets the strength metres, and adaptation cycle, which replaces the weakest rules with new ones. The selection cycles occur about every 20 or 30 cycles, to allow time for trying out the new and obviously weakest rules. This balance between the operations and selection cycles allows for learning, adjusting and improving the cumulative performance of the cluster by changing and selecting the best rules.

The operations cycle of the genetic algorithm goes from environmental sensing, to selecting and using rules, and back to the environment. Information from the internal environment is sensed, and a sample of these messages are combined with information from memory (i.e. "posted" messages from the previous period) (see Figure 1). The information, like the condition of the rules, contains 1, 0 and #, which allows for ambiguity and incomplete information. Using this information, the program checks for possible relevant rules from the global rule set, simply looking for matches of the messages to the condition section of the rules. This matching process creates a set of possible rules. The set of possible rules is then filter in a two step process. First, the rule is given a bid for action, which essential gives it a number of lottery tickets based on specificity, i.e. how well the rule matches one or more messages, strength, i.e. how often the rule has been used in the past, and support, i.e. the number of messages activating the rule. Then the filter process generates one or more random numbers, and the rule(s) with the winning numbers are used. Using the rule means that the message section of the rule is posted to the next period(s), and the action section of the rule acts on the environment.

The selection cycles, which we have called the attribution and adaptation cycles, are the key to the overall adaptivity of the cluster of genes. The attribution cycles award the fitness of the gene with strength metre points, and the adaptation cycles replace a number of weak genes with new creations. Both the selection cycles are often run at the same time, but they can be run on independent schedules. For example, some programmers prefer to run the attribution cycle continuously, along with the operations cycle. The program usually fixes the timing of the selection cycles at regular intervals, after a specified number of operations cycles. Recent research has shown that running selection cycles asynchronously has unexpected results. We refer to the timing of the selection cycles as 'cycle time', where the unit of measurement is a single operations cycle.

The attribution cycle, which recognises and credits valuable rules, can use a number of different methods. The essential point is that the system measures actual performance against desired performance, and based on that awards strength to rules used. The strength metre on the rule, then, is reset by the accumulator which sums all the awards given to a rule since the previous attribution cycle was run. The accumulator combines rewards from positive outcomes (successful actions on the environment), the number of times the rule was triggered by messages and incoming information, and the number of rules it in turn triggered. The system remembers linkages between rules, and from action to performance. It should be said, however, that the system does not need to hold this memory centrally - the rules of a chain of action usually award back down the chain directly and individually (even before getting a performance reward itself), and so memory is decentralised. Technically, the accumulator can run periodically or continuously (which would affect the bidding process in the operations cycle), but usually the strength metre is updated at regular intervals.

The adaptation cycle, which replaces weak rules with new ones, also can uses a number of different methods. The first step is rank ordering the rules in the global rule set, according to their strength metres. The weakest are then eliminated. Note that the global rule set is always a limited, pre-set number, and that the rule for elimination is also usually set by the model builder - typically a set number of the lowest ranking rules. These spots are then filled by new rules which are created through various types of adaptation mechanisms. The primary means of adaptation are simply reproduction of the strong rules, and cross-over of strong rules. Reproduction is an exact copy of the strong rule, while cross-over trades complete sections of C, M or A between two strong rules to create new ones. Secondary means of adaptation include mutation, which is a random change of one or more position in the rule, inversion, which reorders all the positions with a single section of the rule (condition, message or action sections), permutation, which reorders some positions within a single section, and more rarely editing, decimation, and encapsulation. At the extreme, random mutation creates entirely fresh rules, which probably will not perform well. Note that the strong rules are used as a part of much of the adaptation.

The Environment and Performance Goals

The desired performance is based on the stated (and changeable) goals of the system, which are specified by the programmer. The performance goals may be stated as optimization goals, or as cybernetic goals, and may also be quite complex. For example, a simple model of a pipeline may specify performance goals in terms of changing demands for the product, as well as numerous safety requirements such as pressure, flow, temperature, etc. Environment is also controlled by the programmer, so that it reflects the specific problem being modelled. In our case, for example, we would model the environment to reflect various business environments, from stable, predictable markets to fast changing, unpredictable ones.

GENETIC ALGORITHMS AND THE ADAPTIVE ORGANISATION: COMMON GROUNDS

Underlying Assumptions of GA Models

Why do genetic algorithms help us understand AEOs? In this section, we will address this point by looking at the assumptions that underlie GA models..

The key advantage of genetic algorithm models is that they are inherently non-deterministic. The assumptions of deterministic models (which more typically are used in business research) as summarised by Koza (1992, 1994) and Holland et al. (1986), describe a world which is radically different from the logic and reality of an evolving organisation or world, where the future is not always predictable nor knowable. Genetic algorithms are based on a different set of assumptions, which we would argue are closer to the reality of business. As with traditional models or organisations, GA models work better when the future is an extension of the past, but they do adapt in the face of major shifts. It is worthwhile to follow the Koza and Holland discussion, and look at the assumptions driving deterministic models more closely. In this context, we then can consider some of the assumptions or "non-assumptions" inherent in genetic algorithm models.

Correctness. They note that deterministic models assume that the solution must be correct. In contrast, GA models are "nearly correct" in the sense of being in the right neighbourhood, finding the local optimum, but rarely finding the optimal solution. Like EOs, GA models move away from the assumption of optimizing performance in the given environment, and move towards realisation of performance (and the environment!). Or, to use Stevenson's (1985) words, the world is not optimal, it is created: the actions of entrepreneurs and other individual actors who jointly, through their actions, determine the course of the future or the potential outcomes.

Consistency. Deterministic models assume the consistency, where GA models thrive on inconsistency through diversity, contradiction and paradox. We observed this type of inconsistency in many EOs with their multiple and even contradicting decision making centres and resource allocation processes (Muzyka et al., 1995; Churchill et al., 1995).

Justifiability. Deterministic model assume justifiability, beginning from premises and using logical inference to find solutions. GA models take many "leaps of faith" in generating mutations, and experimenting with them in real time. In parallel, EOs have a bias to simply trying out new products and processes, rather than engage in excessive processes focused on justifying action.

Causal certainty. Deterministic models operate in a universe of causal certainty, while GA models are non-deterministic, and therefore have a lack of clarity in causal links. Interestingly, the links between rules (who is causing whom to act) are carefully maintained in a GA model. The question of whether the rule "causes" performance is less important. High performance is rewarded, of course, but if the rule's action did not cause but was merely correlated with high performance, than poor performance in later periods will correct attribution. Certainty is not required because hypothesized causal links are checked and updated regularly.

Orderliness. Deterministic models assume orderliness, proceeding in a tightly controlled and synchronized way. In GA models, Koza (1995, p. 6) observes:

> "It is unsettling to think about numerous uncoordinated, independent, and distributed processes operation asynchronously and in parallel without central supervision. Untidiness and disorderliness are central features..."

GA models use the disorderliness of creating and maintaining diversity, to ensure adaptability. The EOs we observed also tried to achieve a controlled measure of diversity.

Parsimony. Deterministic modelling is designed to achieve parsimony, using as simple a set of variable and causal arguments as possible. GAs are more concerned with resource-adequate solutions that fit the problem, and therefore are only parsimonious if the environment is resource poor.

Decisive. Finally, deterministic models are decisive, in the sense that the problem solving process eventually reaches an obvious termination point, when a clear result is achieved. GA models, on the other hand, "go on and on" (Koza, 1995, p. 6). We observe and analyse the models at various stages of evolution, but that does not imply we have reached the end.

The assumptions driving GA models appear to be much more realistic for all economic organisations, especially at this time Business is a continuing, open-ended process of interacting with markets, competitors, technology, etc., with a continuing need to achieve satisfactory performance. In particular, we find the GA model even more appropriate to understand the decentralised, messy, and chaotic world of an adaptive entrepreneurial organisation. In addition to the assumptions above, we need to highlight an important assumption which is crucial to the effective evolving of the GA model.

Clarity and availability of information. The GA models do make a reasonably strong assumption about clarity and availability of information, which in fact we found was also true in EOs - often through free observation from any point in the organisation. One important result of this (non-centralised) clarity is that attribution or recognition of contribution is maintained. Clarity and availability of information does not include completeness of information. GA models, likes EOs, sample the relevant environment, and impute the rest. In this context, we define organisational learning as the modification and reinforcement of the organisation's knowledge base.

Additional Observations

Before developing our propositions, the broad parallels between the GA model and the Adaptive entrepreneurial organisation should be sketched out. It is possible to go into detail, discussing parallels between each element of the model and the EO. Out of respect for space constraints, however, we here focus on only the major elements, preferring to spend more time on developing propositions.

Our need to interpret the model in terms of the adaptive entrepreneurial organisation made us consider the intuitive advantages of considering the GA rules as people-routines within the organisation. Three main arguments can be developed to support this assertion. First, we recognise that people originate new ideas, and remember and enact old routines. A person with an excellent, high-performance routine may, in fact, be the right person to experiment with a new "mutation". Second, limiting the parallel of GA rule to an individual suggests that people can't change if their work is not adequate, and are therefore eliminated. We would prefer to suggest that people imitate successful routines, or learn to improve their own routines. It is routines that should be eliminated, not necessarily people. Third, limiting the parallel of each GA rule to organisational routines suggests that specific routines in an organisation are consciously recognised and given credit for performance, which in fact is unlikely. Most routines, especially new untried ideas based on largely tacit knowledge, are attached to people or groups who eventually get credit for the performance. The process of attributing recognition for performance is an important ingredient in the overall effectiveness of the GA model, and therefore a people orientation should drive our explication of the analogy between organisations and the GA model.

The term "routine" comes from the Carnegie School (e.g. Cyert and March, 1963) and population ecology (esp. highlighted by Nelson & Winter, 1982). The Carnegie School observed that most action in organisations was not explicitly thought-out responses to stimuli, but rather routines were enacted semi-automatically. Because these routines were not explicit, and often individuals enacted parts of the routine without being aware of the rationale or link of their part to others, this makes the routines difficult if not impossible to change. Early population ecology went further, arguing that the organisations could not learn, and therefore were victims of the quality of their routines. Both these observations are valid for many organisations - the change management literature has hosts of examples of companies who have difficulty changing effectively. In the context of adaptive entrepreneurial organisations, we extend these perspectives on organisation routine by linking routines to people and the external environment. Thus we identify both the learning agent - a person - and the motivation for changing behaviour - performance rewards and punishments from the environment. In our propositions we show how the adaptive entrepreneurial organisation is able to leverage learning and incentives effectively.

The rules or people routines engage in two types of cycles, the operations and selection cycles. The operations cycles are parallel to the daily production and information flow of an organisation, which people engage in as they get their work done. The selection cycles are the processes of giving rewards and process of reflecting, learning, and creating. As noted above, the rewards refer to both direct performance rewards, and also to linkage rewards which recognise the relationships between people and routines. The rewards can be done continuously, or at longer intervals. The more interesting implications and underlying assumptions of these processes are developed below.

IMPLICATIONS OF GENETIC ALGORITHMS FOR ECONOMIC ORGANISATIONS

The adaptive system as described by a genetic algorithm model depends on a number of key behaviours and structural characteristics of the system. We have developed these points into propositions which can be tested in adaptive entrepreneurial organisations (AEOs). To illustrate each proposition, we have selected examples for our previously reported and on-going field research into adaptive entrepreneurial organisations, and identified other research which is relevant. In some cases, the propositions derived from the GA model confirm or illuminate results reported by other researchers, and in other cases new areas of inquiry are introduced.

Key Behaviours of the Adaptive Entrepreneurial Organisation

Mindfulness. The processes of sensing from the environment, remembering messages from previous periods, and matching these to routines, all suggest that a high degree of attentiveness is required. This attentiveness is directed both to the environment and to the organisation internally. An organisation cannot absorb and react to all information, so some basis for selection, for directing the attention of people to specific elements in the environment or organisation is required. Normally, the cognitive frames of employees is not managed, and so over time the blinders of routine limit both the type and the amount of information sensed by the organisation, often simply confirming assumptions about the state of the world and avoiding possible contradiction (e.g. Porac et al., 1989).

Proposition 1: EO leadership actively manages the cognitive frames of employees, usually towards identifying potential business opportunities.

Research by Fuller indicates that heightening people's mindfulness before exposing them to information categories actually increases the amount of information they process, as well as helping them identify the relevant information. The problem within adaptive organisations is that managing cognitive frames might suggest a rigid, non-adaptive framework for filtering information. On the contrary, work on learning in the audit environment has shown that even the wrong' framework can help novice auditors structure their experience, and they move to an accurate and efficient understanding of audit procedures more quickly than those who received no prompting.

Christensen (1987) has studied several projects within organisations designed to increase overall entrepreneurial activity. One specific technique called for bi-weekly meetings throughout the company, with the simple agenda of customer frustrations.' Not only do these meetings provide a structure for discussing the customers' orientations, it also regularly re-enforces the mindfulness' of every employee towards finding and resolving customer frustrations.

Attribution. The importance of attributing performance to specific people or routines cannot be overstated. The attribution of effective performance both outside the organisation ("the environment") and within the organisation (the linkages between people and routines) influences the selection of specific routines for action and drives the adaptation processes by identifying good performance and eliminating low performers. Learning or adaptation is goal directed by the attribution process, because the strength of the rules or routines is a measure of how their actual performance matches the desired

performance. Thus, the contribution of routines and chains of routines to performance must be recognised clearly. In the typical GA model, the linkages between routines are rewarded by the routines that are activated by the linkage, but those activating routines are not directly rewarded by the external performance. GA models do not require some centralised overseer to recognise all the links of performance to rules, and the linkages between rules. To summarise, we need correct, accurate, and complete attribution of credit, but not centralised. For the system to work, there must be a clear, local understanding of the performance goals, as well as local measurements of actual performance.

It would be tempting to follow the common pattern of simplifying performance rewards to financial incentives, but that is impractical given other aspects of the attribution system. Attribution of external performance to specific routines occurs at a system level, which we can compare to the centralised reward systems of most companies. When rewarding linkages between people and routines, a more decentralised attribution is needed because of the limitations of centralised information processing. The practical necessity of direct attribution between routines suggest that attribution must operate primarily as a public recognition of contribution, and not directly as financial incentives. Clearly, the culture in such an organisation must value people who share credit for successes generously. We might expect a long run correlation between recognition for individual contribution and the financial rewards, but not in the short run. We also can argue that excessive financial rewards for external performance (i.e. winner-take-all rewards) can distort people's incentives to share credit, and thus lead to a breakdown in the attribution system.

Proposition 2: Adaptive entrepreneurial organisations provide clear, accurate and complete recognition for performance to employees.

An example from 3M may illuminate the value of the attribution process for an adaptive entrepreneurial organisation. During our research we observed that people in 3M have a remarkable memory of "what and who". If asked about some action or information, most people will identify the person who initiated it. The value of this memory is that potential responses to future problems are not 'lost'. The detail of why and how the decision was arrived is retained by the actors themselves, so remembering who did or said what becomes a crucial organisation memory tool for remembering potentially useful information. This approach to remembering routines limits the excessive inflexibility of responses written up as standard operating procedures and divorced from their creators, and also allows for future access to the underlying logic of the routine and subtle tacit knowledge.

Most of the other companies observed have taken steps to ensure that the question of "who" actually delivers. ABB and Siemens-Nixdorf have designed their control system to ensure not only transparency but attribution. Trinkhaus and Burkhardt (the bank) have designed their internal control and monitoring system to ensure they know who is delivering and who is not so that they may directly reward them. Each of these organisations, as shall be seen shortly, are looking for performance, but they move beyond

this to ensure that attribution is achieved. In the end, attribution appears in the model and in practice to be a key to adaptation.

The Innovation Spectrum: Recombining to Splicing to Mutating. GA models provide a useful paradigm for understanding the fundamental nature of innovation within the context of a history of action. Within a GA model, the degree of innovativeness in the adaptation is deliberately varied, along a continuum of virtual replication of high performance routines to random generation of new routines. From the perspective of the adaptive entrepreneurial organisation, the scale covers incremental improvement and copying high performance routines into new domains (recombination), to more creative combination of high and medium performance rules with some random elements (splicing), to the completely unusual output of the independent eccentric (mutation). The adaptations hold high performance and diversity in tension -- very appropriate to the business environment, where profits must be achieved if the organisation is to survive, and yet where flexibility to meet changing environments is increasingly important. In reflecting on this scale, we found it helpful to separate recombination from splicing and mutation, although we recognise that the boundaries between these two categories are somewhat fuzzy.

In many working GA models, recombination is used to create the dominate proportion of new routines introduced into the system. This allows the system to converge towards high performance rules quickly. In other works, most adaptation is based on demonstrably strong routines. The organisational equivalent of this emphasis is experimentation with new approaches close to and using the core competencies of the business. If learning is goal-directed, i.e. to achieve high performance, it seems quite logical to emphasize improvement and experimentation with proven strengths, rather than investing to heavily in randomly generated new possibilities. A review of cases of failed corporate venturing suggests that often the problem was exactly that the companies were unwilling to open up the core of the company to learning. The corporate sandboxes satisfied an externally imposed need to experiment, but protected the core of the company from that experimentation. Of course, recombination is based on reasonably accurate self-awareness in identifying core capabilities or knowledge, and in recognising and preserving good patterns of decisions and actions at both the organisational and individual level. Splicing is likewise based on self-awareness, because useful patterns or ideas are used in the new strings, although it involves less copying, and more analogy type connections.

Proposition 3: Adaptive entrepreneurial organisations demonstrate a strong willingness to leverage successes, including technical knowledge, functional activities, and people.

Splicing and mutating are more divergent forms of innovation, introducing new elements into the organisation. If adaptation was solely oriented to finding and replicating high-performance rules, the system would not maintain the requisite diversity to deal with future environmental changes. Splicing uses segments of established knowledge (from high and medium performance routines), recombining them creatively, and perhaps adding a random element. In addition, there is a small element of random adaptation. The well-known example of 3M's 15% rule, which allows every employee the option for independently directed 'play' time, is one way in which the adaptation of strong

performance can achieved. Another example is temporary cross-functional teams to work out a specific problem. Research shows these teams achieve better results when high performance employees are seconded to the teams.

Research in GA modelling has shown that mutation is generally overrated - contributes little to the performance of the system - but does provide the necessary spark which prevents the system from going 'dead'. In simulations of GA models, where performance goals and the environment have been change radically during the simulation, the diversity created in the adaptation cycle was enough to allow the model to quickly learn and adapt to the new regime. In fact, many companies do avoid projects which required both new technical knowledge and marketing competence (i.e. new customers or new channels), judging them too risky.

Proposition 4: Adaptive entrepreneurial organisations continually create new ideas, largely using the existing knowledge base to explore new related possibilities.

Our observations in EOs show similar attention to multiple levels of innovation, and to using performance-proven knowledge as the basis of most innovation. Rarely do the EOs we have observed focus solely on one level of innovation (i.e., incremental or breakthrough). They invest in incremental innovation (which finds a close parallel in recombination) and understanding and building off of core skills and abilities. They also permit some fundamental innovation, which may involve creating radically new concepts and knowledge. This is akin to "mutation" which works to provide random new forms. This could be seen in 3M, Trinkhaus and Burkhardt, DEC, and MHZ. Rhone-Poulenc and Siemens-Nixdorf were in the course of developing such variety in the system. It is clear, especially in the case of the last three companies, that it is important to develop the ability to continually analyse what is working, and to understand the degree to which it is actually working. New combinations were sought through extensive networking and sharing of customer needs, opportunities and basic knowledge.

Experimentation. The key to achieving the balance between performance and flexibility is not just in the mix of adaptations, but also some patience to experiment with the new forms. Experimentation, in our usage, refers specifically to the use of adaptations, not the adaptation process itself. Adaptations, then, are hypotheses for action, which are experimented with during the daily process of doing business. The need to experiment with new forms means that the selection cycle cannot be run too frequently, or else new rules will simply be eliminated before they have had a chance to be used. Even when the new adaptation is good, several operations cycles are needed for the attribution process to correctly assess the strength of the new rule.

Proposition 5: Adaptive entrepreneurial organisations do actively experiment, or more specifically, permit new knowledge and/or ideas to be tested and evaluated in practice.

The patience to experiment, to try out the adaptations in the real world setting, and to observe the performance results perhaps for several trials, is essential for learning. Sastry (1994) found in her simulation of responsive organisations that her simulated

businesses quickly spun out of control in high change environments, because they changed too quickly after each environmental change. To make the simulation work, she introduced 'learning time', where no new changes were introduced for several periods, regardless of the environmental conditions. The GA models suggest the same result. From our research cases we know that adaptive entrepreneurial organisations allow for experimentation, even failures, before cutting off projects. In fact, a key managerial challenge is judging how much patience is required: striking a balance between an escalating commitment to a failure and patiently nurturing a potentially big innovation is not easy. In the GA model, the selection cycle is fixed at the same length for all the rules. Thus, a 'long-shot' would be quickly eliminated - but it could also reappear in future adaptations.

A strong feature of entrepreneurial organisations, observed in all of them, is the willingness to experiment or test new combinations in practice. This is coupled with the constant challenge to strive and create new value. It is also coupled with the acceptance that in striving, failures will occur.

Termination. The complementary of adaptation is termination: when routines are simply not good enough, they must be replaced by other adaptations. By implication, this means that unsuccessful routines, whether they were formerly high performing or not, must be terminated. An AEO undertakes an active pruning of process, products and procedures. This insights echoes Schumpeter's argument that destruction (termination) is necessary to free up resources for new combinations. The idea of termination is implicit in life cycle theories of products. Yet the inertia of people and organisations makes termination particularly difficult. Numerous case histories in the transformation literature show how difficult to let go of traditions and habits.

Proposition 6: Adaptive entrepreneurial organisations demonstrate an ease with identifying and removing low performance products, processes and procedures from the portfolio.

The Failures Room underneath Lego's design department is a interesting example of how termination might work. The Failures Room is a museum of products which did not sell well. Every designer in the company has a key to this room, and is encouraged to visit it. The idea is not just to terminate bad products, but also to not repeat them. These products have been actively pruned from Lego's broad and frequently updated product line, terminated and yet the lessons learned are not completely forgotten.

Other examples of managing failures through constant evaluation of activities and termination of clear failures as early as possible could be scene in ABB and Siemens-Nixdorf. This is an important element in the evolving but strong culture of measurement and control in these two organisations. 3M has a regular, incremental review process for projects that makes strong "suggestions" regarding termination of activities that the review board feels are not making progress. The audit process of Rhone-Poulenc is similarly directed at incrementally evaluating progress in projects and pruning those that make little sense to continue. In the case of 3M and Rhone-Poulenc, it is known that termination of projects does not necessarily mean termination of the people involved -- they

move on to other projects and activities. DEC used to practice such failure termination, at least in the first 25 years of the company. It is carried out more informally than at 3M, but was frequently as decisive. This is one of the elements in what was formerly an entrepreneurial corporation that broke down and led to great difficulties for the company in the early 1990s.

Performance. The AEO must be clear and expressive at a given point in time of what constitutes 'good' or 'bad' performance, and are regularly evaluating performance. The performance standards may change over time, but are more commonly relatively stable, evolving slowly over time. Note that for the AEO, what is significant is both the nature of the performance goals and the nature of measuring performance. Both must be decentralised to maintain adaptiveness, which means goals must be local, consistent, specific or appropriate to goals and tasks. Each person must clearly understand the performance goals, if they also must measure actual performance.

The performance goals which can be specified in a GA model are inherently attractive to a business student. Like most businesses, GA models are driven by multiple goals, which can evolve and emerge with the system. The complex goal structure can operate as cybernetic goals (i.e. bad performance is above or below an acceptable range), or as maximization goals. GA models are not typically designed to optimize, however. As DeJong (1993, p. 13) notes:

> "[GAs] rapidly locate the region in which a global optimum exists, (but) they don't locate the optimum with similar speed. If one looks at this behavior from a "maximizing cumulative returns" point of view, it makes perfect sense".

Thus, we can structure our adaptive organisation to achieve multiple performance goals, and we can change those goals over time, without affecting the integrity of the GA adaptive system.

Proposition 7: An AEO relies on individual or local measurement of actual performance, based on well understood and agreed upon performance standards and goals.

Each of the entrepreneurial corporations studied was very keen on ultimate performance. Performance goals and criteria were shared. Each individual was made to feel responsible for the ultimate performance of the organisation. Some corporations, particularly Siemens-Nixdorf, make more individuals and groups ultimately responsible for profit and loss by measuring it at multiple points. 3M has a more delicate sense of performance, but there is a strong shared sense of performance the need for it. Everyone at 3M realizes the consequences of corporate non-performance. All of the entrepreneurial corporations reviewed, with the possible exception of DEC during its first 30 years, have very strong measurement systems which are constantly monitoring performance. The organisations, however, differentiate between measurement and control. They only control on agreed values.

The Structural and Process Characteristics of the AEO

Most of the key characteristics of the GA-style adaptive system deal with information accessibility. Information or knowledge is the glue and the energy which drives the adaptive system. Whether the information is sensed from the environment or retrieved from the memory, whether it is complete or not, it is the necessary condition to activate routines, generate inferences about the state of the world and useful action, judge performance, and direct the specific adaptations of routines. Information access is improved by transparency, common language, and good memory. In addition, to information accessibility, we acknowledge the importance of slack and self-awareness of strengths and weaknesses.

Transparency of Information. Given that information is important, and yet the adaptive organisations we are describing are not centralised, we quickly realise that transparency of information throughout the system is essential. Hidden files and restricted knowledge would quickly distort the fluidity of the system, for example, by limiting the potential number of routines activated, and by distorting performance judgments. By transparency of information, we mean that all of information is available to everyone. From a managerial perspective, transparency of information must be managed by actively removing barriers to information. Regular staff meetings where people are prompted to talk about what they know, even to people who seem irrelevant to the task, is one way of improving transparency. Widely distributed manual of valuable corporate services, for anyone in the company to read, is another way of increasing transparency. Sometimes management must deliberately destroy a social or cultural restriction against asking superiors for explanations before complying with directives, or challenging their information and inferences. AEOs do not assume that information flows freely to where it is most needed or can be used. Rather, they assume that barriers to information flow must be continually managed, whether they are barriers of habitual patterns, structures or politics.

Proposition 8: Entrepreneurial organisations use specific structures and processes to improve free access of information to all organisational members.

MHZ, a German company began the transformation into an entrepreneurial organisation five years ago, discovered that the practical benefit of a policy of information transparency is not achieved by simply supplying access. MHZ's management first made all information available to everyone, but discovered it did not change people's habits sufficiently. They then decided to encourage active response to the available information by teaching people to ask questions. Rather than assume they knew what the information implied, managers were taught to conduct staff meetings through questioning the assumptions and inferences and relationships in the information available. This step was crucial to achieving true transparency, because until then employees did not feel free to evaluate and use information from their own perspective.

In all of the organisations reviewed, the fact that information was to be made available was a fact of corporate life. Networking for information gathering is an accepted norm in such companies as 3M, DEC, ABB and Siemens-Nixdorf. There is very little impacted information. This does not suggest that some private information (e.g. personnel

records) does not exist, but that the norm and bias is to share information. Transparency, in practice, simply means people can see through organisational barriers to capture the information they need. It does not necessarily imply a distribution of every fact in the organisation. The rule might better be stated: seek and ye shall find. On the other hand, Trinkhaus and Burkhardt makes performance information very public as do other organisations interviewed.

Common Language. A point related to the last is equally as important: access to information does not guarantee that someone else understands it. Another prerequisite for effective access to information is a shared language within the organisation. Within a GA model, this is simply achieved by keeping the length and content of message and the condition segments of the routines similar. An organisation cannot impose this simplicity on its world, but rather must actively translate information into a common language. We would therefore propose that AEOs are more actively concerned with developing a common language.

Proposition 9: Adaptive entrepreneurial organisations use a common language for information flow across internal boundaries.

In our previous papers on adaptive entrepreneurial organisations, we developed a knowledge cycle, which noted that between sensing information, and analysing it for action and learning, there must be encoding. By encoding we mean that information is sensed or created by people in different parts of the organisation, and these people must 'translate' the information into a language that is widely recognised within the organisation. For example, Dougherty's work on thought worlds has highlighted the barriers to communication between functions. Adaptive entrepreneurial organisations try to create a broader common language for passing information which transcends the functional (and other) barriers to communication.

Our case research in MHZ, DEC, Rhone-Poulenc, Siemens-Nixdorf and others have all shown a great deal of attention to building and/or maintaining a common reference language. Beyond transparency, the entrepreneurial corporations interviewed suggested that it was necessary not only to have access to information, but to have the information expressed in such a way that it could be reinterpreted quickly. Simply stated, people need to communicate. The ability to adapt is degraded if there are difficulties "encoding" and transferring encoded information. The AEOs reviewed tried to maintain both a common language around key topics but also diversity in the language to permit encoding of difficult concepts. Developing a common language for "encoding" and ensuring that information will be properly understood is key. This required a common language or way of expressing what is observed. Inside most corporations, short-hand phrases develop. Within the entrepreneurial corporations, language is richly embellished by folklore (e.g., 3M, Trinkhaus and Burkhardt). This folklore also becomes part of the short-hand language. In addition, clear protocols for communicating develop (e.g., at DEC and MHZ).

Memory. The GA model for adaptation makes considerable use of memory in the system, albeit decentralised memory. For the attribution cycle to work, actions taken and linkages

between routines must be remembered at a rather specific level. The linkage of routines through posting messages also indicates that information is being processed and remembered and reprocessed through a chain of routines. For the adaptation cycle to work, we must remember which routines delivered high results, and which ones are too weak to retain, and which ones need engineering. Likewise, the adaptive entrepreneurial organisation requires a strong and specific memory, both for information and actions. This does not imply an amazing centralised computer system, but it does mean that individual people make an effort to encode, share and retain information, even when that information does not directly impinge on their task.

Proposition 10: Individuals within adaptive entrepreneurial organisations remember past actions and information in detail.

In our previous work we noted that adaptive entrepreneurial organisations did not actually see more information that other organisations, but that they seemed to forget less (Muzyka, et al., 1995). MHZ formally captures such information in computerized databases which are available to all of management. DEC and 3M systematically have attempted to retain information flexibly in the network. ABB, Rhone-Poulenc, and Siemens-Nixdorf are attempting to create memory in the system as well as through the human network. Some of the information, especially concerning management performance and administrative history, is captured in folklore. Anthropologists would be impressed by the rich oral history some of these organisations.

Self Awareness. An organisation which plans to leverage its high performance routines, and replace low performance with new ideas, must be explicitly aware of its own activities. The need for self-awareness with an adaptive entrepreneurial organisation operates at all levels of analysis, from the organisation as a whole, to individual members. Leaders of the adaptive entrepreneurial organisation must be clear on high performance routines, so that resources and people can be leveraged to create more high performance. This is fundamentally similar to core competency or resource-based diversification arguments, where resources are directed to areas where the company already has demonstrated competitive advantage. At the individual or team level, people must be self-aware of their own routines. This self-awareness, in the sense of explicitly identifying the value of activities, is the first step to both attributing performance rewards and to generate improvements through adaptation. The self-awareness at all levels takes on a dual character of praising successes, thus focusing more energy to those areas, but also of criticizing even the strongest people and routines so that learning and improvement can proceed.

Proposition 11: Adaptive entrepreneurial organisations explicitly identify and discuss activities in terms of performance and possible improvements.

In the entrepreneurial corporations interviewed, self-awareness was a key feature. Individuals were challenged to be self-disciplined in their activities and aware of the ways their activities impact others. One way in which self-awareness at an organisational level was enhanced was that individuals and groups are charged with ensuring communication with units potentially affected by their work. Top management did not assume this

information coordination role, with the result that individuals and were much more aware of their own position and influence within the whole, as well as the work of other groups and individuals. 3M and DEC contain many examples of this behaviour in their history.

Slack. The GA model assumes that adaptation resources exist and are generated within the system. From the perspective of organisations, we know that resource availability is important for change and innovation, but by no means automatic. The resources are partly required to create new forms, new routines for action and information processing, but are even more essential for the experimenting with the new forms. The slack exists to enable long-term adaptability of the organisation, and therefore long-term profits. We will discuss specific aspects of innovating and experimenting later, but here we simply acknowledge that adaptive entrepreneurial organisations have necessary slack built into the system.

Proposition 12: Adaptive entrepreneurial organisations allow slack to stay in the management and budgeting systems.

3M is a common example of how slack may be structured into a business organisation. Besides the 15% rule, which allows employees to self-manage their time investment in projects, 3M also provides decentralised financial resources for projects. The practical implication of these multiple sources of funds is that projects can apply for support several times under several different decision criteria. 3M does not tightly nor centrally control investment into innovation projects, which by implication means that slack exists in the system. We also observed slack management' at DEC during its first 30 years. The phrase "you worry about your opportunity, we will worry about the funds" was an often heard phrase. Slack is actively ensured in entrepreneurial organisations, though it isn't always visible.

A PROPOSAL FOR A CONTINGENCY MODEL FOR ADAPTATION IN ORGANISATIONS

Based on our explorations of adaptive systems, we would like to propose a simple contingency model which may also be a useful output from this mapping of findings from genetic algorithms on organisational adaptation. The major interest in further developing a GA model of the adaptive entrepreneurial organisation is that it allows us to explore several key relationships. Specifically, we are interested in the profitable survival of the firm under various environmental conditions. Environmental changes range from small incremental changes to discontinuous shifts, and the rate of change can be invisibly slow to hyper-accelerated. In addition, the firms performance goals might shift in emphasis or dimensions.

Within a given context, we wish to explore the relationships between what we perceive to be the key managerial levers, i.e. the relative mix of adaptations from the recombination to splicing to mutation continuum, the percentage of routines replaced in each selection cycle and the rate or timing for the attribution and adaptation cycles. Our hypothesis is that in a low change environment, cycle time is less relevant, but performance would benefit from an emphasis on recombination. Nothing is gained by

excessive experimentation in this context. On the other hand, with a higher rate of change, a fast cycle time would require a strong mix of recombination and splicing, to ensure benefitting from any performance improvements. Probably, though, the system would benefit from slow cycle time, and a higher mix of splicing and mutation. It is ironic, yet logical, that discontinuous change requires a slower, not faster cycle change. The actual adaptations of splicing and mutation are dramatic and swift responses to the changes, but slower cycle time is necessary to allow for experimentation with the new adaptations -- thus we see a fast-slow response to change (see Figure 2 for a graphic display).

At this point in the research program, our contingency hypotheses are still rather speculative. We believe, however, that our approach to the question can help resolve some of the conflicting prescriptions for responses to hyper competition. Rather than producing yet another list of things that should have been done yesterday, we are using GA models to explore the fundamental mechanisms for organisational adaptation. This theoretical basis will then drive our future development of testable contingency hypotheses.

SUMMARY AND NEXT STEPS

This has been a first attempt on our part to link work in adaptive models from artificial research with research into how to build more adaptive organisations. If we are to build more entrepreneurial corporations, a major element is adaptation. We believe, based on this first effort, that there is an analog between these adaptive models originally designed to simulate human decision process and the key elements required to create an adaptive organisation. The ability to create an internally consistent model, even if the model is not a simple combination of deterministic linear equations, which might be able to describe the essential elements of an adaptive organisational system may be within our grasp.

The suggested organisational implications of GA models for organisations, as evidenced by the propositions above, find a strong parallel in the reality of entrepreneurial corporations. While we have not provided every example that validates the propositions in practice, we see a significant amount of data supporting the propositions. We also believe the GA models provide a new framework with which to organise the observations. During the course of the conference, we will be happy to provide additional information.

REFERENCES

Ackenhusen, M., Muzyka, D. F., & Churchill, N. C. (1996a). Restructuring 3M for an Integrated Europe, Part One: Initiating the Change. European Management Journal, 14(1), 21 -36.

Ackenhusen, M., Muzyka, D. F., & Churchill, N. C. (1996b). Restructuring 3M for an Integrated Europe. Part Two: Implementing the Change. European Management Journal, 14(2), 151-159.

Block, Z., & MacMillan, I. C. (1993). Corporate Venturing. Boston, MA: Harvard Business School Press.

Burgelman, R. A. (1983). A Process Model of Internal Corporate Venturing in the Diversified Major Firm. Administrative Science Quarterly, 28, 223-244.

Christensen, P. S. (1989). Strategy, Opportunity Identification and Entrepreneurship: A Study of the Entrepreneurial Opportunity Identification Process. Aarhus: Institute of Management, Aarhus University.

Churchill, N. C., De Koning, A. J., & Muzyka, D. F. (1995). Entrepreneurial Organizations: What They Are, Why They are Important, and How to Become One. In ENDEC Conference, . Shanghai:

Cyert, R. M., & March, J. M. (1963). A Behavioral Theory ofthe Firm. U.S.A.: Prentice Hall.

De Koning, A. J., & Muzyka, D. F. (1996). The Convergence of Good Ideas: How do Serial Entrepreneurs Recognize Innovative Business Ideas (Working Paper No. 96/16/ENT). INSEAD.

DeJong, K. A. (1993). Genetic Algorithms are NOT Function Optimizers. In L. D. Whitley (Eds.), Foundation of Genetic Algorithms 2 (pp. 5-18). San Mateo, CA: Morgan Kaufmann Publishers.

Holland, J. H. (1975). Adaptation in Natural and Artificial Systems. New York: Harper Collins.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1986). Induction: Process of Inference, Learning and Discovery. Cambridge, Mass: MIT Press.

Koza, J. R. (1992). Genetic Programming. Cambridge, Mass.: MIT Press.

Koza, J. R. (1994). Genetic Programming II, Automatic Discovery of Reusable Programs. Cambridge, Mass: MIT Press.

Muzyka, D. F., de Koning, A. J., & Churchill, N. C. (1995). Adaptation and Transformation. EMJ(December).

Nelson, R., & Winter, S. (1982). An Evolutionary Theory of Economic Change. Boston, Mass.: Harvard University Press.

Pinchot, G. (1985). Intrepreneuring. New York: Harper and Row.

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(4), 397-416.

Sastry, A. (1994). Problems and Paradoxes in a Theoretical Model of Punctuated Organisational Change. In

Stevenson, H. H., Roberts, M. J., & Grousbeck, H. I. (1985). New Business Ventures and the Entrepreneur. Homewood, Il: Irwin.

Story, J. (1993). The New Europe: Politics, government, and economy since 1945. Oxford: Blackwell Publishers.

Tushman, M., & Anderson, P. (1986). Technological Discontinuities and Organization Environments. Administrative Science Quarterly, 31, 439-465.

Van de Ven, A. H., Angle, H. L., & Poole, M. S. (1989). Research on the Management of Innovation. New York: Harper & Row.

Figure 1
The Basic Operational Cycles

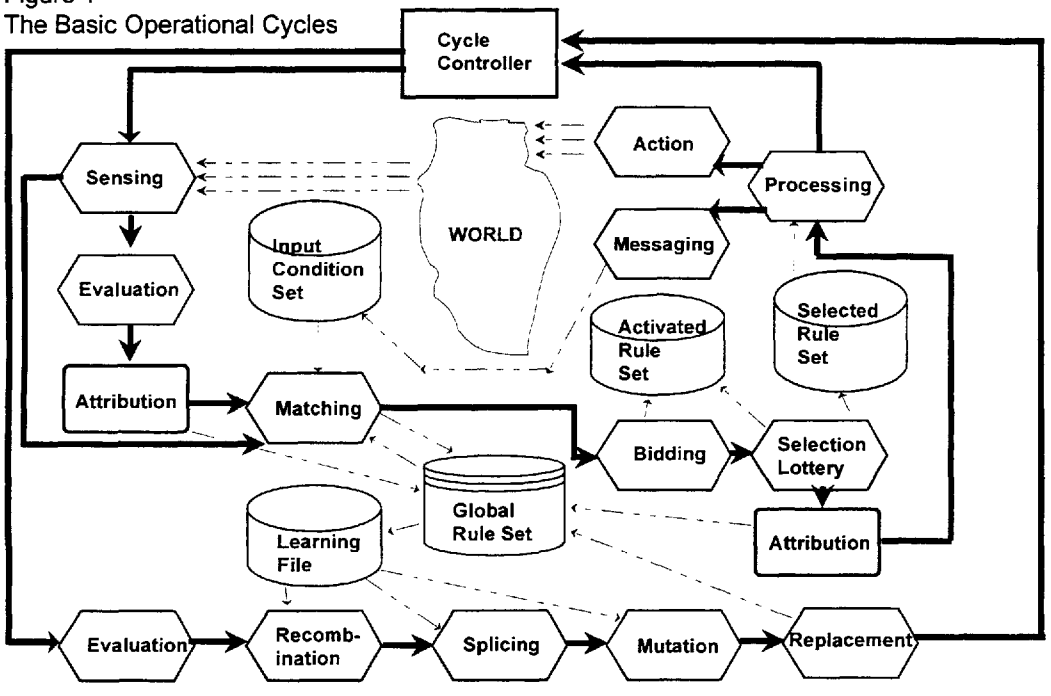

Figure 2
Contingency Framework for Adaptation Strategy

		Control Variable	
		Short Cycle Time	*Long Cycle Time*
State Environment	*Low Slow Change*	**Emphasis on Recombination; Low/No Mutation** **Low % Replacement**	**Emphasis on Recombination; Low/No Mutation** **Low % Replacement**
	Fast High Change	**Lots of Recombination; Lots of Splicing** **High % Replacement**	**Lots of Splicing; Some Mutating** **Medium % Replacement**