

Evolution, Emergence, and Learning in Complex Systems

Peter M. Allen & Mark Strathern

Is there a “thermodynamics” that characterizes change in social and organizational networks? How do they evolve and change over time? Is there some “overall trend” in something (specific entropy production?) associated with their evolution, and hence with the passage of time? The answer cannot be too simple, because we know that civilizations rise and fall, most species that have existed on the planet are extinct, and most business organizations fail eventually. But has something been accomplished nevertheless? In other words, is there some overall indicator, some macroscopic variable that is changed in one particular direction over time? If there is, then we could possibly interpret some actions or strategies as working with history, and others as fighting against it. Without a “scientific” theory of change and transformation, social and organizational change could only be driven by trial and error and by people’s accumulating experience and confusion. However, this cannot be enough to prove anything about “progress,” since it would have no theoretical framework within which to make such a claim. It would be purely subjective and a matter of opinion and taste.

Today the situation has changed, because we do now have a scientifically based theory of change and transformation, namely complexity theory or complex systems science (Nicolis & Prigogine, 1977). It is therefore legitimate to ask whether this changes anything, whether it allows us to know how to “improve” system adaptation and learning, and whether this in turn provides us with any simple indicator of overall progress. What may be more important for the topic of complexity and organization, however, is that it provides a new basis on which to consider

the link between any overall goals, desires, and motivations of people and organizations and the strategies and behavior that might help to bring them about. It clearly offers a new basis on which to consider knowledge, strategy, and action.

KNOWLEDGE

In considering these questions, the first step we must take is to reflect on the nature of knowledge and how it is informed by complexity science. The first remark must be that it surely modifies any “grand vision” of science as being an accumulating story of successively more general truths. While this idea was a valid basis for scientific progress in its classical phase, we now see that such an idea is simplistic when it comes to dealing with evolved, organized entities and structures. The theory of evolutionary complex systems underlines the fact that there are really *two kinds of knowledge*:

- ❖ One is the traditional, classical type of knowledge concerning the physical laws of nature that are in no way affected by us knowing them. For example, Newton’s law of gravity allowed all sorts of predictions and actions to be taken, but these laws themselves were not affected by this. They also appear to be (as far as we can tell) “eternal” and “unchanging,” and therefore science could be seen as a process of knowledge accumulation.
- ❖ However, knowledge concerning people’s behavior, or the values placed on something, is changed by the “knowing.” So either by communicating the knowledge or by taking action in response to it, the situation and the knowledge of that situation are changed. In human and social systems then, much knowledge is simply the internal cognitive patterns that have been generated linking stimulus to a person’s particular response. As these patterns change and spread through the system, they undermine their own validity, requiring the regeneration of new patterns of stimulus and response. Subjective knowledge is only part of the evolution of the system.

The difference between knowledge that is independent of its knowing and knowledge that isn’t reflects two different situations, passive and active. A mechanical system is passive in that it merely runs according to the laws of interaction that govern it. But, human and indeed biological systems are the result of an evolutionary process in which an “active”

behavior has been selected for by differential success in the past. In other words, the structures that have grown preferentially have internal mechanisms that perceive their surroundings in some way, and respond to those perceptions in a way that past experience has found to be better than random. So, that is all that is required for a basis of “rational behavior,” since it implies merely that the actions (probabilistic) of an entity will be on the whole in agreement with what it has “learned” to avoid or seek. This does not claim that what it avoids or seeks, its goals, are “really” the correct things for a long and happy existence, but merely that the behavior of the entity is correlated, even weakly, to its perceived environment.

Of course over time, as evolution and learning proceed, the differential success of experimenting entities may find ways of improving and refining their modes of perception and of interpreting what they perceive. They will therefore possess “better” means of sensing the environment, and also will modify their internal mechanisms (or models), or “sensemaking” patterns. So, on a short timescale “knowledge” really corresponds to the internal models (interpretive frameworks) that people construct, which tend to determine current behavioral responses. On a longer timescale these internal models will need modification as others alter their behavior and as the environment changes. And clearly, the “environment” for any individual is made up of some physical reality and also of the different individuals and organizations that are part of the individual’s social and organizational networks. In this way, the permanent updating of internal “models” and the “sensemaking” patterns that translate perceptions into response (input into output) are now simply part of the evolving system, emerging through self-reinforcing experiences and either transforming themselves appropriately over time or being eliminated in the system.

The questions that any individual or organization wants to address are how best to update their understanding of the world around them, and how to relate their desires for success and prosperity to their possible strategies, policies, and actions. It is important to stress the “possible” because it essentially moves us into the future, and into the domain of prediction and predictive knowledge. What we want to know is fourfold:

- ❖ What possible options do I have?
- ❖ What might be the outcomes of these?
- ❖ What are my criteria of preference?
- ❖ Which possible choice is therefore preferable?

Classically, for a mechanical system, we could have analyzed the system, constructed a model, and run it forward to explore what would happen under each possible option. Indeed, this is clearly the way in which many operational decision support systems work. The “decider” is outside the system and is merely trying to decide what to do in relation to the system if it is running. The “decider” simply considers which of the options seem to be most advantageous (e.g. cost–benefit) and chooses accordingly. In such a case, there would not really be any “choice” because the outcome would simply result from the preferences of the decider. However, we could also imagine a slightly more complex case when the design of the mechanical system itself is open to choice, and this is the decision that must be made. Then, the different possible systems could be run in order to see how each of them performs under different possible design options, and again the decision maker’s preferences will simply lead to a decision.

But, what if the “system” concerned involves people? We may be making a mechanical product, but what about potential customers whom we hope will buy our new design? Can we include them in our “mechanical” model? And what about the responses of possible competitors or partners in deciding on a new design or a new pricing strategy? What if we are discussing how to organize and manage a process? How can we proceed then?

The traditional answer really has been by trial and error. If 1,000 different decisions are made, some will survive and some will not. This will serve, for some, to define which are “better,” and the system will evolve over time, retaining some experiments and rejecting others. This is the principle behind a belief in the effectiveness of “free markets” to bring economic success, which is simply that a regional economy can develop providing that there are entrepreneurs willing to make economic experiments to discover which products, organizations, and strategies work and which do not. However, in such a blind process it will not be at all clear which options correspond to choices that are generalizable to other situations and which are not. There will be no transferable knowledge. Obviously, the whole process will be affected by the richness and diversity of the entrepreneurial explorations, and also the willingness to take risks. For any individual, there is still a real desire to be among the successful rather than the failures. There is therefore a desire to try to understand better what outcome may await a possible action rather than merely trying it out. And clearly, everyone will try to use their internal models, their interpretive frameworks, to help decide which experiment they believe is more likely to succeed.

Equally clearly, the degree of success that will actually be encountered will depend on how well the chosen action of one actor finds favor and reinforcement in the subsequent responses of the surrounding individuals and organizations. This will depend on their internal models and interpretive frameworks, and whether they see the action, opportunity, or product as positive or negative in their own attempts to achieve their aims. Complexity arises because we want to include in our internal models the internal models of others, internal models that may therefore have some representation of our internal model—and so on. Since this infinite regression cannot be dealt with, we tend to return to trial and error, to simple categorization of situations, rules of thumb, or faith in the latest management fad or guru.

In this article we want to try to consider how complex systems science can throw an important light on the evolution of social and economic systems, and provide a useful new basis for our reflections.

CO-EMERGENCE OF KNOWLEDGE AND STRUCTURE

As a first example of the kind of insights that complex systems thinking can bring, let us consider a simple logistics problem. This concerns the distribution of photocopiers in the UK from a central point of production. It is a simple model cunningly set out on 100 Excel cells.

In Figure 1 we see the average distribution of daily demand.

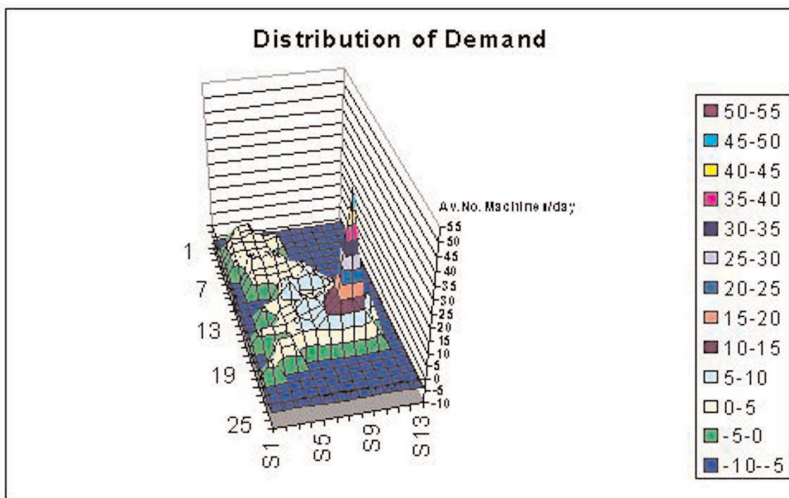


Figure 1 The distribution of daily demand for photocopiers in the UK

The distribution was achieved by the company setting up one major depot and four subdepots at different locations in the country. It is based on some research conducted by U. Seppala and the author in 1995.

Now, the question arises as to whether this is a good choice of subcenters and if they are in the best places. One way of studying this problem is to build a model that “self-organizes” the subcenters and chooses how many there should be and where they should be located. To do this some simple mathematics is used to represent the decisions of potential customers in any zone concerning where they would tend to buy their photocopier from. This, we shall assume, would be probably from the place with the least transportation costs. We also assume that the costs of running a distribution subcenter have both fixed and variable parts, so that the cost per photocopier decreases with the daily volume. With these simple rules, we try out a dramatically simple idea. Suppose initially that every single zone is a subcenter, so that any customers in that zone can get their photocopiers from the zone they are in and indeed from any zone around it.

However, as the simulation proceeds, zones with very little traffic have higher unit costs than zones with more, and so there is a tendency for customers to prefer to buy from the zones with lower costs and higher volumes. The simulation therefore leads to the suppression of small subcenters of distribution and the growth of a number of main centers. The increased throughput decreases cost and increases customer attraction,

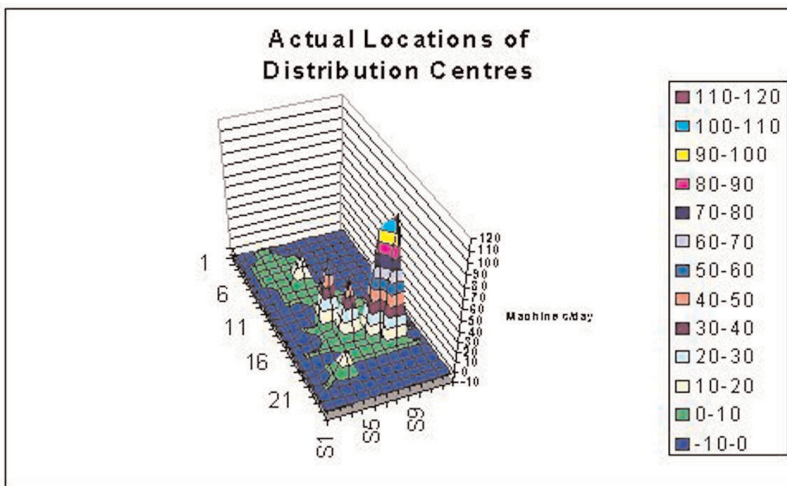


Figure 2 Distribution and throughput of subcenters of distribution

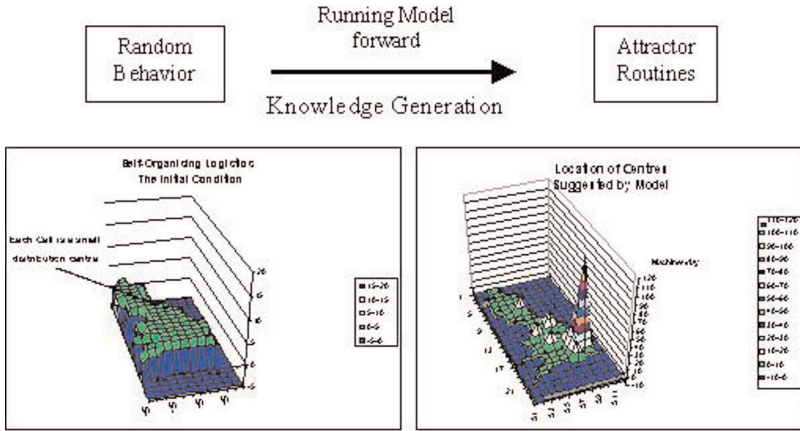


Figure 3 An effective distribution structure co-emerges over time with customers' knowledge of where to buy

which in turn leads to higher flows and still lower costs, as detailed in Figure 3.

The customer behavior changes during the simulation from being essentially random to being highly directed to a large center that is nearest. A particular pattern of centers emerges from the many possible, but the precise structure depends on the small differences in size of the initial distribution of centers.

What we see is that our experiment presents us with the co-emergence of structure and knowledge of that structure.

The centers emerge because customers choose them disproportionately as a result of their price advantage, and customers' knowledge of the advantage increases with use leading to the price fall, thus reinforcing the growth. At the end of the simulation we have customers who "know" exactly where to go, and a structure of distribution centers that is entirely logical. Structure and knowledge have co-emerged over time.

However, if people go always and only to the appropriate center so that it becomes routine, they may in fact forget why they do this. Without comparisons of the cost of alternative behaviors, the knowledge "map" will erode. They would not even know if some center started offering free photocopiers unless they talked to people from other affected zones. But in an efficient system, why would they do this? More importantly, if the pattern of demand, the nature of photocopying, or the pattern of supply were to change, then the pattern of subcenters and the knowledge of customers would both have to change. Our model would succeed in model-

ing that adaptation, provided that we did have some “exploratory” customers who were not totally locked into their previous routines. It would take these “switchers” to demonstrate the benefits of some new choice and to set off a change in the structure. This again reveals the importance of micro-diverse behavior that allows information to be generated about what it is like to behave in a nonaverage way. This tells us something about the real world and how learning really takes place.

What is also important for our discussion is that running an artificial, mathematical model actually created knowledge. It was able to find fairly effective structures and make choices that were better than “random.” It is true that if in reality 1,000 firms had all randomly tried different structures, we would possibly have found an even better structure. However, this would have required the bankruptcy of all the firms except one. One firm equipped with our self-organizing model and the data about costs as a function of volume would have been able to ensure a response that was far better than random. Overall, for society it would have been far better to have firms that make such calculations rather than simply relying on thousands of failed experiments, with all their costs in wasted materials and personal effort. A “modeling” firm would have a much better chance than a nonmodeling firm, particularly with this kind of model that allows for further adaptation and change to be anticipated.

Of course, if only “modeling” firms survive, then unless they are solving quite different problems it will lead inevitably to a competition concerning the speed and the method used in the modeling. This corresponds to evolution moving up a level from merely finding effective ways of doing things, to finding more effective ways of finding more effective ways of doing things—and so on.

This simple little exercise in complex systems thinking has already shown us that if we want our organization to survive, we must not only improve how it does what it does, but also improve the way it learns how to improve what it does.

EVOLVING ECONOMIC MARKETS

In a second example, we will consider the more general problem of a mathematical model of an economic market. In the model we will include several competing firms and their potential customers. The structure of each firm is modeled as shown in Figure 4. Inputs and labor are necessary, and the cost of these, added to the fixed and startup costs, produces goods that are sold by sales staff who must “interact” with potential

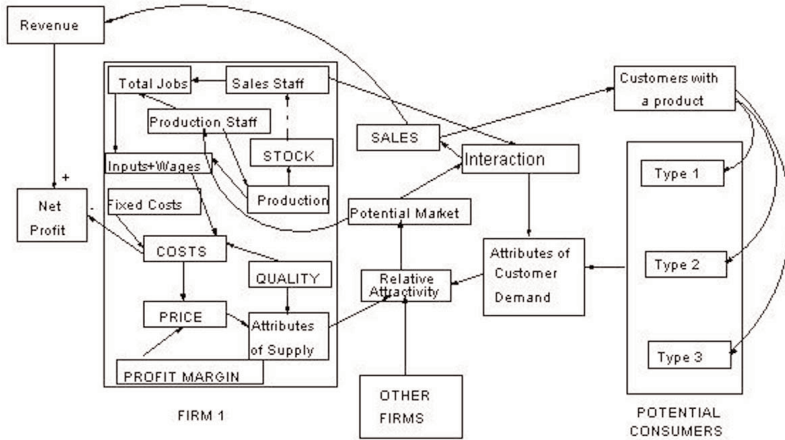


Figure 4 A simple system dynamics model of the dynamic interaction of demand and supply in a market

customers in order to turn them into actual customers. The potential market for a product is related to its qualities and price, and although in this simple case we have assumed that customers all like the same qualities, they have a different response to the price charged. The price charged is made up of the cost of production (variable cost), to which is added a markup. The markup needs to be such that it will turn out to cover the fixed and startup costs as well as the sales staff's wages. Depending on the quality and price, therefore, there are different-sized potential markets coming from the different customer segments.

When customers buy a product, they cease to be potential customers for a time that is related to the lifetime of the product. For high-quality goods this may be longer than for those of low quality, but of course many goods are bought in order to follow fashion and style rather than through absolute necessity. Indeed, different strategies would be required depending on whether or not this is the case, and so this is one of the many explorations that can be made with the model.

In the model, each firm has a "credit limit" that it must not exceed, a mechanism that was inserted as a result of building the model. This is because the initial idea had been to build a model in which firms respond to above-normal profits by expansion (expressing an intentionality of making profits) and reducing production if they are making losses. However, with these rules the model would not start running, because firms all began with an investment—that is, a negative profit—and immedi-

ately shut down. Clearly then, the behavior of the firms could not be represented by such a simple rule, and the next idea was to consider a rule that would link growth and decline in production to a firm's "expected" profits.

Here we see how *building* a model (even without actually running it) allows us to anticipate some of the real problems that would be faced by a participant in the real world. For example, we are also faced with the problem of calculating what profits an entrepreneur might "expect." Obviously, the impossibility of doing this calculation ties in with the fact that some firms will fail although the entrepreneurs involved would not have started off with an "expectation" of failure. We know that the actual outcome will depend not only on what the entrepreneur does, but also on what his competitors and potential customers do. So to write down a mathematical expression for his expected sales and profits under different assumptions concerning the quality and markup on his product is impossible. Therefore it is impossible for an entrepreneur to know this in the real world, and therefore entrepreneurs must act without this knowledge. Evidently, expectations can be right or wrong as the attributes of their strategies are revealed over time.

To represent this more realistic idea, the model supposed that firms had certain "strategies," represented by the chosen parameters of quality and markup, and that they simply "discovered" whether that strategy led to growth or decline. In our model, therefore, we have assumed that managers want to expand to capture their potential markets, but are forced to cut production if sales fall. So, they can make a loss for some time, providing that it is within their credit limit, but they much prefer to make a profit, and so attempt to increase sales and to match production to this.

The model was then used to explore the relative successes of various strategies when pitched against each other. In this way the "modeler" may be able to discover a robust and successful strategy for use in real life. The model thus allows the modeler to have an idea of how different strategies interact, although it assumes that the strategies are *fixed*. It does not tell the modeler what would really happen over time, since a firm that could see its fortunes failing would probably attempt to change its strategy. Instead of mechanical players simply churning through their script to either death or glory, we have adaptive, intelligent players that can respond creatively. They may simply try a new strategy of the same type, but, more interestingly, they may move into a new dimension, bringing new attributes and factors into the game.

As a first step toward an evolutionary model, we can modify the model described elsewhere (Allen, 1994, 1995) to allow firms to learn and to discover successful strategies within the given strategy space. This still does not try to capture intentionality on the part of the entrepreneurs, but merely the will to test their beliefs about a strategy. In this case we consider a model in which six firms interact with initially random strategies. As the model runs forward, some firms succeed, making sales and profits, and the profits serve to provide investment in increased production, lowering production costs. Other firms decline, and eventually exceed their credit limits and are closed down. In this model, however, these firms are then immediately relaunched with new, randomly chosen strategies. In this way, over time an evolutionary process is simulated, in which gradually the market structures into a few successful strategies, corresponding to niches.

A typical long-term simulation is in Figure 5. This shows the 2D space of markup and quality, and the positions of the various firms. The rows at the top show the strategy, price, profit, present balance, and sales of each firm, and the state of the market is shown in the lower left. The simulation shows us that using purely random searches of possible strategies leads eventually to a fairly sensible distribution of the firms in the space of “possibilities.” Essentially, with the particular parameter values used,

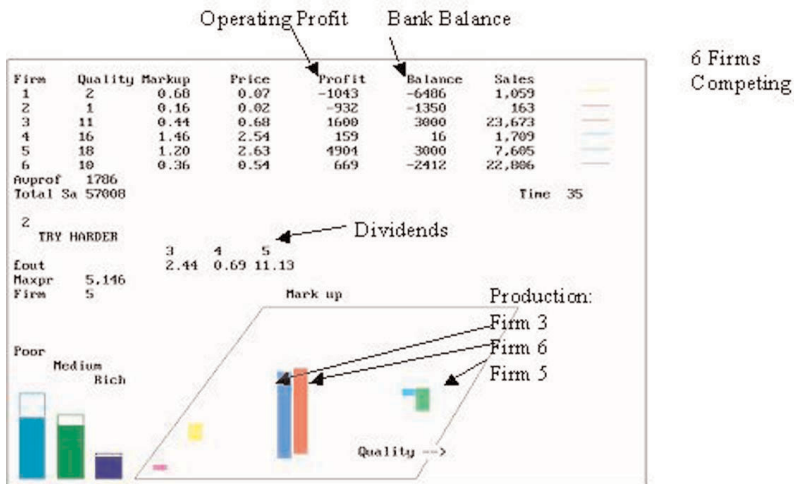


Figure 5 The situation of six competing firms after a long time. The firms have tried out various strategies. Whenever a firm fails it is relaunched with a new strategy picked at random.

there are two approximate niches that can coexist. They are around quality 10 and markup 40 percent, and quality 17 and markup 110 percent. The system takes time to find these out but eventually does so, and provides quite reasonable customer satisfaction with this market structure. However, this assumes a random strategy search for entrepreneurs, which they would certainly not consciously admit to. Indeed, it would seem to throw into doubt their “right” to be entrepreneurs if they flipped a coin to decide what to do.

The evolutionary model of Figure 5 has a kind of “Darwinian” evolutionary mechanism that allows entrepreneurs to explore the “possibility space” for products of this kind. The payoff achieved by any one firm or entrepreneur depends on the strategy (product quality and markup) used by the other entrepreneurs present. The real evolutionary sense arises if we admit that no products have only a single dimension of “quality.” There will always be such factors as performance, weight, efficiency, style, color, noise, flexibility, and so on, as well as simply price. Entrepreneurs would therefore explore different technologies, organizational forms, and production factors, as well as different types of product. This process would lead inevitably to a broader coexistence of different firms, since the multiple dimensions of attribute space produce a greater spread of consumer preferences and less intense competition in general.

The simulations can be used to examine the effects of a firm actively “hill climbing” in profit space by exploring systematically the gains or losses of higher or lower qualities and higher or lower markups. This is very successful for a firm, as we see in Figure 6, which shows the effect of inserting the ability for firm 1 to hill climb while the others do not. Firm 1 rapidly moves into profit by climbing whatever gradient it encounters. It succeeds in paying a dividend to its investors.

However, if we allow all the firms to hill climb, their mutual interaction reduces the advantage of learning. There are also fewer bankruptcies in attaining this market structure. We see the “limits to learning,” in which the speed of learning and the frequency really matter and affect the ability to survive and prosper. In Figure 7 we see the situation at $t = 40$, when all the firms of Figure 6 can hill climb.

In Figure 7 the whole market evolution is different, as all the firms hill climb in profit space, moving overall to higher qualities and higher markups. However, for several firms this very sophisticated strategy, involving careful testing of experimental variations in markup and quality, does not bring success. In that case, it becomes important to see what other strategies are possible, and whether they can result in better outcomes.

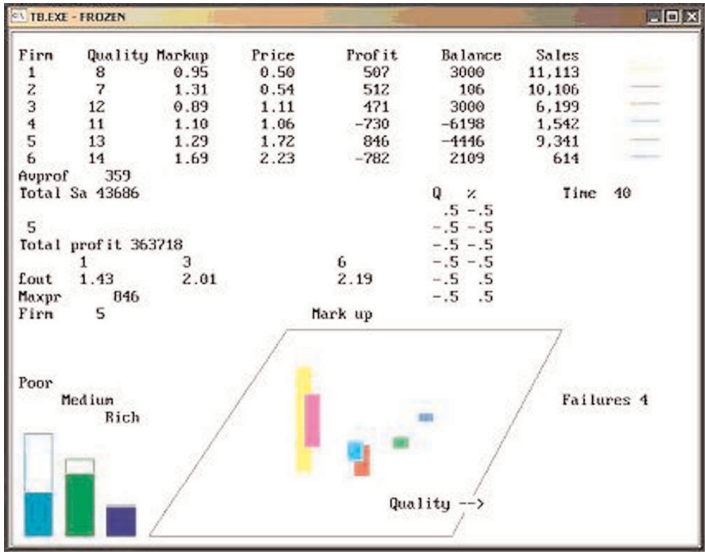


Figure 6 This is an identical simulation to that in Figure 5, except that here Firm 1 tests the profit gradient and moves Q and % accordingly uphill. It does much better, making a profit and paying a dividend to investors

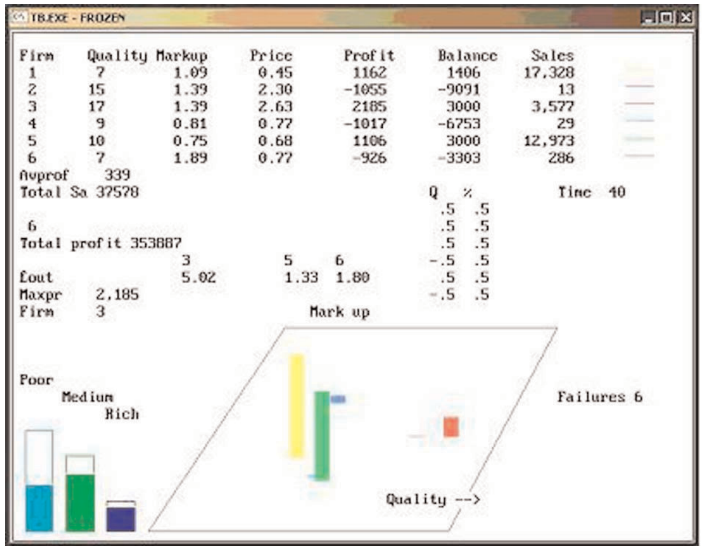


Figure 7 The whole market is completely different when all firms can climb hill. Firm 1 does just get into profit, but not enough to pay a dividend to investors.

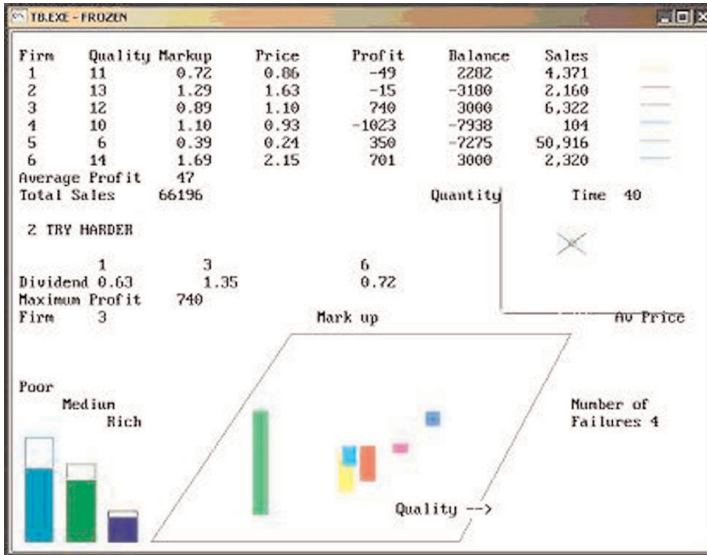


Figure 8 If Firm 1 imitates whoever is currently winning, this can lead to success. It is in the mid-range niche and so will probably survive. It avoided the bankruptcy that occurred in Figure 5.

One such strategy is for a firm to monitor the market carefully and adapt its production as rapidly as possible to copy whichever firm is currently making most profit. The model then allows the imitator to move toward its target. Of course, it may not have the same economies of scale, but nevertheless its presence clearly increases competition at that point in strategy space and changes the outcome for the market as a whole. In Figure 8 we see the result for an identical simulation to that of Figures 5, 6, and 7, except that here Firm 1 discovers which firm is most profitable and imitates its strategy.

We see that the number of bankruptcies “required” to shape the market varies for the different runs. For the “Darwinian” strategy of Figure 5 it is five up to this point, and for Figures 6, 7, and 8 it is four, six, and four respectively.

In the next simulation, we consider the impact of the idea of imitating the firms making up the market. Rather than taking the risk of finding out whether one’s own, individual strategy will really work, it seems tempting to imitate whichever strategy is making maximum profit. At least the “decision maker” is not going to be alone, and obviously it must be a good strategy, since it is already making more profit than any other. So, the idea can be tested. What happens if all the players decide to imitate whoever

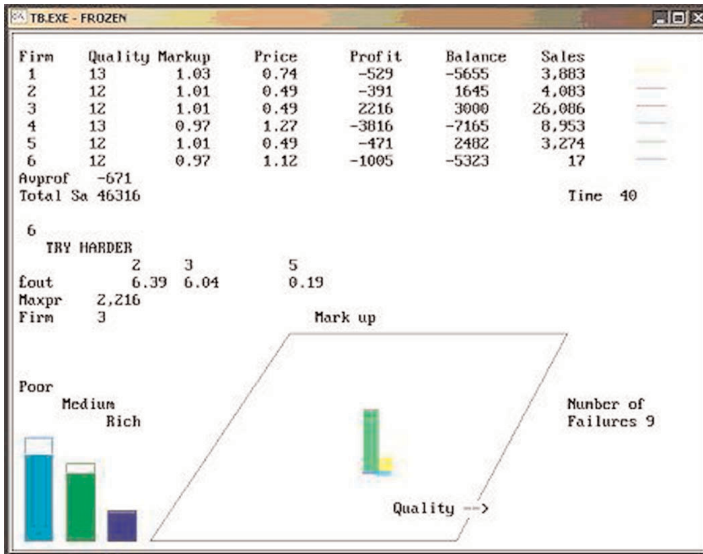


Figure 9 If all the firms imitate a winning strategy, the outcome is bad. In trying to avoid risks, they actually increase them. There are nine failures up to this point.

is making the most profit? The answer is shown in Figure 9. It shows us that all the firms move to the same place in strategy space, and in so doing increase the degree of competition that each feels. As a result, there are more bankruptcies (nine) than in any of the other simulations.

What might have seemed a “risk-averse” strategy turns out to be the opposite! To imitate in a market of imitators is highly risky.

In a previous article concerning the emergence of different strategies among fleets of fishing boats (Allen, 1998), it was shown that what mattered was that an ecology of strategies emerged. Instead of having agents that are susceptible to adopting the same strategy, what really matters is that there should be real micro-diversity such that whatever happens, there will be a diverse set of strategies being played out in the collective system. This is equivalent to insisting on the opposite of “best practice.” In the next simulation of Figure 10 we see the outcome for a situation in which some firms imitate winners and others hill climb.

We see that it is not particularly successful for Firm 1, but that it has actually produced an overall market structure that has been the largest in total profits, and has only suffered three bankruptcies.

This introduces an interesting point about the different levels at which we can look at market evolution—that of the internal capabilities

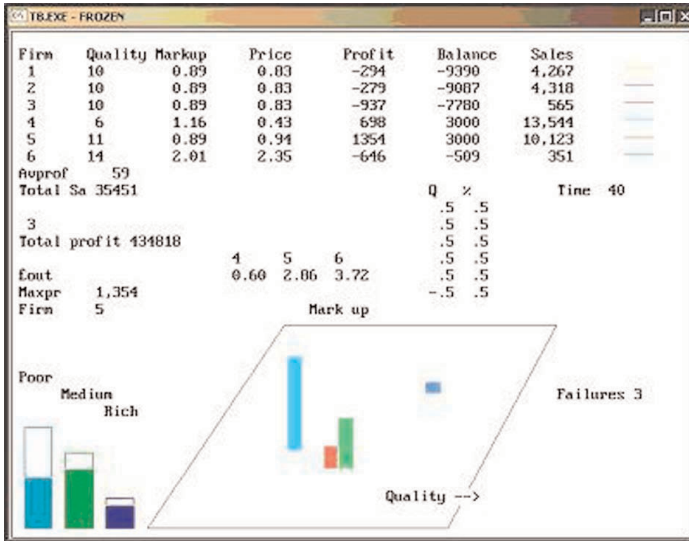


Figure 10 The situation at $t = 40$ for a set of mixed strategies. Firms 1 to 3 imitate and Firms 4 to 6 hill climb. There are only three bankruptcies and good overall profits.

of firms and of their products, the strategy of one firm relative to the others, and finally the overall outcome of the different capabilities and strategies adopted by participating firms. In some ways, for public policy what matters is the level of customer satisfaction and the level of overall profit for the sector. In our accounting for overall costs we need to include that of bankruptcy, since every time it occurs in our model, the social system, other firms, and so on lose 10,000 units. In the real world the costs can be more devastating still to those involved and could even lead to a serious limitation on the willingness of actors to innovate.

We can examine the question of the overall outcomes for the “industry” of different strategies. In order to look at this, we have calculated the overall profits of the whole market, and included the costs of bankruptcies, in which often a loser takes trade away from others in an attempt to keep going, but eventually crashes with debts. In Figure 11 we show the overall outcome for four different learning strategies:

- ❖ Darwinian (random strategies of quality and percentage profit, no learning).
- ❖ Old strategy (if profit less than half average, cut prices by reducing percentage profit).

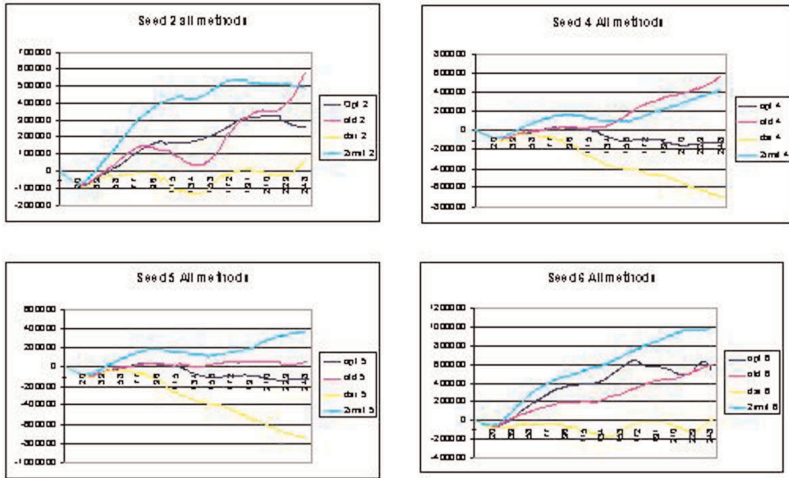


Figure 11 Darwinian learning can be very inefficient, but diverse strategies of hill climbing and imitation of success are best.

- ❖ Hill climbing (test the “profit slope” in quality and percentage profit space and move uphill).
- ❖ Firms 1 to 3 are imitators of winners, while 3 to 6 are hill climbers.

The comparative results for the overall profit profile for the entire market are shown in Figure 11. However, here we have also performed four different runs for different sequences of random numbers, implying simply a different sequence of chance events.

The important result that emerges is that in general, hill climbing in profit space is a good strategy, but a system that mixes this with some firms that imitate success seems even better. However, what is really significant is that “luck matters.” Even for identical potential demand and potential supply, very different overall markets can emerge simply depending on the particular luck of who tries which strategy. This shows us that for the same potential demand, for the same technology, the same strategies, and the same interactions, chance can still allow great variation in market structures to emerge, some very favorable and some very unfavorable, and this tells us that the “structural attractors” of economic markets are diverse and of very different overall efficiency. The invisible hand seems to be highly capricious.

These simple evolutionary models allow us to explore the importance of “learning” and sensemaking. Knowledge about strategies needs to evolve as the use of past knowledge invalidates itself. This evolution can

be achieved by constant monitoring of the marketplace, followed either by imitation or hill climbing. It is interesting to reflect that our modeling (an entirely artificial activity) can nevertheless tell us things about resilient strategies that we did not necessarily know. This means that even if we dispute the precise accuracy of the representation of a particular market system, reflection and modeling can tell us how to learn. In other words, by running such models and testing out firms with different rates and types of response mechanism, we can move toward understanding not only the emergent “behavioral rules” for firms, but also the rules about “how to learn” these rules—that is, how much to experiment and with which parameters and whether any new dimensions of attribute space can be invaded.

This is a key issue, since real innovations concern new dimensions or attributes, and this can confer a temporary monopoly on a firm if it can move into some aspect of quality space that has hitherto been neglected. In this way an “ecology” of firms and products will eventually form, by taking up niches in a multidimensional attribute space, which may still display occasional restructuring. In this model only the quality (1 dimension) and the markup are considered to be decision variables. However, we could use other dimensions of quality space, the parameters of the equation governing the sales force, and also the “research and development” parameter, which can also lead to a change in the firm’s performance parameters.

Returning to our original question concerning the overall effects of changes that occur in socioeconomic networks of actors, we see that evolution and learning do tend to give rise to products and strategies that fulfill people’s requirements. However, we also see that the success of this emergence depends on the nature of the “search mechanisms” that underlie the adaptive processes. Darwinian search can be so “wasteful” in bankruptcies that it can lead to a “net negative” value for the entire market sector, for a considerable period of time. The emergence of a “sensible” outcome depends on the existence of micro-diversity among individual agents, so that “search” is sufficiently broad and effective to discover bundles of activities that can pull in resources from the environment—from customers. Overall, we see that the effect of the evolutionary market models, and of the emergent structure, is to fashion products and production in ways that customers will buy. This in turn ensures that the sector dissipates energy and matter in greater quantities than it would otherwise. However, if firms merely imitate each other, through fear of risk taking, the product exploration is very poor and the market structure not sensible.

Knowledge is only created if the search process at the level below is broad. *Mechanisms that lead to micro-diversity lead to emergent “knowledge” at the level above.* But here we see a conflict between the requirements for “success” for an industry and for individual firms. This really mirrors the result already obtained from a model of Canadian Fisheries, in which successful fishing required exploration by “stochasts” as well as efficient exploitation by “cartesians” (Allen, 1998).

The point is that no individual firm wants to be the “explorer” that is guided by random chance. Instead, it wants to “anticipate” the outcome of evolution and “be there already.” It wants to know how to create a learning organization that will be able to generate a sensible product design or strategy, and to transform this over time as circumstances and opportunities change. The answer is through an evolutionary process quite like the market model described above, in which possible designs are explored and a conscious selection process internal to the firm is used to decide which to adopt. This selection process would aim to ensure that the new products would fit the firm’s view of its own identity and strategy.

DISCUSSION

In reality, each agent in a social or economic network tries to make choices that are “better than random,” and no CEO could be seen flipping a coin in order to resolve some strategic choice. This is the essence of knowledge: an interpretive framework that enables us to make a “better than random” choice. This process is summarized in Figure 12.

In this figure, we show the mechanisms necessary for a firm to respond “intelligently” to possible threats or opportunities in the environment. This will primarily arise from customer tastes and needs and from the behavior of other firms in the market. There are three main parts to the diagram. The first is the environment itself, upper left, where potential customers, competitors, and collaborators exist. The second is the pool of new ideas seen on the right of the diagram. That corresponds to “raw, random fluctuations.” It is perhaps like the subconscious mind, with unrestrained novelty and action suggested. The rest of the diagram of boxes and arrows constitutes the “guidance” system of the organization seeking to filter out unsuitable, irrelevant, and ineffective ideas. These concern the establishment of benchmarks, of target attributes and criteria of evaluation to retain suitable and successful new designs, innovations, or changes. A proposed new product is then put before the board,

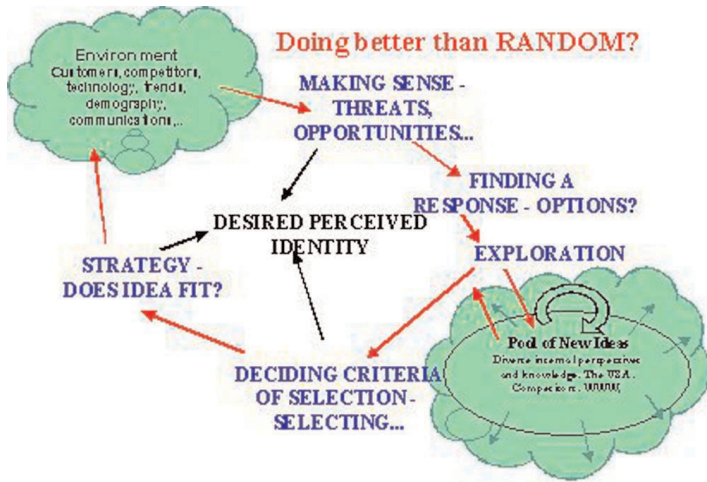


Figure 12 The clouds and arrows constitute an attempt to guide the innovations that are actually launched into the environment from an initially random selection. The desired perceived identity is really the firm’s current strategy.

who accept or reject it depending on whether it fits the “strategy”—the firm’s desired perceived identity.

It is important to realize that this picture really operates on multiple timescales. At a microscopic scale, subcomponents will be designed and judged by how they fit into the desired performance capabilities of the new product. At the next level, a suggested new design or product will be evaluated using performance tests, and overall the project will be selected by seeing whether it fits the desired perceived identity of the firm—its current strategy. On a longer timescale, this strategy (desired perceived identity) itself will be subject to possible renewal and new conjectures will be put forward. The same diagram will apply, but instead of a new product design being the focus, some new possible strategy of the firm will be considered and judged against the long-term overall measure—the expected financial viability of such a strategy. Ultimately, survival will be the long-term selection criterion for the whole enterprise.

We still rely on the occurrence of diverse novel ideas, but we hope to “anticipate” the fate of possible new products when launched in the market, by devising “tests” that will allow us to launch only successful new products. Again, in order to do this perfectly we would need to have a model of the different actors in the market, and also what they might

launch or need. We find the same paradox as before, that in order to know what a firm should do, we need to have a model that includes knowledge of what all the other firms and customers would do. In effect, we must fall back on the same kind of mechanism as the industry, by having micro-diversity inside a firm that produces a broad search, not a narrow one, and uses the differential success of the prototypes to guide product choice.

So, a firm creates knowledge by a wide search of possible products and an evaluation that is characteristic of its opportunities and inner resources. The industry gains knowledge by the diverse searches of its diverse firms. Everything rests on the hierarchy of structure and its micro-diversity. Within firms, departments need internal diversity and multiple perspectives among individuals to follow a healthy evolutionary path. Advantage goes to the individuals that can form the more effective interpretive frameworks to guide their selections. *They still need the initial random ideas, however. Similarly, firms that have a good interpretive framework of the marketplace succeed in making the right choices about the evaluation criteria for their innovations.*

At a still higher level, societies that can understand what attributes the activity and products of an industry may have can regulate the market interaction in order to maintain some ethical criteria that are considered important. Of course, this may be achieved only *post hoc*, after the unacceptable consequences have been declared. Increasingly, this socially contextual “meta-market” level is becoming important as “price” alone is perceived as inadequate information to guide the outcome.

The real problem is a multidimensional one, with the multiple attributes affected by the choices made by the agents. Almost no improvements can be made assuming *ceteris paribus* (all other things remaining the same), because many attributes are emergent properties that arise from the interacting components that make them up. Since some of these attributes may only declare themselves over time, possibly long after the decision to “adopt” has been made, the full implications of any particular design change or innovation cannot really be known at the time the decision to adopt must be made. In reality then, we can only hope to do better than a random choice, but we can never know absolutely what the effects of a particular decision will be.

If we reflect on the evolution of organizational forms, a similar set of ideas emerges. For example, if we consider the different characteristics that can be discerned in manufacturing organizations (McKelvey, 1982, 1994, McCarthy, 1995, McCarthy *et al.*, 1997), then according to

McCarthy *et al.* there are about 53. However, there are 16 distinct organizational forms that have been identified and these are made of particular clusters of “synergetic” characteristics. The cladistic analysis of these different organizational forms produces a particular evolutionary history that leads to an “ecology” of possible organizational forms. In other words, instead of evolution simply producing “progress” expressed through successively “better” structures, we see a diverse evolution of different types of organization with different emergent capabilities, many of which coexist. So, rather than simply demonstrating “improvement” within a given niche, we see the discovery and emergence of new niches. In a recent article, Allen *et al.* (2004) show that the emergent organizational forms observed are ones in which synergy of internal practices is high. In this way, it is the emergence of better functional capabilities, and higher performance, that leads to the selection of new organizational forms. Once again, new random explorations are made, and structures with emergent properties that are desired by the environment are selected.

This demonstrates the generality of the idea that evolution is all about the discovery and emergence of *structural attractors* corresponding to emergent capabilities and properties. It shows us that for a system in which we do not make the assumptions of average types and average behavior (Allen, 1992) that would take out the natural micro-diversity and idiosyncrasy of real-life agents, actors, and objects, we automatically obtain the emergence of structural attractors. These are complex systems of interdependent behaviors in multidimensional space whose attributes are on the whole synergetic. They have better performance than their homogeneous ancestors (initial states), but are less diverse than if all “possible” behaviors were present. They correspond to a generalization of the idea of hypercycles in the work of Eigen and Schuster (1979), but recognize the importance of emergent collective attributes and dimensions. The particular attractor that emerges results from the precise history of search undertaken. In other words, a structural attractor is the emergence of a set of interacting factors that have mutually supportive, complementary attributes.

If we define “knowledge” as being something that structures possible action from random to highly defined, then clearly self-organization and evolution are creators of knowledge. A set of nonlinear equations representing the behaviors of multiple agents will “run” from initially random choices to behaviors structured strictly on the emergent attractor. There may have been several attractors possible, but any particular “run” will

lead a particular attractor and its corresponding set of behaviors or “routines.” In order for learning to continue, however, these routines must not become totally dominant. Micro-diversity of behavior must be left in the system to some degree. This, together with sufficient communication, maintains the existence of a learning network ready to adapt and change.

If we return to the original question about whether the coevolutionary processes produce any perceptible overall trend, we should go back to an article of 1976 (Allen, 1976). In this, an evolutionary criterion was developed that looked at the mathematical condition that must be fulfilled for any “new behavior” to invade a system. By applying this criterion to simple ecological systems, it was shown that although at each interaction of “predator and prey” there was a potentially balanced arms race, nevertheless the benefits of improved prey behavior were passed on to the predators. Over time, with evolution, the “center of gravity” of the biomass in the ecosystem tended to move “up” the metabolic chain. This may correspond to the observable fact that economic gains tend to be higher further up the economic chain. So, for example, it would suggest that options traders will often earn more than commodity farmers, and this does indeed appear to be true.

The successful development of learning networks is about systems whose internal structure is not reducible to a set of mechanisms. In particular, it is about connected complex systems, for which the assumptions of average types and average interactions are not appropriate and are not made. Such systems coevolve with their environment, being “open” to flows of energy, matter, and information across whatever boundaries we have chosen to define. These flows do not obey simple, fixed laws but instead result from the internal “sensemaking” going on inside them, as experience, conjectures, and experiments are used to modify the interpretive frameworks within. Because of this, the behavior of the systems with which each system is coevolving is necessarily uncertain and creative, and is not best represented by some predictable, fixed trajectory. This takes some steps toward the “postmodern” point of view. However, as Cilliers (1998) indicates, the original definition of postmodernism (Lyotard, 1984) does not take us to the situation of total subjectivity where no assumptions can be made, but rather to the domain of evolutionary complex systems discussed in this article.

Instead of a fixed landscape of attractors, and of a system operating in one of them, we have a changing system, moving in a changing landscape of potential attractors. Provided that there is an underlying potential of diversity, creativity and noise (supposing that they are different) offer a

constant exploration of “other” possibilities. Some of these possible behaviors mark the system and alter the dimensions of its attributes, leading to new attractors and new behaviors, toward which the system may begin to move but at which it may never arrive, as new changes may occur “on the way.” The real revolution is not therefore about a neo-classical, equilibrium view as opposed to nonlinear dynamics having cyclic and chaotic attractors, but instead is about the representation of the world as a nonstationary situation of permanent adaptation and change. The picture we have arrived at here is one that Stacey *et al.* (2000) refer to as a “transformational teleology,” in which potential futures (patterns of attractors and pathways) are being transformed in the present. The landscape of attractors we may calculate now is not in fact where we shall go, because it is itself being transformed by our present experiences.

The macro-structures that emerge spontaneously in complex systems constrain the choices of individuals and fashion their experience. Behaviors are being affected by “knowledge,” which is driven by the learning experience of individuals. Each actor is coevolving with the structures resulting from the behavior and knowledge/ignorance of all the others, and surprise and uncertainty are part of the result. The “selection” process results from the success or failure of different behaviors and strategies in the competitive and cooperative dynamical game that is running.

However, there is no single optimal strategy. What emerge are structural attractors, *ecologies* of behaviors, beliefs, and strategies, clustered in a mutually consistent way and characterized by a mixture of competition and symbiosis. This nested hierarchy of structure is the result of evolution and is not necessarily “optimal” in any way, because there are a multiplicity of subjectivities and intentions, fed by a web of imperfect information and diverse interpretive frameworks. In human systems, at the microscopic level, behavior reflects the different beliefs of individuals based on past experience, and it is the interaction of these behaviors that actually creates the future. Nevertheless, often the outcomes are not what was expected, and often there is not enough evidence to know why what happened happened. The people concerned then ascribe events to whatever they believe and either retire “hurt” or try again with some new hopes and beliefs. Evolution in human systems is therefore a continual, imperfect learning process, spurred by the difference between expectation and experience, but rarely providing enough information for a complete understanding.

Although this sounds tragic, it is in fact our salvation. It is this very “ignorance” or multiple misunderstandings that generates micro-

diversity, and leads therefore to exploration and (imperfect) learning. In turn, the changes in behavior that are the external sign of that “learning” induce fresh uncertainties in the behavior of the system, and therefore new ignorance (Allen, 1994). Knowledge, once acted on, begins to lose its value. This offers a much more realistic picture of the complex game that is being played in the world, one that our models can begin to quantify and explore.

Returning to our initial question about the overall effects of evolution and change, we see that the emergent capabilities to which complex systems necessarily give rise will pull the energy and matter necessary to maintain them into the new structure. Dissipative structure therefore increases. However, at a longer timescale, the evolutionary process may well improve the functionality and performance of this structure so that it provides the capability it has for less energy and less cost. But this will generally lead to the greater abundance of such structures, and also to the emergence of new capabilities and structures, operating in different dimensions and possibly including the old systems within new, more complex entities.

Nevertheless, we should not adopt the naïve view that complexity tells us that things get better and better. In reality, it tells us that although structure will often increase in complexity and dissipation, it can also collapse. There is no guaranteed steady improvement to better and better worlds. Progress can happen from some points of view, but so can catastrophe. Complexity is for adults, in that it means that what we do matters, and trying to gain better knowledge to guide our decisions is certainly worthwhile, though not necessarily easy.

NOTE

This work was partially carried out under the ESRC NEXSUS Network program. The idea of “structural attractors” occurred during a discussion with J. Mcglade and M. Strathern.

REFERENCES

- Allen, P. M. (1976) “Evolution, population dynamics and stability,” *Proceedings National Academy of Science*, 73(3): 665–8.
- Allen, P. M. (1992) “Evolution: Persistent ignorance from continual learning,” in R. Day & P. Chen (eds), *Nonlinear Dynamics and Evolutionary Economics*, Oxford, UK: Oxford University Press: 101–12.
- Allen, P. M. (1994) “Evolutionary complex systems: Models of technology change,” in L. Leydesdorff & P. van den Besselaar (eds), *Chaos and Economic Theory*, London: Pinter: 1–19.

- Allen, P. M. (1995) "Intelligent, self-organizing models in economics and in finance," in S. Goonatilake & P. Treleaven (eds), *Intelligent Systems for Finance and Business*, Chichester: Wiley.
- Allen, P. M. (1998) "Evolving complexity in social science, in G. Altmann & W. A. Koch (eds), *Systems: New Paradigms for the Human Sciences*, Berlin/New York: Walter de Gruyter.
- Allen, P. M., Strathern, M. & Baldwin, J. (2004) "Computational social science: Modelling evolution and emergence in social systems," submitted for publication in the *Proceedings of the National Academy of Sciences*.
- Cilliers, P. (1998) *Complexity and Post-Modernism*, London/New York: Routledge.
- Eigen, M. & Schuster P. (1979) *The Hypercycle*, Berlin: Springer.
- Lyotard, J.-F. (1984) *The Post-Modern Condition: A Report on Knowledge*, Manchester, UK: Manchester University Press.
- McCarthy, I. (1995) "Manufacturing classifications: Lessons from organisational systematics and biological taxonomy," *Journal of Manufacturing and Technology Management: Integrated Manufacturing Systems*, 6(6): 37–49.
- McCarthy, I., Leseure, M., Ridgeway, K. & Fieller, N. (1997) "Building a manufacturing cladogram," *International Journal of Technology Management*, 13(3): 2269–96.
- McKelvey, B. (1982) *Organizational Systematics*, Berkeley, CA: University of California Press.
- McKelvey, B. (1994) "Evolution and organizational science," in J. Baum & J. Singh (eds), *Evolutionary Dynamics of Organizations*, Oxford, UK: Oxford University Press: 314–26.
- Nicolis, G. & Prigogine, I. (1977) *Self-Organization in Non-Equilibrium Systems*, Chichester, UK: Wiley Interscience.
- Seppala, U. (1995) "A self-organising model for logistics networks," M.Phil thesis, Cranfield, UK: Cranfield University.
- Stacey, R., Griffen, D. & Shaw, P. (2000) *Complexity and Management*, Complexity and Emergence in Organisations series, London/New York: Routledge.