Developing Adaptive Organizations, Part I: Research Findings

Our last newsletter concluded that, despite their best efforts to make decisions in the face of uncertainty, managers will at some point inevitably find themselves surprised by an unexpected change in their environment. Given this, a key question they face is how to ensure that their organization will be able to adapt to such a change in a manner that ideally furthers its success, and at minimum guarantees its continued survival.

In the management literature, there are two main schools of thought on this issue. The largest body of work on the subject is generally optimistic, and asserts that most organizations are capable of successfully adapting to changes in their environment. However, writers belonging to the "population ecology" school are much more pessimistic. They believe that in most organizations there are substantial obstacles to adaptation. As a result, most of the response to environmental change takes place at the "population" level, in the form of old firms dying (e.g., via bankruptcy or merger) and new ones being born through entrepreneurial action.

At the macro level, evidence with respect to these two schools is conflicting; there are plenty of case histories of successful change, as well as some powerful data on the strength of population level effects (e.g., only 29 of the firms listed in the first Fortune 500 in 1956 are still in existence in 1998). At the micro level, as any manager can attest, there is no shortage of claims (usually by consultants) as to the efficacy of various programmatic approaches to "change implementation", "transformation", and "organizational learning." However, a 1997 analysis of three of the most popular of these approaches (downsizing, reengineering, and total quality management) by the National Research Council (a joint undertaking of the National Academies of Science and Engineering) concluded that "none of these approaches appears to have a consistently positive relationship with organizational effectiveness...Spotty evidence of their impact on effectiveness exists, but evidence also exists that each approach is useless, if not harmful...[However], the reasons for the wide-scale failure of these approaches remain unclear."[1]

Where does this leave a manager faced with the challenge of ensuring the adaptability of his or her organization? For the past four months we've been looking at this question, searching for answers that are grounded in data and analysis rather than assumptions and assertions. This issue of Management Insights (and the next one) summarizes the results of our explorations.

Our starting point is the relatively recent body of scientific findings about the nature of non-linear systems.

Linearity and its Discontents

"We are essentially linear creatures. Whether this is the native mode of humanity or whether it is primarily the result of acculturation is open to question...However, it is unarguable that our society fosters and rewards linear behavior and performance from kindergarten on. Our educational system teaches it and grades on it, our workplaces hire, fire, and promote on it, and our governmental programs are designed and executed on it...Often associated with the name of Sir Isaac Newton, Newtonian or linear science has become a powerful philosophy to both describe and ultimately control nature.[2] The key features of linear systems include the following:

- The dominant feedback loops in a linear system are negative ones which act to dampen the effect of any change. As a result, changes in a linear system are proportional: small changes in inputs lead to small changes in outputs, while big changes in inputs lead to big changes in outputs. Given this, the natural state of a linear system is equilibrium. Any disorder that exists is caused by outside shocks to the system, and is transitory.

- Cause and effect in a linear system are readily understood and demonstrable; as a result, one can safely predict that an action taken under a given set of conditions will always produce the same result.

- A linear system is additive: the whole is equal to the sum of its parts. This leads to reductionism, or the practice of taking a large, complicated problem and breaking it down into more manageable pieces, analyzing them, and then combining the results to arrive at an optimal solution.

- From an organizational perspective, the purpose of control in linear systems is to gain certainty (i.e., minimize surprises). Control in linear systems tends to be a top down process: Managers are said to be "in control" and situations (or organizations) "under control." Control, in this sense, is achieved by the use of negative feedback loops that compare the outcomes of behavior to pre-existing targets, and trigger responses to minimize deviations from plan.

By way of contrast, non-linear systems have very different properties:

- The dominant feedback loops between the variables in a non-linear system are positive -- rather than dampening the impact of a change, they magnify it. Non-linear systems are therefore disproportional: small changes in inputs can product big changes in outputs, and vice versa. As a result, disequilibrium is the natural state of affairs in a nonlinear system, and can persist for long periods of time even in the absence of outside shocks.

- Cause and effect in non-linear systems are widely separated, and generally difficult to identify with each other. Predicting the future behavior of non-linear systems is therefore extremely difficult.

- Non-linear systems are not additive; the whole may be equal to more or less than the sum of its parts. This usually makes reductionist analysis of nonlinear systems impossible – they must be understood at the level of the whole system, rather than the individual parts. Because of the difficulty in doing this accurately, managers solving problems in nonlinear systems typically seek a solution that is "good enough" rather than optimal, because the latter is inherently unknowable.

- As we will further discuss later on , in a non-linear system control is not something imposed by management; rather it is a quality of the system as a whole. The aim of control, in this sense, is to maximize the adaptiveness of the overall system by avoiding the two extremes of steady state equilibrium and purely random activity. This is achieved through simultaneous bottom up and top down processes that employ both negative and positive.[3]

"Linearity's real power has been in its ability to produce technology -- systems designed to behave predictably...On every other front, however, linearity has not come close to meeting with the success it has enjoyed in technology...When linear thinking has met the multitude of non-linear phenomena that exist in the real world, "it has established its domain by either imposing non-linear substitutes or making excuses. In the former case, the linear has invented the calculus, statistical techniques, and operations research. These methods work well in linear and mildly non-linear environments. But beyond a certain point, these methods give way to rounding, 'good enough for government work', experts who contradict each other, and the unintended results of well intended actions (Murphy's Law). The results of applying linear thinking to nonlinear problems are all round us in painful and embarrassing profusion."[4]

By now, it should be clear that it is well worth our while to better understand non-linear systems. It is to this subject that we will now turn, beginning with simple (deterministic) nonlinear systems, and then moving on to complex (adaptive) ones.

Deterministic Nonlinear Systems

A deterministic nonlinear system is a network of agents (or variables) whose behaviors are determined by a common set of fixed, unchanging rules. Agents in this system do not adjust their behavior in light of their consequences -- there is no learning of any kind.

A simple deterministic nonlinear system is described by the following, which is often referred to as either the "logistic" or "predator-prey" equation: $X_{t+1} = rX_t(1-X_t)$. For example, consider a valley populated by just two species, sheep and wolves. Migration in or out of the valley is impossible, and the food supply (grass for the sheep) is fixed. The percent of sheep in the valley's animal population next year (X_{t+1}) is equal to the percentage of sheep this year (X_t) times their net growth rate (r, or the difference between sheep birth and death rates) limited by the total food capacity available to support more sheep ($1- X_t$).

When this equation is run over fifty periods, it displays some very interesting behavior as the value of r increases. Given a starting value of $X_t = .5$ (ie., half the valley's population is sheep), when r is less than 2, the value of X declines to zero, and sheep become extinct. When r is between 2 and 3, X reaches a steady state equilibrium at increasingly higher percentages of sheep in the valley's total population. Up to this point the system's behavior has been linear. However, as r increases beyond about 3, the system begins to act in new way. At first, it oscillates between two values (or equilibriums) for X. Then, as r increases further, it oscillates between four, then eight, then sixteen equilibrium values for X. Moreover, each of these new "bifurcations" (ie., doubling of potential equilibrium values) comes more quickly -- the length of time between each doubling is approximately 22 percent of the previous interval. As r increases beyond about 3.7, the bifurcations increase to the point at which it appears that X takes on random values in each period as the sheep population either explodes or collapses.

The state of behavior to which the system (i.e., the sheep population) converges is technically called an "attractor." In other words, an attractor is the pattern of behavior into which a system ultimately settles in the absence of further changes to its control parameters (in this case, r). It is a possible outcome that is contained within the simple set of rules that govern the operation of the system. A stable equilibrium attractor is when the system is drawn either to a single point (a "point attractor") or a periodic cycle (a "periodic attractor"). Deterministic nonlinear systems are attracted to a stable equilibrium when their control parameters are low. When the control parameters are turned up high enough, the system is drawn to an "unstable attractor", which is another name for completely random behavior.

However, and critically for our purposes here, stable and unstable attractors aren't the only ones that exist in a deterministic nonlinear system. Studies have uncovered the presence of another attractor that lies on the border between stability and instability. This third attractor typically goes by either of two names: "strange attractor" or "chaos". It is an attractor in the sense that it has definable boundary dimensions (like the stable attractor, and unlike the random attractor), but it is strange in that unlike the case of the stable attractor, tiny changes in the initial value of X_t will generate extremely different values for subsequent Xs within those boundary dimensions. This phenomenon is known as "sensitive dependence on initial conditions." Therefore, a strange attractor defines a bounded range of outcomes into which the system may move, yet within which specific trajectories cannot be predicted because of sensitive dependence on initial conditions. The behavior of the system at this point neither random, nor stable. Chaos is a paradox. This bounded unpredictability introduces instability and disorder into the system, but also opens it up to immense variety without pushing it over the edge into disintegration (randomness). In this sense, chaos is therefore the archetype of creativity and innovation.[5]

At this point, another example is probably in order. Consider the nature of an ice cube: the hydrogen and oxygen molecules that comprise it do not move; they are stable. At the other extreme, consider steam: the molecules that comprise it move randomly. Now consider water, in which the molecules are capable of moving in an infinite variety of different ways (e.g., a dripping faucet, waves, rapids, etc.) while always remaining as water. Assuming you were looking for creative inspiration, which would you rather stare at: ice, steam, or water?

Okay, you say, point taken, but what does this have to do with making my organization more adaptive? Two things. First, many attempts at organizational change seem to be based on two implicit assumptions: (1) the environment will change in ways we can predict ahead of time, and (2) if we get the rules right, we can predict that the organization will behave in a way that leads to success in that environment. This looks a lot like a deterministic system. And that means that if the environment changes in a way we hadn't anticipated (i.e., r goes higher than we had expected), the result can be anything but predictable behavior.

Second, the idea that a system's maximum creativity (and adaptability) occurs at the "phase change" between stability and instability is critical. This will be explored further in the next section.

Adaptive Nonlinear Systems

Unlike the agents in a deterministic nonlinear system, the agents in a simple adaptive system have a purpose, and adjust their behavior in light of its consequences for achieving that purpose (i.e., they are capable of type one or single loop learning). Simple adaptive systems can produce remarkably complex behavior. Consider a famous example. In a simulation model called "boids", agents are given three simple rules: keep up your speed, avoid obstacles, and maintain a certain distance from the agents around you. When the simulation is run, these three simple rules produce flocking behavior. In other words, a new behavior (flocking) "emerges" from the interaction of agents using these rules, rather than being specified at the outset of the simulation. This of "self organizing behavior" is critical characteristic of all adaptive nonlinear systems, and is an important phenomenon to keep in mind as we go forward.

In a complex adaptive system (or CAS), agents are also able to adjust the rules governing their behavior (i.e., they are capable of type two or double loop learning). Much of the learning about the behavior of these systems has come from evolutionary biology. In this context, biologists have replaced organisms' genetic code with computer code, and then run simulations to see how a complex adaptive system evolves over time. Because their findings are central to this paper's conclusions, it is worthwhile looking at them in detail.

The most common tool used to simulate the evolution of complex adaptive systems is called the NK model. In this model, an organism has N number of genes, each of which can have a value of 0 or 1. As we know from life, some genes are particularly valuable to have; they make us "fitter" relative to the other organisms with which we compete for the resources needed for survival. However, the "fitness value" of a given gene may depend not just on its own value, but also on the value of other genes. For example, the value of extremely good eyesight in a baseball player is compounded if she is also highly coordinated, fast, and strong. In modeling terms, the fitness value of each of an organism's N genes depends not only on its own value, but also on the value of K other genes. Finally, the overall fitness value of an organism is logically a function of the fitness value of each of its genes. As NK models are usually structured, each gene has a fitness value between 0 and 1,and the fitness value of the total organism is defined as the average fitness value of its individual genes.

In terms of nonlinear systems, the K value in the NK model is analogous to the "r" value in the deterministic nonlinear system described above. It is the critical control parameter that determines whether the nature of the system is stable, random, or in the phase transition between the two. At this point, we need to make a short point about terminology. In the study of deterministic nonlinear systems, the phase transition between stability and randomness is known as "chaos." Clearly, this is not what most people have in mind when they hear this term. To make matters more confusing, in the study of complex adaptive systems randomness is called "chaos" (i.e., in this case, it means what most people have in mind when they use the term), while the phase transition is called either "complexity" or the "edge of chaos". If nothing else, this shows that the people doing this work are pure minded academics rather than slick marketers!

Back to the properties of an NK model. Consider a model with ten genes (N=10), which can be of either of two types (0 or 1) and whose fitness value depends only on themselves (K=0). Different combinations of genes (i.e., N_1=0, N_2=1, N_3=1, etc.) will yield different fitness values for the organism as a whole. In total, there are 2^{10} different possible combinations of genes. One of these combinations will produce a fitness value that is higher than all of the others. The best metaphor for visualizing this is that of a landscape. Looking down on such a landscape, one would see the ground rising up around a single peak. In the language of NK models, such a landscape is termed a "smooth" because the heights of different points on it are highly correlated with one another.

The simplest form of adaptation in an NK model is called mutation. In a mutation operation, a new organism is created by switching the value of one of the first organism's genes. The resulting organism may have a higher or lower fitness value. In the next period, both organisms mutate again, so that the model now has four organisms. As this process continues, the resource constraint is reached on the environment's ability to support further new organisms. When this happens, organisms with a fitness value below a certain level are eliminated. This process is known a selection. As this process continues over time, it will result in all organisms having the combination of genes that yields the highest possible fitness. They will all end up on the landscape's single high peak.

The second form of adaptation in the simple NK model is recombination. In this case, a new organism is formed by copying half the genes from each of two other organisms. Because more genes are changing each turn, recombination reduces the time required for all organisms to reach the highest peak on the landscape.

Now consider the opposite case, where N still equals 10, but now K = 9. In other words, the fitness value of any gene depends on the fitness value of every other gene. In this case, the fitness function is highly nonlinear, and the correlation between peaks on the landscape is equal to zero. Rather than a single "global" peak, local peaks are scattered randomly, with no relationship between them. Mathematically, the number of local peaks in this model is equal to ($2^N/N+1$, or 93). This very rugged landscape has some interesting characteristics.

First, the average value of all the peaks is much lower than in the K=0 case. Stuart Kauffman, the biologist from the Santa Fe Institute who has popularized the NK model, has termed this the "complexity catastrophe." As he puts it, "because so many constraints [K] on the value of each N are in conflict, there is now a large number of rather modest compromise solutions rather than a single superb solution. There are, in other words, many local peaks of a rather low height."[6]

Second, because there are more local peaks, the expected length of time required for an organism to reach a peak is much shorter than when K=0 and the landscape has a single high peak. However, it is now possible that a mutating organism will become "trapped" on a local peak that is lower than the highest peak on the landscape. Consider an example. The simulation begins with three organisms (A, B, C) located on different parts of the landscape. Each organism goes through a process of single gene mutation. Because the landscape is completely uncorrelated, succeeding generations of each organism will "climb" the peak that is closest to them. Over a number of periods, this results in all the A organisms being concentrated on a local peak whose height (or fitness value) is .5, all the B organisms on a local peak with a height of .6, and all the C organisms on a local peak with a height of .8. After this point, succeeding rounds of mutation and selection will result in the elimination of organisms A and B because they were trapped on relatively low local peaks in the rugged landscape. Moreover, on a completely uncorrelated landscape, recombination doesn't offer a way out of this trap. Producing a new organism via recombination is just as likely to result in a lower fitness value as is producing one via a one gene mutation.

The third point is that on a truly random landscape, the number of directions that lead uphill is cut in half with each step. In other words, when only one gene is changed each period, the probability that the next change will produce a higher fitness value is cut in half with each successful mutation. Consequently, the time required to find a higher fitness value doubles with each improvement.

Having looked at the two extreme cases, where K=0 (a perfectly correlated landscape with a single high peak), and K=(N-1) (a perfectly uncorrelated landscape, with randomly distributed low peaks), let us now consider the most realistic case, where K is somewhere between 0 and (N-1). What have simulations taught us about the properties of this type of complex adaptive system?

As K increases, the time required to find a higher fitness value via mutation (also called "local adaptation") will grow at an increasing rate until, once K is around 8, it plateaus at a factor of about two (the time doubling situation described above).

Imperfectly correlated landscapes are characterized by different "mountainous regions", within which the height of various peaks is positively correlated. However, these regions can exist in widely separated parts of the landscape. "Recombination on such landscapes can take one of two forms: with other organisms located in a correlated or local region, or with organisms from outside that region. In the case of the latter, every time a "long jump" (that is, a recombination with an organism from outside the local region) results in a higher fitness value, the expected number of long jumps to find a still fitter combination doubles."[7]

"This give rise to the phenomenon that the development of organisms on a correlated landscape should show four time frames [or, in other words, a predetermined life cycle]. Assuming the initial organisms start on a peak of average height, in the early stages of the simulation, many long jumps will be successful, giving rise to many different types of organism. As long jumps become more difficult, the population will begin to climb up local hills to improve their fitness [i.e., mutation will produce the most improvement in fitness]. A sharp slowing in the rate of fitness improvement will ensue, as progress up local peaks becomes more difficult. Finally, selection pressures will eliminate organisms trapped on relatively low local peaks."[8]

Up to this point, we have assumed that the fitness function that (along with the size of K) defines the landscape has remained fixed. In other words, the evolution of the organism we have been describing has occurred in isolation. In the real world, however (and even in our deterministic nonlinear model of sheep and wolves) this isn't the case. The fitness value of a given combination of one organism's genes is often a function of the combination of one or more other organisms' genes. This means that a change in one organism's combination of genes can have a substantial effect on another organism's fitness. In other words, changes in one species' genes can cause the landscape faced by another species to "deform", or take on a new shape. When this is the case, the organisms in question are said to be "coevolving."

Biologists have expanded the NK model to simulate coevolution by adding two new variables, "C" and "S". In the NKCS model, the fitness value of a given gene (N) in a given species depends not only on its own fitness value, but also on the fitness value of K other of its own genes, and on the fitness value of C genes in S other species. Think of it as a system made up of linked NK landscapes (or a multidimensional chess game).

In this model of coevolution, the critical control parameters are K, C, and S, which affect the degree of "connectivity" in the system. Depending on the level at which they are set, the model will function in one of three states.

The system will be in equilibrium (known as an "evolutionary stable strategy", or "ESS") when either (a) K is high for each species (i.e., they all have lots of local peaks to become trapped on); (b) C is low, so landscapes don't deform (change) too much when other species make adaptive moves; or (c) S, the number of other species interacted with is low. In the ESS average fitness values throughout the system are low.

At the other extreme, the system will function randomly (or be in the "chaotic" zone) when (a) K is low for each species (i.e., they are all trying to move up to a few peaks); (b) C is high, so the location of these peaks frequently shifts as each species makes its own adaptive moves; or (c) the number of other species interacted with (S) is high. In this situation, known as the "red queen effect", average fitness values are also low, because each species is chasing high peaks that move too rapidly to be climbed.

When organisms in each species are allowed to vary their values of K, C, and S, simulations show that the system as a whole will naturally converge on the values for K, C, and S that produce the highest average fitness values for all the coevolving species. Not surprisingly, these values occur precisely in the system's phase transition between order (ESS) and chaos (red queen). More importantly, as in the case of the boids' flocking behavior, these values for K, C, and S are not programmed into the system beforehand; rather they naturally emerge as the system "self-organizes."

This property of "self-organization" in complex adaptive coevolving systems is the one that has most fascinated complexity scientists, suggesting as it does the presence of very powerful and heretofore unimagined laws that govern our existence.

Beyond this, the NKCS model also suggests the logic behind another phenomenon that has been observed in the fossil record of life on earth: punctuated equilibrium, or the tendency of extinctions to occur in a pattern that follows a power curve (i.e., a large number of small events, and a small number of large events). In the NKCS model, when species are allowed to invade each other's landscapes (e.g., reproducing a variant of species A and letting it loose in species B's landscape), it becomes possible for a species trapped on a relatively low peak on a rugged landscape to be driven extinct by a fitter invader (which enters the landscape in a different region with higher peaks). Depending on the nature of interconnections between the extinct species and other species, this can trigger a subsequent sequence of extinctions in other species that is either small or large. Typically, these subsequent extinctions are small, but occasionally they are very large. As in the fossil record, they tend to follow a power law distribution. Unsurprisingly (by now) the probability of large extinction events occurring in a given system is minimized when it is operating in the phase transition between order and chaos.

The Relevance for Business

Are the findings from the NKCS model relevant for business? In some important ways, business is very different from the NKCS world. First, consider the difficulty in trying to describe a business in terms of its genes. One could take either of two approaches. On the one hand, you could set the value of each gene equal to a simple decision (yes or no), and then define each N as a major policy decision that each organization operating in an industry has to make. Assuming the decisions cover, at minimum, marketing, production, finance, human resources, systems, corporate development and R+D considerations the number of genes (Ns) involved becomes very large (akin, perhaps to human DNA), and the range of possible combinations huge (e.g., 2^{35} assuming "just" five key decisions in each of these seven areas). On the other hand, you could allow for a larger number of values for each gene, and use fewer Ns overall -- however, you would still be left with a very large number of possible combinations. Finally, the reality is that in business (and unlike the model), the total number of variables that define the landscape containing all possible combinations open to a manager can never be known with certainty. Depending on your approach, there will either be new Ns or values of N that you don't know about.

Without going any further, one can immediately see an application to business: given the number of possible "gene combinations" in a typical industry, a systematic qualitative examination of each in search of the "best one" is out of the question. Moreover, assuming that more than four of these decisions are interrelated, the system itself will be highly nonlinear, and therefore not susceptible to optimization using any type of mathematical algorithm. Once again, a comprehensive search is out of the question. In short, when defining their strategy (for which searching for the highest possible peak in a rugged landscape seems a very good metaphor), managers cannot and should not spend their time trying to identify "the best one"; instead they should look for one that is "good enough."

The adaptive processes utilized in the NKCS model are also different from those utilized by a business. Unlike the model's mutation operation, businesses can and do change more than one gene at a time when they are trying to improve their fitness (e.g., their market value). In this sense, their adaptation process is much more akin to recombination. However, it does not occur randomly, as in the model. Businesses evaluate other organizations' performance before changing their own (e.g., via benchmarking or strategic alliances). They attempt to adapt intelligently. This difference and its implications is the subject of a recent working papers by Jan Rivkin of Harvard Business School and Daniel Levinthal of Wharton.[9] Both Rivkin and Levinthal note that many management approaches emphasize the importance of fit between an organization's strategy and structure -- in other words, they stress the need for K to be relatively high. Consequently, industry landscapes will tend to be relatively rugged (i.e., only moderately correlated), which should result in a normal distribution that leaves most firms on relatively low fitness peaks. In such a situation, one would expect to see many firms attempting "long jumps" to improve their performance. In the real world, however, such jumps have to contend with three factors that are absent from "biological" versions of the NKCS model: barriers to entry, uncertainty, and cost.

First, unlike the NKCS model, some recombinations are unavailable to the firm because they are protected by strong barriers to entry (e.g., scale, patents, sunk costs, etc.).

Second, uncertainty in the real world comes in two forms: first, whether or not the firm accurately perceives the value of the trait (N) it is trying to copy, and second, whether or not it can accurately replicate that trait in its own organization. On a rugged landscape, these uncertainties raise the possibility that the change you are attempting will land you not on a higher peak, but rather in a valley surrounded by peaks (which you can eventually climb) that are lower than the one you left. In other words, a long jump (multiple N) change can result in a decline in overall fitness (market value) rather than an increase. Examples of this abound in the real world. Moreover, the probability of this happening increases with the number of strategy and structure variables that you change (in other words, the probability of reaching a higher peak declines with the number of variables you are changing).

Looked at from another perspective, it is also easy to see how a highly integrated strategy (e.g., high K) in and of itself provides a barrier to successful imitation. As Rivkin notes, this explains “why some firms continue to earn higher returns than competitors long after the causes of their success have been documented and publicized.”[10] However, as Levinthal points out, the flip side of this is that it is as hard to change such a strategy (if the landscape deforms) as it is to imitate it.[11]

Third, in the business world, change is not costless. The very act of changing may therefore have an impact (though a transitory one) on an organization’s fitness value (when that is defined as its market capitalization). Practically, this is another barrier to long jump changes, as they would appear to be much more expensive (and less certain) to undertake than local adaptation (i.e., climbing up local peaks). Taken together, these factors may help explain why major business transformations (i.e., attempts at long jump change) are quite rare and not always successful.

On the other hand, some phenomena from the NKCS model track quite closely with the real world of business. Lifecycles characterized by an early profusion of organizations and product variants, followed by the emergence of a few dominant architectures and firms, and then by a period of gradually slowing incremental improvements are observable in many businesses. It is also true that businesses in any given industry coevolve with other firms located on at least seven other linked landscapes: customers (assuming they come from only one other industry), suppliers (from at least four other landscapes supplying capital, people, regulations, and information or materials), substitute products, and complementary products (e.g., Microsoft and Intel). Episodes of punctuated equilibrium that result from coevolution are observable throughout business history (think of the consequences of the invention of the automobile, airplane, or computer). Moreover, as coevolution theory would lead one to expect, in many industries innovations that have resulted in substantial improvements in fitness have first appeared in peripheral regions of the industry landscape. One can also think of coevolving business systems that have exhibited both exceptionally stable behavior for long periods of time (e.g., the regulated natural gas electric power, and telecommunications industries), as well as exceptionally unstable “red queen” behavior (e.g., some parts of the high technology industry). The power laws that describe adaptive walks up local fitness peaks also seem to describe both learning curve phenomena and the diminishing returns associated with many operational improvement efforts. Lifecycles and power laws also help to explain why organizations tend to display more inertia as they get older, particularly in industries that are not dynamic (e.g., those with low C and S).

At the aggregate level, one can observe a number of business behaviors across a range of industries that align quite closely with the self-organization to the edge of chaos phenomenon observed in the NKCS model. As you recall, the key control parameters in that model that determine whether it is operating in equilibrium, chaos, or the phase transition are K (the degree of interconnectivity between elements of an organization's strategy); C (the degree of connectivity between elements of an organization's strategy and elements of organizations' strategy in other industries), and S (the number of other industries with which an organization is connected). Across a large number of industries, we observe the following phenomena over the past decade:

- Reengineering, delayering, downsizing, and other efforts whose effect probably (in NKCS terms) has been to reduce connectivity within these organizations (K);

- Increased efforts to build closer relationships with customers, suppliers, investors, and employees (i.e., to increase C);

- And an explosion in the number of strategic alliances between firms in different industries (i.e., an increase in S).

Graphically, one can see how taken together, these trends suggest the movement of a large number of coevolving systems away from the stable zone and towards the phase transition that generates maximum fitness.

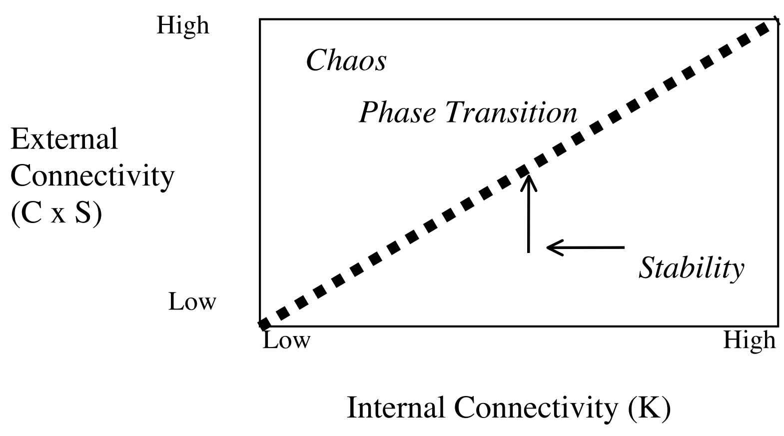

In this regard, the NKCS model leads to a conclusion which, by the standards of the traditional management literature, is quite radical: in a coevolutionary system, any attempt to separate the concepts of environment, strategy, and organization is at best artificial, and at worst dangerous. Far more important than the traditional reductionist analytical tools are a holistic point of view and a talent for recognizing patterns in the behavior of the system as a whole. Rather than planning extended sequences of intentional actions, leaders should instead focus on managing short term performance and medium term connectivity (i.e., K, C, and S) while utilizing feedback to ensure that the system of which they are a part remains in the phase transition state where the highest organizational fitness is possible.

Finally, perhaps one of the strongest endorsements of the validity of a set of conclusions is when two different people using different methodologies arrive at them independently. With that in mind, we will briefly review the findings of an investigation by professors Andrew Van de Ven and Marshall Poole, as published in the Academy of Management Review.[12] Having examined over twenty different theories purporting to explain the process of organizational change, Van de Ven and Poole concluded that they were all some combination of just four basic "motors" or drivers of change.

The first motor of change was the lifecycle: change was driven by each organization following a predetermined path through a sequence of stages (e.g., entrepreneurship, growth, maturity, decline). As we have seen, the NKCS model demonstrates the evolutionary basis for the existence of the lifecycle phenomenon. However, left unspecified by this theory are the events that trigger movement from one lifecycle stage to another.

The second motor of change was termed "teleology": a purpose or goal drives change, as the organization passes through a cycle of defining a goal, formulating plans to reach it, monitoring progress, and changing goals and/or plans on the basis of the results achieved. In the NKCS model, purposeful recombination and local adaptation are the analogs to teleological change. This theory, however, left unspecified the process through which the goal is set.

The third motor of change was termed "dialectical": it assumes that organizations exist in a pluralistic world where oppositions (i.e., thesis and antithesis) repeatedly come into conflict to produce either stability (in the case of a balance of power between them), or change (a synthesis). In the NKCS model, this is referred to as coevolution. However, this theory left unspecified how antitheses arise.

The fourth motor of change was termed "evolution": it assumes that change is produced via a sequence of organizational variation and selection based on competition for scarce resources. The NKCS model makes just the same point. Left unspecified by this theory, however, is the origin of variation.

As you can see, each motor of change is inherently incomplete. Moreover, the four motors operate over different time frames (longer for life cycle and evolution, shorter for teleology and dialectical) and on different levels (within a single firm for lifecycle and teleology, and within a population of firms for dialectical and evolution). Taken together, however, they form a clear picture of the complex adaptive process known as organizational change. As Van de Ven put it, "we speculate that underlying the process of development observed in organizations there exists a relatively simple system of nonlinear relationships among the four motors of change we have identified."

At the theoretical level, then, it appears that there is an underlying logic that governs organizational adaptation. However, because this process is a nonlinear one we cannot use these findings to produce simple Newtonian prescriptions for how to change an organization. At best, we can describe the nature of the system as a whole. In our next newsletter, we will look at a number of specific approaches that attempt to apply these research findings to produce improved organizational performance.

Notes

1. Enhancing Organizational Performance, Daniel Druckman, Jerome Singer, and Harold Van Cott, editors. National Academy Press, 1997

2. This section is drawn from Complexity, Global Politics, and National Security, by Daniel Alperts. National Defense University Press, June, 1997, and Coping With Bounds: Speculations on Nonlinearity in Military Affairs, by Thomas Czerwinski. Published by the Assistant Secretary of Defense for Command, Control, Communications and Intelligence, May, 1998.

3. Ibid.

4. Ibid.

5. Complexity and Creativity in Organizations, by Ralph Stacey. Berrett-Koehler Publishers, 1996.

6. At Home in the Universe, by Stuart Kauffman. Oxford University Press, 1993.

7. Ibid.

8. Ibid.

9. Adaptation on Rugged Landscapes, by Daniel Levinthal. Management Science, Vol. 43, Number 7, July, 1997 and Imitation of Complex Strategies, by Jan Rivkin, Harvard Business School Working Paper, 1997.

10. Rivkin, Op. Cit.

11. Levinthal, Op. Cit.

12. Explaining Development and Change in Organizations, by Andrew Van de Ven and Marshall Poole, Academy of Management Review, July, 1995.

Additional References

- Chaos, by James Gleick. Penguin, 1987
- Hidden Order, by John Holland. Perseus Books, 1995
- Complexity, by Mitchell Waldrop. Touchstone, 1992
- How Hits Happen, by Winslow Farrell. HarperBusiness, 1998
- Leadership and the New Science, by Margaret Wheatley. Berrett-Koehler, 1992
- An Evolutionary Theory of Economic Change, by Richard Nelson and Sidney Winter. Harvard University Press, 1982
- Shifting Paradigms: From Newton to Chaos by Toby Tetenbaum. Organizational Dynamics, Spring, 1998
- “The Self Organizing System”, by David Stamps. Training, April, 1997

- "Complexity, Organization, and Stuard Kauffman's The Origins of Order", by F.H. Westhoff and B.V. Yarborough. Journal of Economic Behavior and Organization (29)

- Last and certainly not least, we strongly recommend visiting the Santa Fe Institute (www.Santafe.edu), which is an extremely rich source of materials on the new science of complexity.

Bristol Partners, Inc. is a general management consulting firm that makes substantial long term equity investments in all its clients. We work with only one company per industry, highly leverage client staff, and help them to understand, integrate, and resolve strategic, organizational, and financial problems. We can be reached on (415) 438-3300, or at info@bristolpartners.com.

Developing Adaptive Organizations, Part II: Practical Applications

In our last issue of Management Insights, we described research findings (largely from the field of evolutionary biology) that seem to have great bearing on the issue of how to design and lead organizations that are highly adaptable and able to successfully cope with surprises. In this issue, we are going to take a closer look at three approaches for putting these findings into practice.

Command, Control and Planning in the U.S. Marine Corps

As defined by the Marine Corps[1] "command and control" refers to "the means by which a commander recognizes what needs to be done, and sees to it that appropriate actions are taken." As the Marines' doctrine publications note, "in war, the defining problem of command and control that overwhelms all others is the need to deal with uncertainty." Historically, military organizations have taken one of two approaches to this problem. As described by the military historian Martin van Creveld, "confronted with a task, and having less information available than is needed to perform the task, an organization may either increase its information collection and processing capacity, or it may redesign the organization and/or the task itself in such a way as to enable it to be performed with less information. The former approach leads to the multiplication of communication channels and to an increase in the size and complexity of headquarters organizations."[2] Best described as "command by plan", this approach to command and control is typically characterized by precise, positive control and highly detailed plans and orders to subordinate units. The object of this Newtonian approach to command and control is to gain certainty, and to impose some degree of predictability on the conduct of war.

Given the nonlinear nature of war, command by plan is almost certain to be an unworkable approach. As the Marines note, "uncertainty in war is not merely a [temporary] environmental condition which can be reduced by gathering more and/or better information. Rather, uncertainty is a natural and unavoidable product of the dynamic of war itself...The fundamental point is than any military action, by its very nature as a complex system will exhibit messy, unpredictable, and often chaotic behavior that defies orderly, efficient, precise control...[Moreover] the widespread belief that information technology will allow us to blow away [this uncertainty] is a dangerous delusion which fails to understand the complex [nonlinear] nature of war."

With this as background, the Marine Corps "starts with a simple model of the command and control process known as the OODA loop, which is an acronym for observation, orientation, decision and action. We first observe the situation -- we take in information about our own status, our surroundings, and our enemy. Sometimes we actively seek information; other times, it is thrust upon us. Having observed the situation, we next orient to it -- by making certain assumptions, analyses, and judgements about the situation to create a cohesive mental image of it. In other words, we try to figure out what the situation means to us and develop 'situational awareness.' Based on this orientation, we decide what to do, whether that decision takes the form of an immediate reaction or a deliberate plan. Then we put the decision into action, which includes disseminating it, supervising to ensure proper execution, and monitoring the results through feedback, which takes us back to the observation phase. Having acted, we have changed the situation, and so the cycle begins again."

"The OODA loop reflects the fact that command and control is a continuous, cyclical process. In any conflict, the antagonist who can consistently and effectively cycle through the OODA loop faster -- who can maintain a higher tempo of operations -- gains an ever increasing advantage. With each cycle, the slower antagonist falls farther and farther behind, and becomes increasingly unable to cope with the deteriorating situation. [This shows that], along with uncertainty, relative speed is the second essential element of effective command and control…The ultimate measure of command and control effectiveness will always be the same: do we act faster and more effectively than the enemy?"

How then do Marines operationalize these principles? They start by thinking of command and control as "an adaptive process in which command is top down guidance and control is bottom up feedback about the effects of the actions that have been taken. It is fundamentally a cycle of reciprocal influence and continuous adaptation. Its objective is not to achieve control in the Newtonian sense, but rather to keep the entire organization surfing on the edge of being out of control, because that is where the system is most adaptive, creative, flexible, and energized."

The key to making this approach work is the "mission order" format used throughout the organization to convey instructions from superiors to subordinate units. In a nutshell, the commander assigns missions, explains their underlying intent, but leaves subordinates as free as possible to choose the manner of their accomplishment.

In simplified form, a mission order contains four key sections, each of which is as brief and simple as possible. First, a description of the situation as it is understood by the person issuing the mission order. Second, a description of the mission itself: the goal to be accomplished, and by when it must be completed.

Third, a statement of the commander's intent: a description of the end state that accomplishment of the assigned goal is intended to produce; in other words, the "why" behind the mission. "Between the mission and the intent, the latter is predominant. while a situation may change, making the specific mission obsolete, the intent is more enduring and should continue to guide initiatives undertaken by those receiving the mission order. An effective intent allows subordinates to act with initiative even in the face of disorder and change."

Fourth, a statement of whom the units receiving the order need to coordinate with. To the extent possible, the required degree of coordination is minimized (ie., K is kept low), and the units involved are left to work out the specifics of such coordination arrangements on their own, guided by their common understanding of the commander's intent.

As described by the Marines, "mission command and control is central to our approach to warfare, because it deals better with the fundamental problems of uncertainty and time. We sacrifice precision and certainty for speed and agility. Mission command and control provides the flexibility to deal with rapidly changing situations and exploit fleeting opportunities. It provides for the degree of cooperation necessary to achieve harmony of effort, yet gives commanders at all levels the latitude to act with initiative and boldness...In fact, it requires them to do so. Being free to act on their own authority, subordinates must accept the corresponding responsibility to take such actions."

On the other hand, "this does not mean that subordinates are free to act without guidance from above. In fact, initiative places a special burden on subordinates, requiring that they always keep the larger situation in mind and act in consonance with their commander's intent."

Nor does delegating authority to subordinates absolve higher commanders of ultimate responsibility. They "have an obligation to frame their orders in such a way that subordinates have sufficient understanding to act in consonance with their intent, while not restricting their freedom of action. [They also have an obligation] not to have a zero defects mentality, which tends to penalize and discourage initiative."

Overall, the Marines believe that "mission command and control has an important psychological effect on the organization. Recognizing what needs to be done and taking the action necessary to succeed is a satisfying experience and a powerful stimulant to human endeavor. People not merely carrying out orders, but acting on their own initiative feel a greater sense of responsibility for the outcomes of their actions and will naturally act with greater vigor. Thus, mission command and control is a source of great strength for the organization, especially in times of crisis." In summary, "rather than increasing the amount of certainty we seek, by mission command and control we reduce the degree of certainty we need to succeed."

What then, in the context of mission command and control, is the function of planning? Marine Corps doctrine states that "we should not think of planning as a scripting process which establishes specific actions to be taken and a timetable for those actions. This approach seeks to narrow possibilities in order to minimize uncertainty and simplify preparations and coordination. Rather, we should view planning as a learning process which helps us to understand how to exploit various possibilities an uncertain future may hold... The sign of a good plan is that it gives you situation awareness, direction, and the flexibility needed to adapt successfully to a changing situation."

Marines believe that they "should plan far enough into the future so that we can maintain the initiative and prepare adequately for upcoming actions, but not so far into the future that plans will have little in common with actual developments. Planning ahead is less a matter of trying to direct events, and more a matter of identifying options and possibilities."

In sum, "the goal of planning is to provide options for the commander to face the future with confidence. The measure of a good plan is not whether it transpires as designed, but whether it facilitates effective action in the face of unforeseen events."

Finally, mission command and control has a clear implication for organizational structure: "Headquarters staff should be kept small. The larger and more compartmented the staff, the more information it will require, and the longer it will take to perform its function."

Competing on the Edge

Competing on the Edge[3] is a recent book by Kathleen Eisenhardt, a professor at Stanford Business School and Shona Brown, a former student of hers. Subtitled "strategy as structured chaos", is intended to help companies compete in "high velocity environments" where change is pervasive and fast. The authors make a number of recommendations for putting into practice ideas grounded in the theory of complex adaptive systems. These ideas are summed up in ten rules:

1. Advantage is temporary. "Managers who compete on the edge understand that competitive advantage is fleeting, and so they focus on continuously generating new sources of advantage."

2. Strategy is diverse, emergent, and complicated. "Strategy is not a single, simple approach to the marketplace. It is a diverse collection of moves that are loosely linked together in a semi-coherent strategic direction. Managers who compete on the edge let strategy emerge, then shape and articulate it. They make a variety of moves, observe what happens, and follow through with the ones that are successful."

3. Reinvention is the goal. Managers who compete on the edge look for opportunities to reinvent their businesses and then let profits follow. They worry about finding new ways to create value, but not necessarily about being the most efficient firm. This shows up in performance measures such as sales of innovative products, market share gains and sales growth relative to leading competitors...The emphasis on reinvention does not imply that profits are forgotten...only that reinvention is the smarter path toward long term profitability."

4. Live in the present. "The present is the most important time frame...The approach to managing today is to minimize structure. Managers who compete on the edge structure their businesses as little as possible...A few strict rules and a few slavishly monitored operating variables...indicate necessary, but minimum structure...These managers use just enough structure to keep things from flying apart."

5. Stretch out the past. Managers who compete on the edge learn more from the past than their counterparts...They keep their product and service platforms in the market longer than others, exploit derivative products more effectively, and extend their offerings into new geographies and customer segments more frequently...Wide use of the past diminishes risk and frees resources to focus on new ideas...However, while using experience, these managers guard against becoming locked into dated competitive models."

6. Reach into the future. "Managers who compete on the edge...[have] a longer time horizon than most others. Driven by a belief that the future is unpredictable, they launch more experimental products and services, create more strategic alliances focused on nascent markets and technologies, and employ more futurists than other firms. Driven by a paranoia that the future is constantly changing, they revisit the future often...Yet at the same time they are not entranced by the future at the expense of the more important present."

7. Time pace change. Managers who compete on the edge pace change in their business with the passage of time as well as by the occurrence of events. They understand that pace is a critical strategic weapon...If they cannot set

the pace, they try to match their tempo of change with the demand of the market and the rhythm of other firms. Managers who compete on the edge understand the power of rhythm to get their businesses into a groove and keep them there."

8. Grow the strategy. "Managers who compete on the edge...do not build all the pieces of their strategy at once...They begin with their current business, then work to incorporate the past and link future opportunities, and conclude with time pacing. These managers never start with the future...they start with the basics of today."

9. Drive strategy from the business level. "The mindset is business level driven strategy and accountability. Business managers control both parts of the strategy: 'where do you want to go?' and 'how are you going to get there?' Managers realize that in high velocity markets, strategy cannot be driven top down."

10. Repatch businesses to markets and articulate the whole. "Any neat match of businesses with markets is all too fleeting in a world where opportunities come and go, collide and divide...This constant flux creates the challenge of continually reexamining the make up of individual businesses and their matches with markets...Continuously realigning businesses with emerging opportunities and articulating and occasionally shaping emergent strategy are the principal responsibilities for managers who compete on the edge...Pattern recognition and the articulation of the essence of patterns are the skills that can be found at the heat of the senior manager's job."

To further elaborate on the patching idea, as described by Stuart Kauffman it is basically coevolution applied at the organism (or organization) level. As he describes it, "the approach is simple. Take a hard conflict laden problem in which many parts interact (eg., high N and high K) and divide it into a quilt of non-overlapping patches (which reduces K). Connect each patch to its immediate neighbors. Then try to optimize each patch. As this process unfolds, the coupling between parts across patch boundaries will mean that finding a good solution in one patch will change the problem being solved by the parts in the adjacent patches...If patches are too few, you will get a stable equilibrium; if there are too many, you will end up with red queen chaos. The goal is to find the number of patches that results in the system operating in the phase transition between stability and chaos."[4]

Crisis and Renewal

Subtitled "meeting the challenge of organizational change", Crisis and Renewal[5] is another recent management book that is based in part on the application of ideas from the world of complex adaptive systems. What makes this book particularly interesting is the background of its author, David Hurst, who was for many years a line manager in a large firm. Unusually in the management literature, Hurst describes change as a continuous cyclical process rather than a sequential linear one. Here is the graphic he uses to describe it:

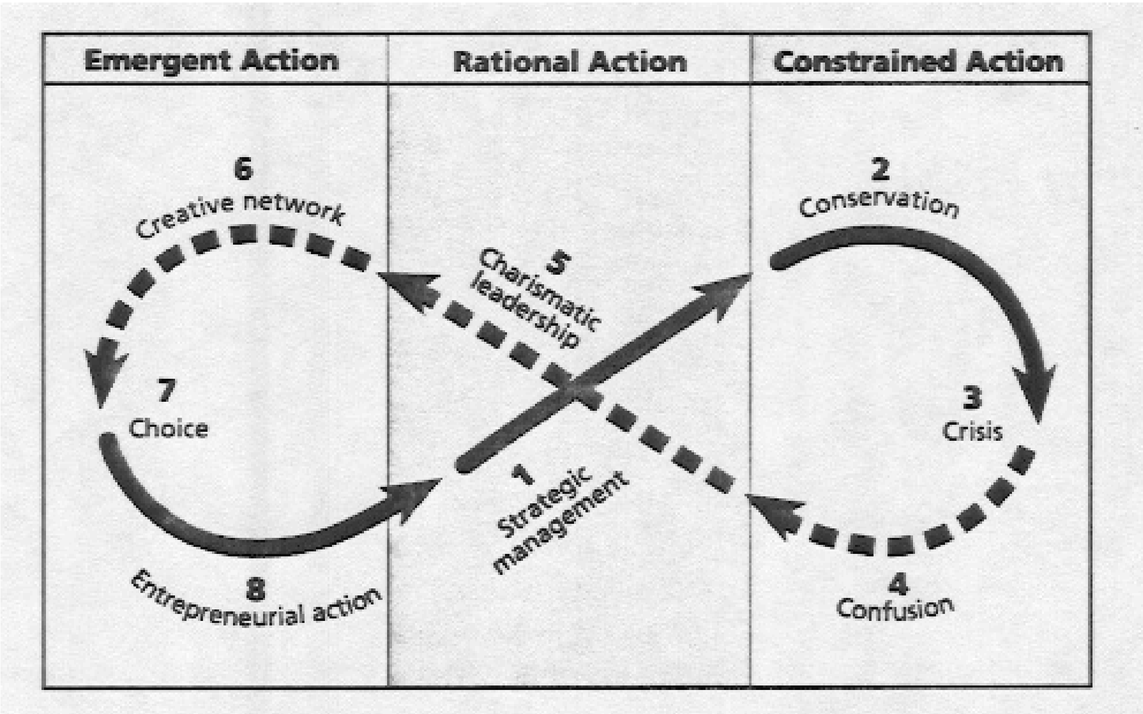

As Hurst describes it, “the first feature of the ecocycle model is that it splits the process of organizational change into two half-loops. The front (performance) loop is the conventional life cycle (points 8 through 3). The back (learning) loop is a less familiar renewal cycle of death and reconception (points 4 through 7).” In terms of the numbering of the phases of the cycle, Hurst made the strategic management phase number one “because that is where most management textbooks start.”

“The performance loop of the organizational ecocycle is identical to the conventional lifecycle... The entrepreneurial phase of the cycle is characterized by the reduction of [successful experiments] into a repeatable formula, which is then extended. Once an organization is established, its conventional life cycle is dominated by a rather lengthy [strategic management phase] during which managers can behave as instrumentally rational actors. During this phase of the cycle, organizational growth may be fairly smooth and linear. The fundamental purposes of the organization are seen as economic and hence calculable. All action is (or should be) a means to economic ends...[During this phase], a manager is continually trying to fit a relatively well defined organization to its environment; the emphasis is on planning and control."

"However, once the growth in overall demand slows, the [conservative phase begins and the] competitive premise becomes 'more of the same'... Managers tend to restrict activities to those that have proved to work. Considerable effort and capital is invested in describing these activities and embedding them in technology and formal organizational procedures to perpetuate their performance ... [During this phase] the organization will specialize, 'stick to its knitting' and emphasize efficiency ... Activities within the system become tightly connected with each other via technology of all kinds, and there is limited variety in the way procedures are performed ... Managerial discretion within the system becomes more and more reduced ... [However], in the process of institutionalizing their successes and pursuing efficiency, however, conservative organizations sacrifice resilience and flexibility and become more vulnerable to catastrophe."

"The next phase of an organization's life cycle is characterized by crises, discontinuities, and wide fluctuations in variables such as sales and prices that have traditionally been stable ... Managers' feelings of fear and uncertainty contrast unfavorably with the feelings of control and even omnipotence that characterized the previous phase of the cycle ... The confusion phase of the cycle breaks the constraints that bound the organization in the final phase of the conventional life cycle (e.g., a compensation system that rewards the wrong kind of behavior, an information system that supplies inappropriate data, and a human resource system that attracts and promotes people who have been selected for their ability and willingness to preserve the status quo, not challenge it)."

"Confusion sets the stage for the emergence of a values-based, charismatic leaders who model the behavior they expect from others. This values based leadership seems to be essential to the attraction of creative people and the development of contexts that nurture innovation and entrepreneurship … Individual efforts, previously self-centered, now become connected with important values and take on new significance ... This allows the formation of a loose network of creative people held together by shared values and an emerging vision of a common purpose. With the emphasis on learning and the options for the future this generates, the renewed organization regains the ability to choose between different courses of action, and moves back into the entrepreneurial phase of the performance loop."

Conclusions

This issue of Management Insights and the one before it have focused on the issue of how to make organizations more adaptive and better able to cope with surprising changes in their environment. While the business world seems awash in unproven "magic bullet" solutions to this problem, we have tried to show that there are, in fact, valid and valuable research findings in this area, as well as a number of approaches to practically applying them. From our extended review of these materials, a number of key conclusions have emerged:

- Uncertainty is not something "out there"; rather it is an inherent part of the coevolving systems in which organizations operate.

- The underlying cause of this uncertainty is the nonlinear nature of these systems, and their consequent tendency toward states of disequilibrium.

- Within this context, organizational adaptation results from not one, but four processes which operate over very different time frames and are subject to varying degrees of management influence.

- Given the nature of the system within which organizations operate, we need to recognize the limitations of many analytical techniques which work best in linear or only slightly non-linear contexts.

- This does not mean, however, that we are doomed to stumble blindly into the future; rather we need to change our expectations about the nature of the foresight we can reasonably expect to develop. While it is not possible to forecast specific outcomes over anything more than a very short time frame, it is possible (eg., through use of simulation or scenario techniques, or the application of power laws) to develop an understanding of how coevolving systems as a whole function, the boundaries that define the range of possible futures they could produce, and the threats, opportunities, and options contained within those boundaries.

- There are also approaches that can help us to effectively manage our organizations in the face of this irreducible uncertainty. In the short term, we can use approaches like mission command and control to maximize organizational initiative and adaptiveness within a goal driven framework. In the medium term, we can manage connectivity, both internally (eg., through patching) and externally (eg., through the structure of our alliances and intensity of our relationships with customers, suppliers, and producers of complementary products and services). In the long term, we can expect to pass through renewal cycles, and take steps to both induce them and to lead them effectively.

- Our approach to control needs to recognize that it is best focused on the system as a whole, rather than on its parts, and is best accomplished via a "mutual influence" process based on guidance and feedback.

- Finally, all of the above make clear that at the personal level, one of the most important functions of leaders is to both absorb and help subordinates to cope with the relatively high levels of personal anxiety that can be created by the irreducible uncertainty that is a dominant feature of the systems in which their organizations operate.

Notes

1. Unless otherwise noted, the quoted material in this section is drawn from Marine Corps Doctrine Publications Plannning, Intelligence, and Command and Control, all published by the United States Marine Corps, and from Complexity, Global Politics, and National Security, by Daniel Alperts. National Defense University Press, June, 1997.

2. Command in War by Martin van Creveld. Harvard University Press, 1985.

3. Competing on the Edge, by Kathleen Eisenhardt and Shona Brown. Harvard Business School Press, 1998.

4. At Home in the Universe, by Stuart Kauffman. Oxford University Press, 1993.

5. Crisis and Renewal, by David K. Hurst. Harvard Business School Press, 1995

Bristol Partners, Inc. is a general management consulting firm that makes substantial long term equity investments in all its clients. We work with only one company per industry, highly leverage client staff, and help them to understand, integrate, and resolve strategic, organizational, and financial problems. We can be reached on (415) 438-3300, or at info@bristolpartners.com.