Emerging Strategies for a Chaotic Environment

Ralph Stacey

The possibility that any member of a system can foresee its future depends upon the dynamical properties of the system. Today's predominant management view is that 'human agents' in an organisation can foresee the future outcomes of their actions sufficiently well jointly to intend comprehensive organisational outcomes. This predominant view is based on the metaphor of an organisation as a machine or as an organism adapting to a given environment. The new science of complexity leads us to see organisations as complex adaptive systems. Such systems are creative when they occupy a space at the edge of disintegration, and here their specific futures cannot be foreseen. The price we pay for creativity and free will is an inability to foresee and intend future outcomes.

IN THE LATE 1970s I was the Corporate Planning Manager for a major company. I recall that colleagues and I wrote papers drawing attention to increased automation and the rapid development of information technology. We foresaw that people would work more at home, that paper forms of communication would decrease and that we would all have much more leisure time. We then speculated on what all of this might mean for our business. We were not alone in these speculations: I recall that most expert commentators were saying much the same thing and visions of machines doing the work while we puzzled about what to do with our leisure time were quite the rage then. We now know, of course, that far from puzzling about leisure activities, the majority of us are working even longer hours under even greater pressure, while a much increased minority is trying to cope with the effect of imposed leisure on a full-time basis. We also now know that the flow of paper has increased and that the great majority of those at work are not working from home. Here our attempts at foresight failed us completely, but this failure did not matter much because we did not take any action on these speculations about the future back in the 1970s simply because, at that time, we could not jointly make up our minds on what it all meant in terms of specific next actions.

On the other hand, a colleague working at that time at another major company was pointing out that the new house purchasing age cohort was set to decline in the 1990s and that this would have a depressing impact on the housing market. Once again many were pointing to the potential impact of declining numbers in the younger age cohorts and we can now see that some of their predictions have come true. Here some attempts at foresight succeeded, but there seems rather little evidence of this success producing more effective action: we seem rather surprised by the problem of recent house buyers trapped in their homes by negative equity and at a loss as to what to do about it.

We can of course multiply examples both of our ability and our inability to foresee the future, as well as our ability and our inability to act effectively upon that foresight or lack of it. For example, in the early 1980s the government seemed convinced that it could foresee enough to set out a medium-term financial strategy. After years of drawing diagrams that persistently showed huge divergences between Public Sector Borrowing Requirement forecast and outturn, the medium-term strategy was abandoned. On the other hand, National Health policy papers pointing to the implications for funding and service quality of major increases in older age cohorts showed con-

Pergamon

S0024-6301(96)00006-4

Long Range Planning, Vol. 29, No. 2, pp. 182 to 189, 1996

0024-6301/96 $15.00+0.00

siderable foresight. The policy response to this ability to foresee is, however, the focus of a great deal of contention at the present time, as we all know.

These examples hint at three important factors that must be borne in mind as we try to gain some insight into the nature and importance of foresight in an organisational setting.

First, such foresight is possible only if the systems we constitute when we interact with each other (the groups, organisations and societies we form) are not only deterministic but also produce determinate outcomes. In other words, agents in a system can exercise foresight only when their system is one in which there are 'in principle' links between a cause and its effects, and in which a cause produces only a limited number of effects making it practically possible to trace the connections between an action and its outcomes over a reasonably long period of time and across a reasonably large number of links in the networks we form. Foresight is not possible to the extent that the future is not determinate but open-ended, i.e. to the extent that causes/actions can produce an infinite number of effects/outcomes. Foresight is possible when events are some kind of repetition of the past or the present, but not when events are completely new. When we talk about foresight we must be explicit about what kind of system we are hoping to practise that foresight in.

Second, our interest in foresight is intimately intertwined with our concern about control. We look for, and greatly value, foresight precisely because it makes it possible for us, as an individual, a group, an organisation or a society, to be 'in control' of our destiny. For example, if we can foresee what new products people will want then we can stand ready to provide them and so be 'in control' of our destiny as an organisation. Given the possibility of foresight whether we are 'in control' or not therefore depends upon how competent and how well behaved we are and these again are matters within our control, in principle at least. Without foresight, however, in a system with an open-ended indeterminate future, we cannot be 'in control' of our destiny: we can only shape and accept what emerges out of what we jointly do. When we talk about foresight we must make explicit the connections we are making with control.

Third, our interest in, and our response to, the twin matters of foresight and control does not lie purely at a rational, intellectual level. Questions to do with control, open-endedness and predictability touch us all in a very direct and personal way: our response to these matters is unavoidably emotional and in that sense irrational. Open-endedness leads to a poor ability to foresee, hence to an inability to be 'in control', and therefore to real fears about disintegration and anarchy that generate high levels of anxiety in all of us. The easiest and most readily available defence against such anxiety is to deny the possibility or importance of open-endedness, unpredictability and lack of personal ability to control outcomes of actions. When we talk about foresight and control therefore we must be very much aware of existential anxiety and how we defend ourselves against it, because this will affect the arguments about foresight that we will be willing to listen to.

This article will focus primarily on the first point, the nature of the system we constitute when we interact with each other, but some reference will be made to the other two points at the end of the article.

Foresight and System Dynamics

A novel and, I suggest, potentially very fruitful way of understanding the dynamics of organisational life is provided by the new science of complexity.[1-11] This science studies the fundamental properties of complex adaptive systems. This article will now take you on a diversion into understanding the nature of complex adaptive systems and then use what we discover to take a fresh look at the possibility of foresight in human systems.

Complex Adaptive Systems

A complex adaptive system consists of a number of components, or agents, interacting with each other according to sets of rules called schemas in such a manner as to improve their behaviour and thus the behaviour of the system which they comprise. In other words, in a complex adaptive system agents interact in a manner that constitutes learning. Complex adaptive systems operate in an environment that consists of other complex adaptive systems so that a system and its environment together form a co-evolving supra-system which, in a sense, learns its way into the future. In this co-evolving supra-system, individual subsystems or agents learn i.e. they alter their schemas during their own individual lifetimes and the system as a whole learns, i.e. it evolves. An ant is one example of a complex adaptive system in that it employs a set of rules to do with locating food and laying trails to the food that other ants can learn to follow. The ant and indeed the colony of which it is part are both complex adaptive systems, as is the ecology with its plants and ant eaters that the colony is located in.

Complexity scientists have so far focused their attention primarily on biological systems, on the evolution of life and the behaviour of chemical and physical systems. The science has been developed by mathematicians and computer scientists, physicists and chemists such as Noble Prize laureates Murray Gell-Mann and Illya Prigogine; Stuart Kauffman, Christopher Langton and John Holland at the Santa Fe Institute in New Mexico; Brian Goodwin at the Open University. However, what these scientists are uncovering is the general properties that apply to all complex adaptive systems no matter where they are found and when you think about it, each of us and each of the organisations we belong to are quite clearly complex adaptive systems.

Each individual has a brain that is a complex adaptive system in which neurons are the agents. Each of us has a mind that is a complex adaptive system in which symbols and images are the agents. Together, in a group we constitute a complex adaptive system in which we, or more accurately our minds, are the agents. All organisations are complex adaptive systems in which groups and individuals are the agents and these organisations interact as agents to form national economic, societal and political systems. Those national systems in turn interact to form a global one, which interacts with natural systems to form an interconnected ecology. All are complex adaptive systems, one fitting into the other. The fundamental dynamic properties that complexity scientists have identified for complex adaptive systems in general must, therefore, apply in some way to human systems in particular, unless we can show that specifically human characteristics nullify them. I will argue later that there is nothing about consciousness, human free will, determination or power of thinking and acting that can overcome the basic properties of the system we are and the system we have no choice but to constitute when we interact with each other.

How do complexity scientists study complex adaptive systems and what do such studies tell us about them?

Method of Studying Complex Adaptive Systems

The most important method of studying complex adaptive systems is that of simulating their evolution on a computer. We can quite plausibly think of members of an ant colony, or members of a human organisation, as agents whose behaviour is driven by a set of rules. In other words it is perfectly reasonable to regard an ant, and in the first instance each of us, as:

- a set of operating rules driving the performance of tasks necessary for survival; plus
- a set of rules for evaluating that operating performance; and
- a set of rules for changing both the operating and evaluating rules, i.e. for learning. These rules might well involve cross-fertilisation between an agent's existing set of rules and the rules of other agents, a kind of mating between the rules of different agents.

A computer program is a set of operating rules and instructions and it is perfectly possible to add a set of rules for evaluating those operations, and for changing the rules of operation and evaluation in the light of their performance, which could quite easily consist of rules for finding another computer program for a mate and then copying part of each set of rules to create a new program. Such a program could quite realistically be regarded as an agent if some of its rules required it to examine the state of other computer programs and adjust its rules in the light of the results. In this way we could build up a population of computer programs that could interact with each other, breed, and so evolve and learn. We would then have a complex adaptive system in our computer consisting of a collection of agents, each of which is a computer program.

Furthermore, each computer program would be made up of a bit string, a series of 0s and 1s, because that is how computer programs are coded. The analogy with our rules of conduct is still very close because our rules of conduct are coded, not into 0s and 1s it is true, but into symbols taking the form of words and images and, of course, the analogy with our bodies is even closer since our physical code consists of genes that can be on or off, more or less. It is therefore quite legitimate to use such simulations to try to uncover the fundamental properties of complex adaptive systems everywhere, whether physical or mental. We will, of course, have to check whether human characteristics such as consciousness, whatever that is, and emotion and so on alter the conclusions we reach, but as a first step the simulations on a computer may, and I suggest do, yield some very important insights.

The Properties of Complex Adaptive Systems

Those who have developed the study of complex adaptive systems have been most interested in the analogy between the digital code of computer program agents and the chemical code in the genes of living creatures. One of their principal questions has been: if in its earliest days the earth consisted of a random soup of chemicals, how could life have come about? You can simulate this problem if you take a system consisting of computer programs with random bit strings and ask if they can evolve order out of such random chaos. The amazing answer to this question is that such systems can indeed evolve order out of chaos and even more amazingly, this chaos, or mess, is essential to the process.

Contrary to some of our most deep-seated beliefs, mess is the material from which life and creativity are built and it turns out that they are built, not according to some prior design, but through a process of spontaneous self-organisation that produces emergent outcomes. If there is a design, it is the basic design principles of the system itself—a system that produces patterns in behaviour. The system is a network of agents driven by iterative non-linear feedback to produce unknowable outcomes which do have a pattern. There is inherent order in complex adaptive systems simply waiting to unfold through the experience of the system, but no-one can know what that experience will be until it does unfold in real time.[12] There are reasons why, in certain conditions, agents interacting in a system can produce not anarchy, but creative new outcomes that none of them ever dreamt of, if they are left to self-organise in what looks like a mess most unlikely to contain within it an implicate order.[13,14] It is simply not true that if we cannot know the outcome and if no-one can be 'in control' we are doomed to anarchy. On the contrary, these are the very conditions required for creativity, for the exciting journey into open-ended evolutionary space with no fixed, predetermined destination. The whole universe, it seems, is indeed lawful and yet it has freedom of choice—the price for the freedom of choice is an inability to know the final destination and an inability to be in control of the journey.

You can see now the profound connection between the nature of complex adaptive systems and our ability to foresee the future. Before this connection is made more explicit, however, consider an example of a complex adaptive system simulation.

An Example of a Complex Adaptive System Simulation

In a simulation called Tierra,[15] the programmer:

- ☐ Determined that the evolution of this system would initially be driven by random mutations.
- ☐ Set some general task for the agents and the requirement that they should seek to survive as a species, i.e. that they should replicate.
- ☐ Set the initial replication rules, a schema common to all agents, requiring agents to clone.
- ☐ Defined a fitness function, a schedule of points awarded for task performance according to which agents were allocated computer time to carry out their replication rules.
- ☐ Introduced a constraint taking the form of scarce computer time which worked as follows. Agents were required to post their locations in the computer memory on a public notice board. Each agent was then called upon, in turn according to a circular queue, to receive a slice of computer time for carrying out its operational and replication tasks.
- ☐ Introduced a further constraint on agent life span. Agents were lined up in a linear queue according to their age and a 'reaper' lopped off some of these, generally the oldest. However, by successfully executing their programs, and so increasing their fitness, agents could slow down their move up the linear queue while flawed agents rose quickly to the top.

The agents, as part of a system, were then left to learn how to perform the task i.e. develop their own survival strategies.

The simulation was set off by introducing a single agent consisting of 80 instructions. Within a short time the computer memory space was 80% occupied by agents, but then the reaper took over and prevented further population growth. After a while agents consisting of 45 instructions appeared but they were too short to replicate. They overcame this problem by attaching themselves to longer agents and borrowing some of their code in order to replicate. This strategy increased their fitness because in their normal operations they needed less computer time and so had more time for replication using borrowed code. In other words, parasites had emerged.

However, if the parasites destroyed too many hosts in order to replicate, they destroyed their own ability to replicate and so declined. In the simulation the parasites suffered periodic catastrophes. One of these catastrophes occurred because the hosts stopped posting their positions on the public notice board and in effect hid so that the parasites could no longer find them. Some hosts had thus developed an immunity to parasites by using camouflage as a survival strategy. But, in hiding, the hosts had not retained any note of their position in the computer memory, so they had to examine themselves to see if their position corresponded to the position being offered computer time before they could respond to that offer. This increased the time they needed for normal operation and reduced the time they had available for replication. However, although not perfect, the strategy worked in a good enough way so that the parasites were nearly wiped out.

Then, however, the parasites developed their own memories and did not need to consult the public posting board. Once again it was the parasites turn to succeed.

Later, hyperparasites appeared to feed off the parasites. These were 80 instructions long just like hosts, but they had developed instructions to examine themselves for parasites and feed off them by diverting computer time from them. These hyperparasites worked symbiotically by sharing reproduction code—they could no longer reproduce on their own but required cooperation. Cross-over replication had thus emerged spontaneously as a strategy for survival without anyone programming it in.

This 'sexual' cooperation was then exploited by opportunistic mutants in the form of tiny intruders who placed themselves between replicating hyperparasites, and intercepted and used hyperparasite code for their own replication. The cheaters could then thrive and replicate although they were only 27 instructions long. Then the hyperparasites found a way to defeat the cheats, but not for long.

What the Simulation Reveals

This system produces periods of apparent stability followed by upheavals as particular strategies for the survival of both hosts and parasites emerge from spontaneous self-organising processes that appear to be very close to market competition. The programmer is not introducing these periods and neither is an agent taking over and formulating any strategy for coping with them. The strategies are emerging unpredictably in a co-evolutionary arms race occurring in a dynamical, somewhat disorderly environment, driven partly by chance. First the strategy is small size, but then parasites change the rules and the most successful strategy becomes feeding off others. Then the hosts change the rules and the better strategy is camouflage. But the parasites change the rules of the game again and the best strategy becomes the development of a local memory. Competition and conflict emerge and the evolution of the system is driven by agents trying to exploit each other, but the game can go on only if neither side succeeds completely or for long in that exploitation.

Furthermore, this system produces diversity in a spontaneous, emergent way that has not been programmed in, and this diversity is vital for the continuing evolution of the system and its ability to produce novelty. Through an internal process of spontaneous self-organisation, the system produces an emergent predator–prey dynamic behaviour that is paradoxically both cooperative and competitive. In a similar manner the system spontaneously produces cross-over replication, or cross fertilisation, again an activity that is both cooperative and competitive at the same time. Adaptive systems learn as a whole system and the introduction of predator–prey and cross-fertilisation dynamics enhances the learning capacity of the total system in the sense that it is capable of generating more novelty and complexity. Both predator–prey and cross-over replication introduce a tension between cooperation and competition, a kind of disorder, but one that clearly assists system-wide learning. Simpler systems that rely simply on random mutations and competitive selection alone cannot produce the same degree of diversity and it is critical levels of diversity that enhance further learning.

A considerable number of simulations and analyses of non-linear feedback networks have shown how this kind of spontaneous self-organising behaviour producing creative emergent outcomes occurs when the control parameters driving a system reach critical points. When information flows, diversity of agent schemas and degree of connectivity between agents is low, the system operates in a stable zone to produce regular predictable behaviour moving toward an equilibrium position—here agents can in principle exercise perfect foresight. When information flows are very rapid, agent schemas very diverse and connectivity between agents very rich, the system moves into an unstable zone where it disintegrates—here too, agents can in principle exercise perfect foresight. At the edge of system disintegration there is a paradoxical state of bounded instability—behaviour is both stable and unstable at the same time. It is in this state at the edge of disintegration that systems are capable of infinitely variable, novel, creative behaviour. In this creative state system-wide outcomes emerge without prior sharde 'intention' on the part of the agents of the system. Here, agents cannot foresee specific long-term outcomes—they can only foresee and recognise qualitative patterns of behaviour.

But does the very nature of being human not alter these fundamental dynamics? Is it possible for human nature so to change the dynamics of human systems that humans can be creative in an intentional, harmonious way? Or have they no option but to rely on spontaneous self-organisation to produce unforeseeable, emergent outcomes?

The Impact of Being Human on Complex Adaptive System Properties

In human systems agents are:

- Affected by emotion and aspiration, inspiration and anxiety, compassion and avarice, honesty and deception, imagination and curiosity: a dynamic of inspiration and anxiety.
- Able to select their own individual mental purposes rather than shared ones for priority, a reflection of the basic struggle all humans have between being themselves and conforming sufficiently to group requirements so as to belong. This aspect of human behaviour can be summarised as the dynamic of conformity and individualism.
- Impacted by power differentials between agents: the leadership–followership dynamic which reflects the basic human tendency to take on omnipotent, omniscient, dominant roles at some times and submissive, dependent roles at other times.
- Capable of systemic thinking, i.e. of observing, reflecting upon and altering behaviour according to their perceptions of the operation of the whole system of which they are a part. This amounts to an ability to reflect upon themselves and take up the role of both participant and observer. It is the property of consciousness and self-awareness.

Consider the first two tensions, those between inspiration, anxiety and its containment; and conformity and individualism. These peculiarly human aspects of agents can only make interactions more complex and unpredictable. These tensions are the paradoxes of human existence, definitions of what it means for humans to be in the boundedly unstable creative space. These factors cannot overcome the sensitivity of complex systems to tiny changes and thus unforeseeable outcomes since they are themselves paradoxes and mechanisms for escalating tiny changes.

However, surely some agents can think more systemically than others, and then acquire more power and so remove the connection between disorder and creativity. Perhaps creativity in the rest of nature has to rely on some kind of mess, but is it possible for powerful human leaders to make things different in human systems so that creativity can come about in an orderly planned manner? The answer to this question seems to be 'yes' and 'no'. Take the 'no' part first.

A powerful agent could remove the need for disorder in the creative process altogether only if that agent could consistently foresee the future. But no matter how powerful, determined or brilliant any agent is, that agent will be unable consistently to foresee the future because it is being determined by others in a very complex system of cooperative and competitive interactions. There will be other powerful agents in competing systems trying to do exactly what the powerful agent in the first is trying to do, namely foresee and be 'in control' of the future. They will make it impossible for each other to succeed in their endeavour. Since all organisations are parts of a competitive–cooperative suprasystem and all wish to survive, they will keep changing the rules for survival, and this will make the outcomes of actions unpredictable. Now take the 'yes' part of the answer.

A few powerful agents could face the disorder required to be creative and shield followers from it. All the complex learning required to produce creative new strategies for survival would then have to be performed by the small number of powerful agents, while the others simply carried out instructions until their schemas were changed by the most powerful. One would expect this to reduce the capacity of the system as a whole to learn, but on the other hand this might not be the case if anxiety levels associated with complex learning and creativity were reduced by this use of power differentials. Nevertheless, it does seem likely that a small group of powerful agents would have more difficulty in escaping from maladaptive learning behaviours, and there will be a psychological price to pay when a few hold, on behalf of the many, all the creativity and the anxiety it provokes.

What we see then is that power differentials cannot remove the fundamental dynamics of a non-linear feedback system, but they will affect its learning capacity in one way or another. In fact the powerful can push their whole system away from the creative space into the stable zone by setting up systems of behaviour for other agents and inspiring or forcing them to obey. Such systems would, however, be incapable of novelty and eventually fall to rivals who change the rules of the game—perhaps the centrally planned societies of the former Eastern Bloc are evidence of this.

In general, therefore, the peculiarly human nature of agents in organisations does not provide grounds for doubting the applicability of the general properties of complex adaptive systems to organisations. What the peculiarly human features do seem to add is further complexity, making the operation of human systems more complex and unpredictable rather than less so.

We can now be quite specific about the nature of foresight in that complex adaptive system which is a human organisation or society.

Links Between Specific Action and Long-term Outcome are Lost

Causal links between specific actions and specific organisational outcomes over the long term disappear in the complexity of the interaction between people in an organisation, and between them and people in other organisations that constitute the environment—small changes can escalate to have unforeseeable outcomes. It follows that specific long-term outcomes cannot be foreseen and hence cannot be intended by people in an organisation—creative and innovative outcomes can only emerge. Members of an organisation, no matter how intelligent and powerful, will be unable to predict the specific long-term outcomes of their actions. They may specify any long-term state they wish to, they may have any dream, fantasy or vision they like, but they will never be able to determine the sequence of actions required to actualise it. Only when their organisation is operating in the stable zone, only when they are conducting ordinary management to reinforce what they already do well,[16,17] will they be able to intend long-term outcomes and then only if rivals stay stable enough for long enough. The conclusion is clear: we can actualise intended long-term outcomes only by chance.

Short-term Outcomes are Predictable

However, because complex adaptive systems are the product of their precise history and because it takes time for small changes to escalate in such systems, their short-term behaviour is predictable.

Critical Values of Control Parameters and the Space for Creativity

The general dynamic progression of any complex adaptive system, including an organisation, is determined by the state of its control parameters. At certain critical points, those control parameters cause a system to occupy what we might think of as a space for creativity. In principle we should be able to identify the conditions required for an organisation to occupy the space for creativity and, in principle, we should be able to predict whether it will or not. The most important of these control parameters for human systems is firstly, the level of anxiety that can be contained and secondly, the degree of power difference and the manner in which that difference is used. The more an organisation is designed to assist members to contain rather than avoid high levels of anxiety, the more power differences are exercised so as to assist members to engage in rather than avoid self-reflection, the more it will be possible to work creatively despite high levels of the other control parameters: rates of information flow, large differences between schemas and rich connectivity between people.

What this implies is that the emphasis on managing long-term specific outcomes is completely misplaced. These cannot be managed but what can be managed is the control parameters, the containment of anxiety, the use of power, the flow of information, the degrees of difference that are tolerable, the extent of the connections across organisational networks. While it is impossible for managers to intend and plan long-term creative outcomes, they can intend and plan to occupy the space for creativity by operating on the control parameters outlined above. From this perspective, managers still need strategic plans, but they relate not to outcomes and actions to achieve them, but to methods of managing anxiety, power, difference and connectivity.

The Predictability of Archetypal Patterns

The simulations of adaptive non-linear feedback networks show how we can predict general, qualitative archetypes even though we are not able to predict the specific actualisations of those archetypes. The Tierra rules produce emergent forms of behaviour of the predator–prey and cross-fertilisation type. There is no basis for predicting from the rules that the agents in the system will behave in this way: their behaviour is emergent. However, once we have run the system we can then safely predict that if we run that set of rules again we will observe the emergence of similar behaviour—those rules contain within them predator–prey and cross-fertilisation behaviour and iteration causes it to emerge. What we can predict, once we know about the emergent pattern, is not the specific pattern that any run of the simulation will display, but simply the fact that there will be predator–prey and cross-fertilisation dynamics. In other words we are predicting the archetype, but the actualization, depending as it does upon precise experience, is unpredictable.

This is the kind of predictability that is possible for human systems too. We can usually predict that certain individuals will become angry in certain conditions but we cannot normally predict just how that anger will unfold. We can predict that in certain conditions a group of people will display highly dependent behaviour but we cannot normally predict just what form that will take. We can predict that under certain pressures an organisation will decentralise its operation but we cannot normally predict just how this will unfold.

This suggests that we might experience more success in predicting the behaviour of organisations if we focus on what kinds of archetypal behaviour tend to be produced by what general kinds of schemas, rather than trying to forecast the specific outcomes of specific actions.

Control and Anxiety

Today's dominant metaphors for organisational systems originate in Newtonian physics leading to the notion that an organisation is akin to a machine with outcomes that managers can foresee and remain in control of. Or organisations are conceived of as organisms adapting to a given environment through the mechanism of competitive selection in the manner postulated by Adam Smith and Darwin.[18] From these perspectives, organisations are moving toward one of a limited number of equilibria, and the management choice is which of these to move toward, unless the organisation is incapable of doing so due to inertia or resources that are too specific. The ability to predict and control is then seen as the key competence of management.

This belief that leaders can foresee the future and are in control of organisational movement toward a more or less known outcome, contains the anxiety generated by the turbulent uncertain conditions in which we lead our lives. Complexity science exposes the simplistic nature of these kinds of assumption about system dynamics. It makes clear that we cannot foresee creative outcomes nor can we retain control of the development of creative organisations over long time frames. This exposes us all to anxiety, and the first defence is to deny the explanations that so expose us. For this reason we can expect the perspectives of complexity science to be resisted by many managers and policy makers.

Conclusion

The ability of people in groups, organisations and societies to exercise foresight depends upon the dynamics of these groups, organisations and societies. This article argues that they are all complex adaptive systems. Such systems are not deterministic systems but rather adaptive systems—they learn their way into an open-ended evolutionary space and through such co-evolution they create their own futures. The future of creative adaptive systems is thus not determinate.

Complex adaptive systems produce order of a changeable and diverse kind that comes about in a spontaneous, emergent way. Such order has not been programmed in and there is no blue print, grand design or plan. There is neither mission statement, nor vision, nor even a charismatic leader in sight. Furthermore this spontaneous self-organising activity, with its emergent order, is vital for the continuing evolution of a system and its ability to produce novelty. However, what form that order takes i.e. the global pattern of behaviour, the system-wide strategies, cannot be predicted from the rules driving individual agent behaviour. In that sense the system is disorderly. Such systems continuously operate in states of bounded instability: periods of stability are followed by upheavals as organisations pursue different strategies for survival.

The evolution of life in the universe, and life in organisations, does not occur primarily through random mutations selected for survival by the forces of competition, but primarily through an internal, spontaneously self-organising, cooperative process that presents orderly forms for selection by the forces of competition. Selection is not made by freely operating competition that chooses between random little pieces, but by a competitive process constrained to choose between new forms emerging from a cooperative process. Life in the universe, and life in organisations, arises from a dialectic between competition and cooperation, not from an unconstrained free market! Nor is the future being determined by any kind of foresight and the comprehensive long-term planning that such foresight makes possible. In fact, complex adaptive systems learn their way into the long-term future and because they do this, they have free will and are capable of creativity. The price of free will and creativity is that agents in creative, free systems cannot have much in the way of foresight and hence cannot be 'in control'. Instead the system produces emergent order and is controlled by spontaneous self-organisation.

This article is drawn from R. Stacey, *Complexity and Creativity in Organizations*, Berrett-Koehler, San Francisco, CA (1996).

References

1. J. Casti, *Complexification: Explaining a Paradoxical World through the Science of Surprise*, Harper Collins, London (1994).
2. J. Cohen and I. Stewart, *The Collapse of Chaos: Discovering Simplicity in a Complex World*, Viking, New York (1994).
3. Gell-Mann, *The Quark and the Jaguar*, W. H. Freeman, New York (1994).
4. B. Goodwin, *How the Leopard Changed its Spots*, Weidenfeld and Nicolson, London (1994).
5. S. A. Kauffman, *At Home in the Universe*, Oxford University Press, New York (1995).
6. G. C. Langton (ed.), *Artificial Life III, Santa Fe Insitute Studies in the Sciences of Complexity*, Vol. XVII, Addison-Wesley, Reading, MA (1994).
7. G. C. Langton, Computation to the edge of chaos: phase transitions and emergent computation, *Physica*, **42D**, pp. 12–37 (1990).
8. S. Levy, *Artificial Life*, First Vintage Books, New York (1992).
9. G. Nicolis and I. Prigogine, *Exploring Complexity: An Introduction*. W. H. Freeman, New York (1989).
10. S. Rasmussen, C. Knudsen, R. Feldberg and M. M. Hindshol, The Core-world: Emergence and Evolution of Cooperative Structures in a Computational Chemistry, *Physica*, **42D**, pp. 111–134 (1990).
11. M. M. Waldorp, *Complexity: the Emerging Science at the Edge of Chaos*, Simon and Schuster, Englewood Cliffs, NJ (1992).
12. D. Bohm, *Wholeness and Implicate Order*, Routledge and Kegan Paul, London (1980).
13. F. A. Hayek, *Law, Legislation and Liberty*, Routledge and Kegan Paul, London (1982).
14. F. A. Hayek, *Individualism and Economic Order*, University of Chicago Press, Chicago, IL (1948).
15. T. S. Ray, in G.C. Langton, *An Approach to the Synthesis of Life*, pp. 371–408 (1992).
16. R. Stacey, *Strategic Management and Organisational Dynamics*, 2nd ed, Pitman, London (1996).
17. R. Stacey, The science of complexity: an alternative perspective for strategic change processes, *The Strategic Management Journal*, August (1995).
18. D. Parker and R. Stacey, *Chaos, Management and Economics: the Implications of Nonlinear Thinking*, Hobart Papers 125, Institute of Economic Affairs, London (1994).