Models, knowledge creation and their limits

Article *in* Futures · January 2005

P. M. Allen
Cranfield University

Available online at www.sciencedirect.com
SCIENCE DIRECT
Futures 37 (2005) 729–744
FUTURES
www.elsevier.com/locate/futures

Models, knowledge creation and their limits

Peter M. Allen*, Mark Strathern[1]

Complex Systems Management Centre, Cranfield University, Cranfield MK43 0AL, UK

Available online 29 January 2005

Abstract

Instead of modelling socio-economic situations as mechanical systems with fixed, predictable behaviour, we now see that socio-economic systems are really complex systems, in which various possible structural changes can occur giving rise to a range of different possible futures. This necessary future uncertainty automatically imposes an uncertainty on the precise pay-off that any particular action or decision that an agent may take. Because of this, the decisions that agents will make are also uncertain and this poses limits to our ability to model socio-economic systems and therefore to the knowledge that we can have at any time about the future. Because of this constant knowledge decay, what matters in real world situations of markets and business is the generation of new, current knowledge. Contrary to traditional science in which the natural laws are independent of who knows them, in social and economic systems, knowledge of system behaviour decays over time, and is in any case used up when it triggers new behaviour in the system. Several examples of evolutionary market systems are presented which demonstrate how knowledge is constantly created and destroyed, and the problem of change, innovation and design are shown to be part of a 'boundedly rational' view in which imperfect search gives rise to 'good enough' behaviour. All of this is a radical departure from the traditional approach that falsely believe in the optimisation of designs, behaviours and profits. Complexity tells us that we must accept risk and uncertainty and work loosely, keeping our options open as much as possible.

1. Introduction

Today we know that we live in a complex world of emergent behaviour and attributes, in which our powers of prediction are limited. This contrasts with the 'classical' views that implicitly supposed a mechanical universe, in which knowledge of the laws of interaction

* Corresponding author. Tel.: +44 123 475 4800; fax: +44 123 475 4804.
E-mail addresses: p.m.allen@cranfield.ac.uk (P.M. Allen), m.strathern@cranfield.ac.uk (M. Strathern).
[1] Tel: + 44 123 475 4803.

0016-3287/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.futures.2004.11.010

would allow us to predict, and to intervene in the world to achieve our ends. Prediction had three basic pillars. First, it could be based on the movement of a frictionless mechanical system along a predefined trajectory as for planetary motion. Secondly, for situations with dissipative forces present the prediction could be modified by adding experimentally determined terms concerning viscosity or friction. Thirdly, for such systems a prediction could be made about the 'final' state of the system as it moved to the predictable state of thermodynamic equilibrium. So, in effect we could predict the behaviour of clockwork systems, either running or run-down.

But, with the study of non-linear systems, open to flows of energy, matter and information with their environments, this simple, psychologically comforting myth was exploded. Systems in the real world, on which the sun shines and where structures and differences abound, can respond to their environment in different possible ways. Their configuration and structure can transform itself, resulting in emergent global properties and functionality, and in addition their internal elements can evolve and change over time. Instead of a clockwork system simply running forwards until it unwinds, the very dimensions of interaction with the outside world can change, new attributes and functionalities can emerge, and the nature and experiences of its constituent components can transform over time. Creativity, innovation and emergence characterise the real, open, non-linear systems that make up our world. It is the new paradigm of co-evolving complex systems, and it implies, accepts and embraces the reality of qualitative and quantitative changes over time [1–3]. In reality this is a much more exciting world than the one related to mechanical motion. But despite this it is often seen in purely negative terms, as introducing inevitable 'uncertainty' (as if this did not exist before), and perhaps meaning that we should not have any long term goals and ambitions.

In fact the opposite is the case: *it is because systems can evolve and transform themselves that we can have hopes and dreams of better things*. Obviously, change can also bring worse things, and our aim must be to try to understand the 'trade-offs' that different actions, strategies or policies may involve. Without such an attempt, or by continuing to adhere to the inadequate ideas of mechanical thinking, we are left with essentially unexpected emergent structural changes, with qualities and trade-offs that are essentially random. Of course, evolutionary principles ensure that the adoption of random innovations by individuals lead to evolutionary changes in the overall system that tend to increase its entropy production, which may have no obvious benefits for the humans contained within it.

Clearly, if we knew that unguided evolution would inevitably lead to an attractive solution, then we could possibly relax, but history suggests that this is not the case, and therefore that our choices, actions, strategies and policies do really matter. Individuals have intentions and preferences as they go about their business, and hence would like to relate their decision-making to the outcomes they will experience. Because of this people are attracted by the possibility of prediction, and are disturbed by the idea that it may not be as simple as their decision support software, or management consultant tells them. We might see the acceptance of the co-evolutionary, complex systems view, with its accompanying acceptance of a level of uncertainty, as being a first sign of adulthood coming upon a reluctant world. The answer will be neither that of 'full prediction' nor that of 'no prediction', not 'total control' nor 'totally out of control'. As most people have secretly suspected, the answer lies somewhere in between. And now science is coming to this conclusion too.

Nevertheless, the acceptance of the new ideas bring far more than just an understanding of the limits to prediction, they allow us to reflect on our social organisations, and our economic and political institutions and reassess them in the light of these new ideas. In particular, the issues of intentionality and of the relation between the individual and society have always been thorny ones. We normally accept that individuals try to achieve their ends but that society should be concerned mainly with helping them to do that. This would naturally see an evolutionary process as being the simple, value-free outcome of individual strivings at a lower level. Yet clearly, there are asymmetries of initial endowments, of power, of information, of relationships, as well as of inherent capacities, and therefore the particular social system does affect outcomes. The ‘winners’ are not simply inherently superior. The outcome is affected by luck and by their precise location and path within the social system. Less controversially, we can consider the question of competing firms within an economic market. The intention of each firm may be to make profits and to grow, but only some will succeed in doing so. The strategies, decisions and product designs of each company will turn out to have either successful or unsuccessful pathways, and the question that we need to address is whether a company can ‘know’ enough to do better than random. It is like buying a stock on the equity markets with the intention of making money. Can we do better than a random choice? Our intention is to do so, but what knowledge can we have that could make it so? Particularly as all the other investors are also trying to achieve similar ends. We immediately see the paradox. If markets are efficient, there can be no privileged individual knowledge, and a random choice among them would be justified. Yet a vast industry of financial management exists and on the whole most of us would prefer to consign our life savings to an established company rather than to a series of dice throws.

Of course, even a financial market consisting of actors that assign money randomly to different stocks would still exhibit a sensible evolution over time, since the companies that received investment and also had viable products and organisation would survive while those that lacked one or both of these would not. And this is interesting because it shows us that with or without intention at the individual level, an evolutionary process can still occur, and produce a seemingly ‘rational’ result. The question really is whether the intentions of individuals can successfully be translated into actions that beat a ‘random’ strategy or not. Of course, we also have the problem that sometimes, just by luck, a random strategy will do really well—better than average—and so we need to define the random strategy in terms of its average performance.

In what follows we shall present an evolutionary model that apparently produces learning and progress for a population or group without necessarily assuming intentionality or knowledge in its micro-actors [4]. Then we shall contrast this with a model and study concerning the individual level that shows how decision-making about an innovation or the adoption of a new strategy take place. In the fusion of these two approaches we should see a holistic view of the limits to knowledge and of our response to it.

2. Self-organisation of firms and products within markets

Consider a simple dynamic model of a market in which firms grow or decline according to their success in making sales to potential customers [5]. The structure of each firm that is modelled is as shown in Fig. 1. Inputs and labour are necessary for production, and the cost of these, added to the fixed and start-up costs, produce goods that are sold by sales staff who must ‘interact’ with potential 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 mark-up. The mark-up needs to be such that it will turn out to cover the fixed and start-up costs as well as the sales staff 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 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.

Let us briefly discuss an initial example in order to see how ‘intentionality’ is treated in the model. In the first run we consider a situation with three firms competing. Each firm starts with an initial size of 1000 units. The firms (using random initial conditions) choose

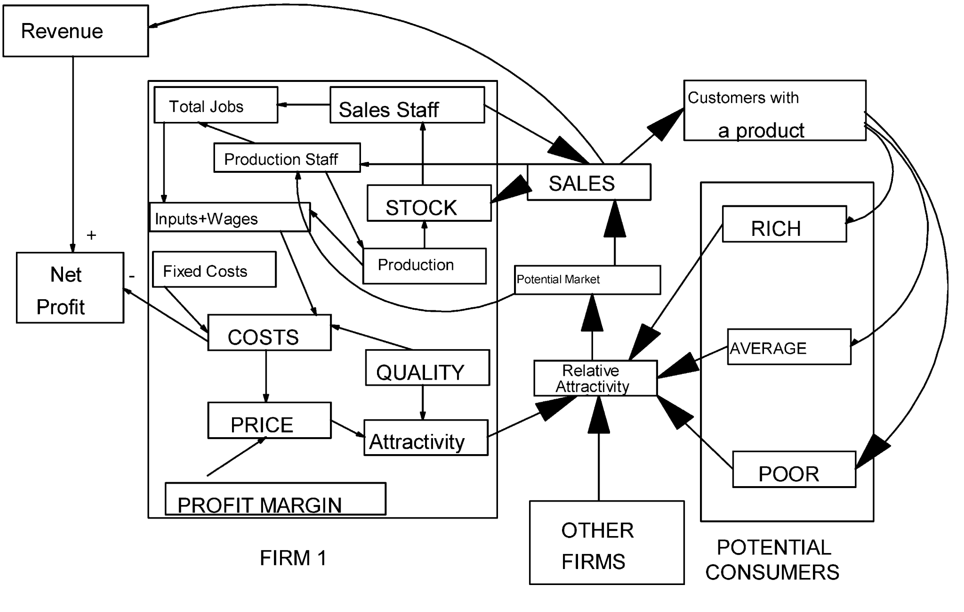

Fig. 1. A simple system dynamics model of the dynamic interaction of Demand and Supply in a market.

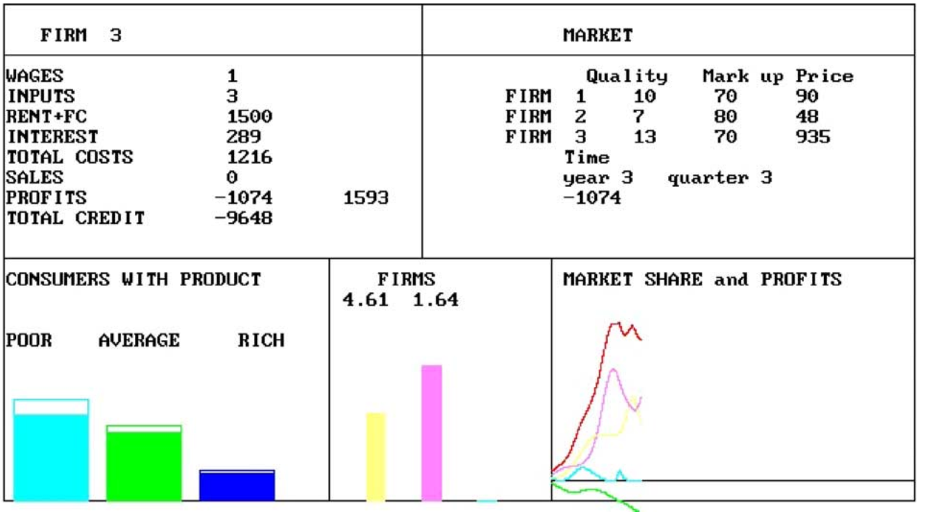

Fig. 2. Firms 1, 2 and 3 produce products at qualities 10, 7 and 13, respectively, and mark-ups of 70, 80 and 70%. Firm 3 fails after three years.

qualities 10, 7 and 13 (in a potential choice of 1–20), and profit margin mark-up of 70%, 80% and 70% (Fig. 2).

Firms 1 and 2 out-compete firm 3, and it crashes after 3 years All firms generated dividends during this run, but following market saturation, firm 3 was squeezed and driven to bankruptcy. In the model, each firm has a ‘credit limit’, which it must not exceed, and this mechanism 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 to reduce production if making losses. However, with these rules the model would not start running—because firms all started with an investment that is a negative profit, and immediately shut down. Clearly then the behaviour 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 ‘expected’ profits of a firm.

But this proved impossible to calculate, owing to the limited (and possibly mistaken) knowledge of entrepreneurs. Some firms will fail although the entrepreneurs involved would not have started off with an ‘expectation’ of failure. Clearly, 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 mark-up, and they simply ‘discovered’ whether that strategy led to growth or to decline. In our model, therefore, we simply 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. There is still no intention inside the model, only perhaps in the ‘modeller’ who may use the information to design a real-life strategy. But this reveals a further paradox, which is that *in order for a modeller to formulate a strategy he needs to understand what strategies the other potential players may formulate*.

3. Strategies interact

The model therefore allows the modeller to get an idea of how different strategies interact with each other. But it assumes that the strategies are *fixed*. It does not tell the modeller what would really happen over time, since a firm that could see its fortunes failing would attempt to change its strategy. Instead of mechanical players simply churning through their script either to 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 new dimensions, bringing new attributes and factors into the game.

As a first step towards an evolutionary model we can modify the model above to allow firms to learn, and to discover successful strategies within the given strategy space. This still does not require intentionality on the part of the entrepreneurs, but it does lead to a market evolution that seems rational. 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 re-launched 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 shown in Fig. 3. This shows the two-dimensional space of mark up 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 does not necessarily lead to a very sensible distribution of the firms in the space of ‘possibilities’. Our model needs to begin to represent the ‘cognitive’ processes of entrepreneurs, who for example, having failed (or seen others fail) with a particular kind of strategy (e.g. low quality, low mark-up) allow this to influence the parameters used for re-launch. Other possibilities could involve imitation.

These simple evolutionary models show us how resilient strategies will emerge from such systems and in the case of particular market sectors suggest how the rules of learning can also evolve. In other words, by testing out firms with different rates and types of response mechanism, we can move towards understanding not only of the emergent ‘behavioural 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 (Fig. 4).

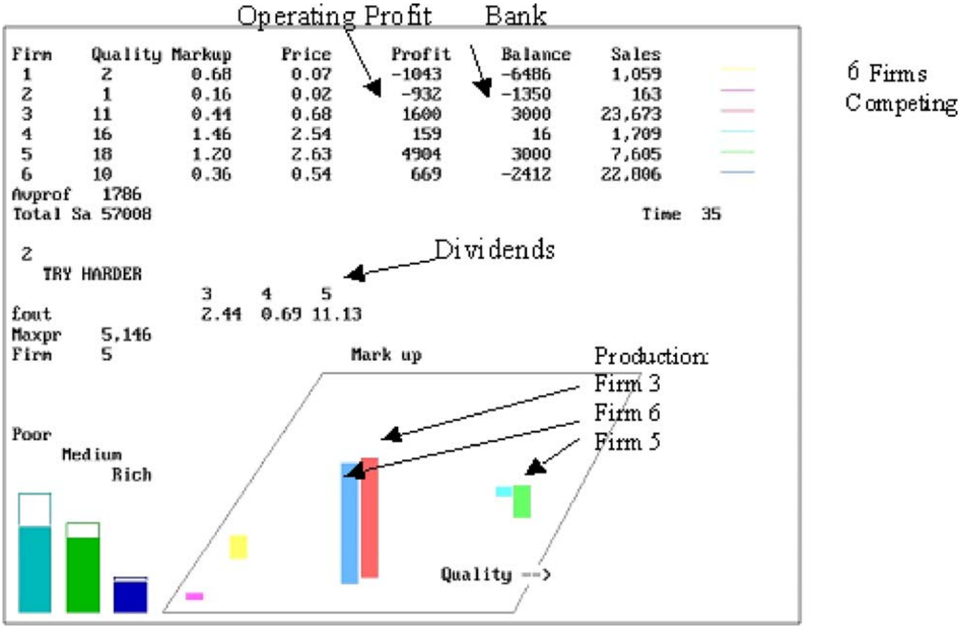

Fig. 3. The situation of six competing firms after a long time. Firms have tried out various strategies. Whenever a firm fails it is re-launched with a new strategy picked at random.

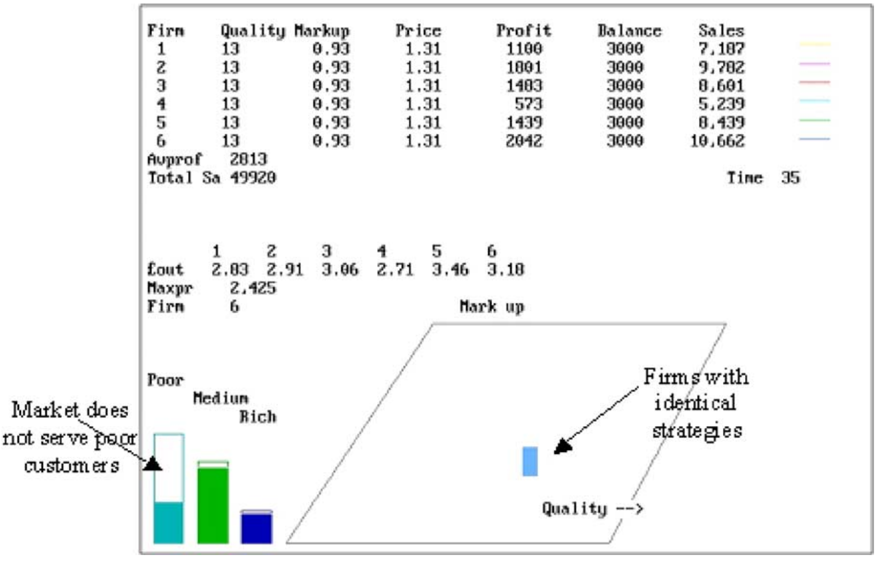

Fig. 4. This is an identical simulation over the same time period. Here, however, in a risk avoidance strategy all firms imitate the winner. The total market is much smaller.

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, that may still display occasional restructuring. Resilience, as a capability, will reside in the nature of the ‘random exploration’ factors that form part of the rules in re-launching a firm. In this model only the quality (a single dimension) and the mark-up are considered to be decision variables. However, we could also use other dimensions of quality space, the parameters of the equation governing the sales force, and also the ‘research and development’ parameter that can also lead to a change in the performance parameters of the firm.

Taking the models described above a stage further, we can envisage an evolving market model in which competing firms learn how to improve their product and sell more of it. Sales increase when a product has, from the point of view of a particular type of customer, the right attributes at the right price. Over the simulation run, the supply networks discover what ‘sells’ and what doesn’t.

In this model the products or services in question could have eight possible attribute dimensions. These attributes are supposed to result from the composition of the product, and so the model runs in two separate, alternate modes. There is an R & D mode in which firms try to improve their product, by offering greater performance in particular dimensions of the attribute space. Particular concepts have different possibilities for trade-offs between the eight different dimensions. Then, in a second phase, the different firms (nine in all) compete for sales to potential customers. There are three types of customer in the model, but clearly this could be adapted to whatever is realistic in a given situation. The customers gauge how attractive they see the products, according to their own needs concerning the eight dimensions, and buy them accordingly. This leads to revenues, and potentially to profits and further investment. If a firm fails to cover its costs, and exceeds its bank credit limit, then it is made bankrupt and a new network is launched with a new concept and initial investment. We shall not describe this model in great detail here, but simply show some views of the simulation at different stages of a run.

In Fig. 5 we see the model at an early stage. This corresponds to internal R & D where the nine firms are exploring how to improve the different attributes of their product. The attributes of the nine firms are shown vertically, together with the ‘concept’ number that is characterising the technological trade-offs they will experience in developing their product. We see the ‘stars’ of attributes of the three different potential customer types, and these are clearly in all eight dimensions. However, it does not follow that their desires are realistic and that it is possible to make a product with all eight dimensions. The nine competing firms have only the first three attributes and they will have to learn about the missing factors that will attract customers, and this will happen over time. However, because the market is competitive, the different networks will encounter different successes in selling their products, and in turn will generate different amounts of revenue. If they are in profit, then they can carry on increasing the production of their product, but if they are making a loss then they will not be able to expand as fast. Because of this, some will go bankrupt and will be replaced by new, hopeful networks.

In Fig. 6 we see the market evolution phase of the model, when the competing firms have differential success in selling their products, and some are growing faster than others.

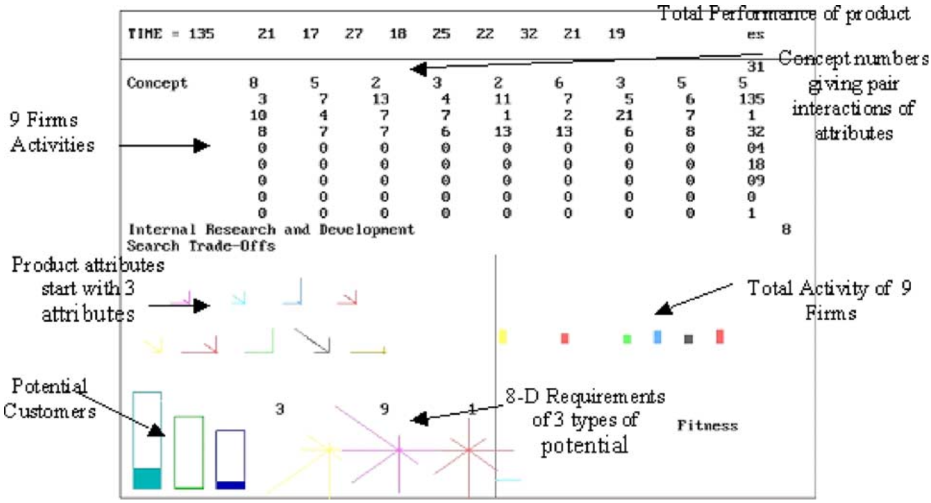

Fig. 5. A moment during one of the early research phases of operation of the model. Products only have three of the eight possible attributes.

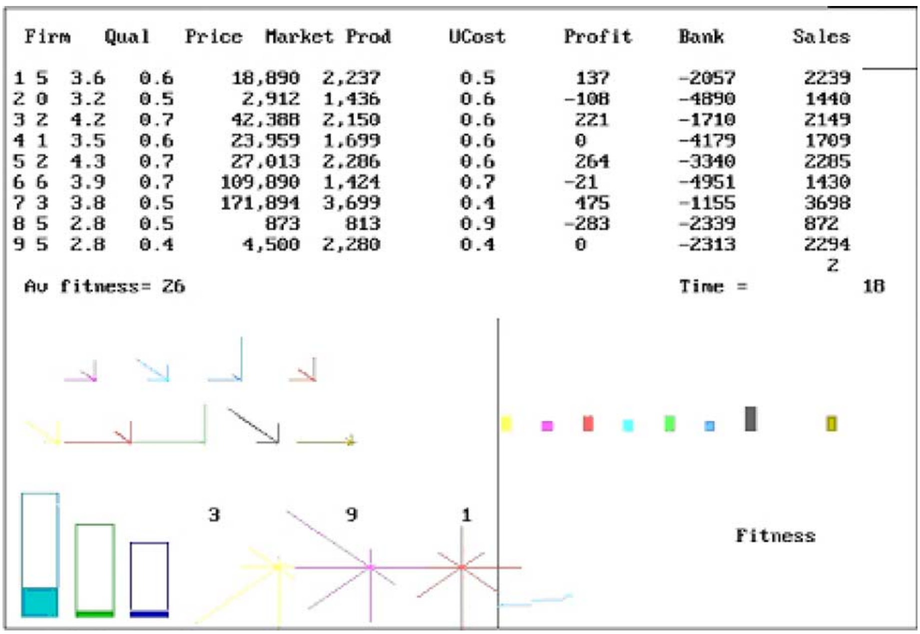

Fig. 6. Here the nine firms are now listed in horizontal rows. We also see the increase in average fitness of the networks to 26. Networks 1, 3, 5 and 7 now have operating profits and can expand faster.

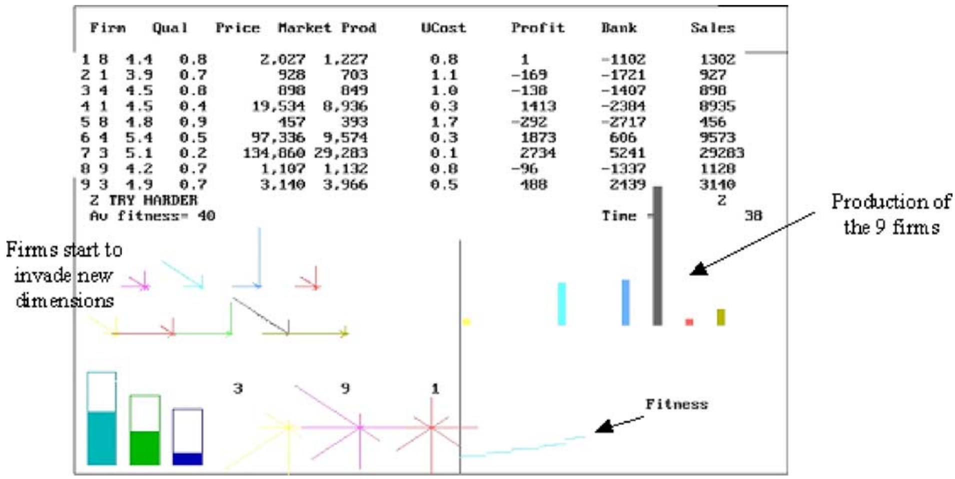

Fig. 7. Here we see that firm 7 is starting to dominate as its product sells and its profits rise. The fitness of the industry has reached 40.

In Fig. 6 we see how the differential amplification of firms is leading to an improvement in the average ‘fitness’ of the industry. These products are better than they were, and also they occupy a larger proportion of the production. If we allow this particular simulation to run on, then we observe a ‘shake-out’ of the industry, and a standardisation as one firm becomes dominant. This is shown in Fig. 7 where firm 7 is starting to dominate. Fitness of the industry (as measured by the ‘fit’ between the products and the customer requirements) has increased to 26.

From this we see how we can model both the ‘internal learning’ of networks concerning the improvement of their products by exploring new products made by reorganised networks, and the external learning of what it is customers will buy. This ‘inner/outer’ learning dynamic is the core issue. Again our model has produced learning on behalf of the industry without explicitly supposing intentionality on behalf of entrepreneurs. The product improvement phase is driven once again by a random exploration, leading to differential amplification within the context of the market.

Also, as the emergent market invades new emergent attributes, an external selection process at the level of society may decide that certain practices (cruelty, environmental damage, etc.) are to be outlawed. Essentially, this would correspond to a society deciding that its own idea of itself does not allow certain types of practice.

We see that the interaction of the different products that are launched leads to fairly sensible structure that does increase the ‘fitness’ of the sector over time. However, this sensible outcome depends on diversity of the product explorations carried out at the lower level. If there is great diversity, then the space is well explored, and the resulting market structure more effective. But if firms merely imitate each other, through fear of risk-taking, then 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 microdiversity lead to emergent ‘knowledge’ at the level above.* Equally, advice to competing firms that leads them to search wider and innovate with more originality will also improve the 'fitness' of the emergent industry. Of course this 'knowledge' is less than perfect. The participants may not know why what they are doing is 'working', but simply that they are. Prediction of exactly what will happen under some new circumstance may also be absent among participants, and yet successful re-adaptation will occur providing there is an underlying micro-diversity that is constantly ready to re-structure the market.

Inside these firms we could now ask the same questions that we had for the market level. How can we create a learning organisation 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 described above. However, explorations into different dimensions of possible innovations could be rejected through a conscious selection process internal to the firm. This would be in order to ensure that the new product would fit the firm's view of its own identity and strategy.

4. Guiding evolution

This section briefly looks at the level below that of the 'system'. In the example of a market then each firm is trying to decide what to do. Instead of throwing a dice to decide on the strategy to use, or the details of a new product, each firm obviously tries to make a choice that is better than random. This is the essence of knowledge—to make a 'better than random' choice. This process can be summarised in Fig. 8.

In this diagram, we show the mechanisms necessary for a firm to respond 'intelligently' to possible threats or opportunities in the environment, Allen [6]. This will primarily arise from customer tastes and needs and from the behaviour of other firms in the market. There are three major parts to this 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 organisation 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.

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, 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 microdiversity 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

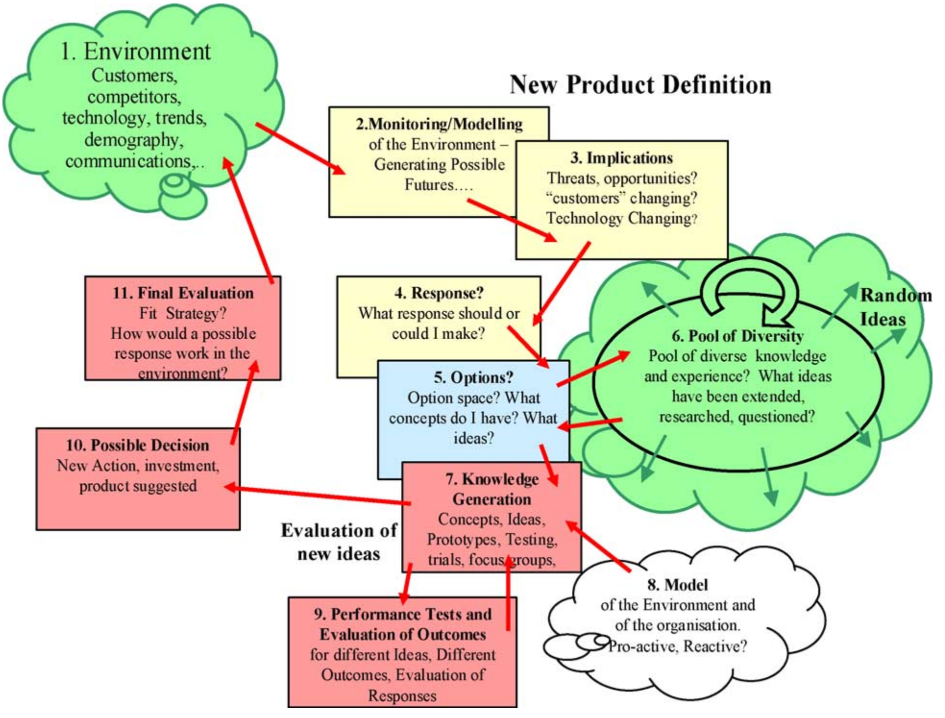

Fig. 8. Guiding evolution. The boxes and arrows constitute an attempt to guide the innovations that are actually launched into the environment from an initially random selection.

by the diverse searches of its diverse firms. All rests upon the hierarchy of structure and its micro-diversity. Within firms, departments need internal diversity and multiple perspectives among its 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 market place 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 multi-dimensional one, with 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). Consequently, the full implications of any particular design change or innovation will only declare themselves over time, possibly long after the decision to 'adopt' has been made. In reality, we can only hope to do better than a random choice, but never to know absolutely what the effects of a particular decision will be. Even if there is in fact a 10% overall improvement in the 'performance' of an artifact, the full series of tests will give a whole series of mixed results, some better and some worse than before. On average there will be a 10% improvement, but this will not be true for each test. Automatically there will be a deviation around this 10% figure that will reflect the technological trade-offs that are embodied in the particular product concept.

A different concept will have different trade-offs. So, we immediately encounter the issue of 'bench marking'. If an earlier design is considered first, then each test dimension can be compared with that of the previous model and from this can derive the idea of a 10% improvement with some features 20% better and others actually worse than the previous model. The real picture therefore is as shown in Fig. 9.

Clearly, the real tests that will be applied to any design that goes on to production and use will have several different stages. Initially there will be the 'bench testing' of various properties. Then will come perhaps problems and issues connected with its ease of production, and finally there will be a series of tests that will affect its overall financial performance concerning its use and impacts in the real world environment. Some of these may be known at the time of design and some may not be. Only when all these have occurred will it be possible to calculate the overall net change in performance. If we perform too few tests, then we may be misled by having accidentally picked either the most favourable or unfavourable dimensions to test. So, we need to know at which point we can consider that we have performed enough tests to successfully discriminate between the performance of any two prototypes.

In our model, we do not know the landscape we are exploring. What we know is simply the performance of the products that already exist. The model allows us to try out various possible strategies for improved performance. The total amount of testing that will be possible is governed by the NPD budget limit. However, this could be divided into different parts, with a wide, shallow sweep across many possibilities, and then a deeper exploration of the promising designs. If we go for a single round, then we may decide to

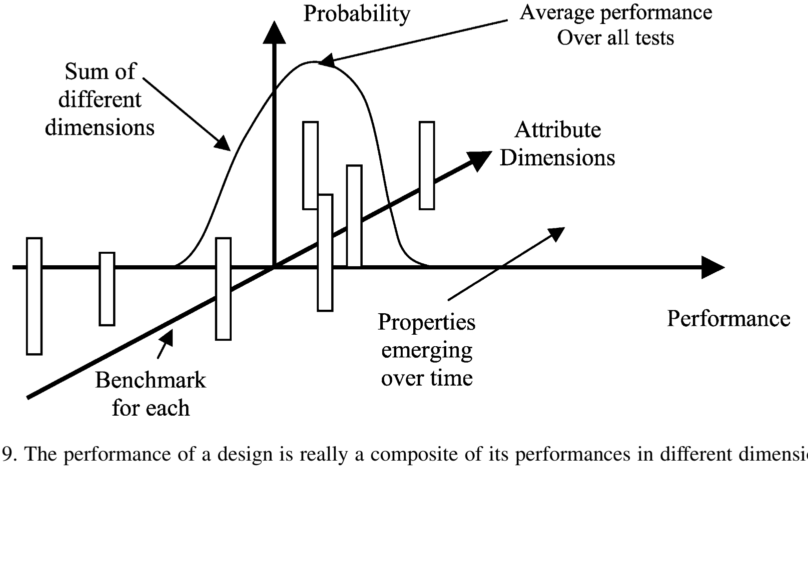

Fig. 9. The performance of a design is really a composite of its performances in different dimensions.

start from the 'best' existing design and see how a concept number close to the best known type behaves. If it is worse, we may then try a concept number in the opposite direction. Otherwise we may believe that we can extrapolate a line of improvement, and try a design that is further along this line. For the concept chosen, and the amount of testing that has been decided, the model will sample the test results from the Gaussian distribution, with a series of positive or negative readings whose length will depend on the money allotted to the test. If the list is short, then the test was cheap but the information possibly misleading.

One is building evolutionary models in which unguided exploration leads to sector evolution and learning and complex organisations. But without any expressed intentionality. Another piece of work concerns New Product Development (but we could say decision-making). This shows that because the 'performance' of a 'product/action/service/idea' only reveals itself over time, as different attributes 'emerge' and give their contribution to overall performance, then a decision must always be made with incomplete information. However, it is made with intention (usually). There is a NPD project that aims at seeing what ideas (random) exist, and in testing out the suitable ones. Suitable with respect to overall strategy that is. The evaluation criteria and testing therefore express the attempt to guide (manage) the emergent processes. Of course, we can only test the easy things (bench testing) but other factors will only emerge over time. We will then be forced to adapt, learn, modify our plans, and to do so in conjunction with others embroiled in the unfolding (more accurately folding) situation. So, to me this shows why the idea of rational decision making is not wholly wrong, in that to choose something deliberately that is bad seems stupid. On the other hand the consequences only emerge over time, as different processes with different time scales are engaged.

The model above shows that successful overall behaviour can be obtained from a competitive evolution of firms that do not necessarily have any 'knowledge', and only learn by random experiences. But what does this mean? By 'running' the system forwards in time it 'evolves' and changes the behaviour of the participating actors from that of random choice to that of structured choice.

5. Discussion

If we define 'knowledge' as being something that structures possible action from random to highly defined, then clearly self-organisation and evolution are creators of knowledge. A set of non-linear equations representing the behaviours of multiple agents will 'run' from initially random choices to behaviours 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 behaviours or 'routines' (Fig. 10).

In our economic models above, initially the firms are chosen with random strategies of profit and qualities, but over time the models run to attractors that are characterised by firms with very particular strategies, as niches develop. If the system is run with a fixed external environment (fixed customer tastes, technology, etc.) and within a set of fixed product dimensions, then eventually it will evolve to one of many possible attractors, and stay there. The emergence of the attractor will correspond to the apparent generation of 'knowledge' in the system, as the model will restrict each agent to a particular, fixed

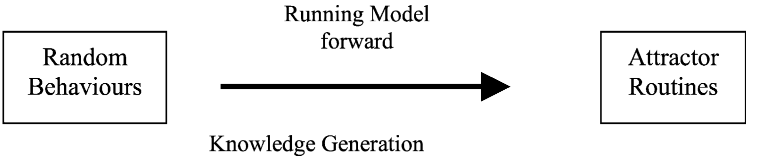

Fig. 10. Over time a self-organising model will generate knowledge as participating behaviours change from random to routine choices.

routine (though possibly cyclic). Reflecting on the meaning of knowledge and of evolution, we might see that the generation of knowledge itself is really simply the ‘running’ of the system. It is the non-linearities of the processes that lead to the amplification and suppression of different possible routines and to the consequent appearance of ‘intelligent’ behaviour, since the system will function to some degree successfully. That is, it will not necessarily find the ‘optimal’ attractor from some particular point of view, but any attractor will be functional given the underlying processes of the system. This view pushes us towards the recent ideas of Ralph Stacey [7], who refers to complex adaptive responses of mutually interacting participants rather than to evolutionary complex systems.

Returning to our reflection of evolutionary models however, we see that running the system automatically leads to the emergence of fixed routines that would appear to embody knowledge. However, this would only be useful for a time, but as the environment changed, the routines and the corresponding attractors would become unsuitable and the particular system would no longer be able to maintain itself. In a changing world, any knowledge that has been generated will lose its value over time and will require renewal. For a system to be ‘sustainable’ over time, it must be able to restructure its behaviours, and to explore the possible potential of new processes. This means that the very boundaries and dimensions that define it must be potentially subject to revision, as fluctuations of behaviour extend beyond any current dimensions. This shows us that although at any particular moment a ‘system’ will appear to have a boundary and to involve certain elements and not others, over time this boundary is potentially moveable, and indeed systems themselves are emergent phenomena. At a given moment, systems are the things that emerge from self-organising processes. Organisations, routines and functions emerge as the non-linearities of processes that interact. Over time, however, they will change as one particular system decays and a new one emerges. In so doing, knowledge is created, used and then replaced.

In order for the elements of a particular system to attain some continuously successful mutual existence, they need to run some way towards the routines of current ‘knowledge‘, but not to adopt them completely. In other words, there must always be some residual exploration of alternative choices and actions, even though these will appear sub-optimal. This will allow the restructuring of the routines as a new attractor becomes desirable. However, it may also be necessary to change the boundaries of the current system, either by including new dimensions within it, or by separating and specialising separate parts. Innovation, re-structuring, mergers and acquisitions are all possible adaptive mechanisms for survival in a changing environment.

What appears to be a fascinating idea is that in running a 'model' forward, we automatically move from a less structured to a more structured distribution of behaviours. It generates a decrease in the entropy of the probability of different decisions, corresponding to an increase in 'knowledge'. In reality, we may not have the equations for the model available, and therefore we must learn by trial, error and inference, but if we can make a conjecture about the equations, then we can run the model to see what knowledge they generate. The Popperian idea of formulating a testable conjecture (by inspiration) and then testing it, can be used either in the practical sense of learning by experience, or in a larger sense by learning whether a conjectured model is sensible.

Limits to knowledge merely admit what has always been true. The future is not determined and the world is irreversible. By ceasing to pretend that we have, or can have, perfect knowledge we can adopt a more sensible exploratory and reflective approach to the world that both responds better to a changing world and also changes the changes that occur. By adopting approaches to decision making and to learning that go beyond the simple 'me too' of benchmarking and of imitation, the creative processes can be broader and can lead to greater micro-, meso- and macro-diversity, reducing ecological and environmental stresses and increasing opportunities.

References

[1] P.M. Allen, Evolution: why the whole is greater than the sum of its parts in: Wolff, Soeder, Drepper (Eds.),, Ecodynamics, Springer, Berlin, 1988.

[2] P.M. Allen, Why the future is not what it was, Futures 22 (6) (1990) 555–569.

[3] Allen, P.M., Evolution: persistent ignorance from continual learning. In: Nonlinear Dynamics and Evolutionary Economics. Oxford: Oxford University Press, 1993, 101-112.

[4] P.M. Allen, Evolving complexity in social science in: G. Altmann, W.A. Koch (Eds.),, Systems- New Paradigms for the Human Sciences, Walter de Gruyter, Berlin, 1998.

[5] P.M. Allen, Knowledge ignorance and the evolution of complex systems, in: J. Foster, J.S. Metcalfe (Eds.),, Frontiers of Evolutionary Economics: Competition, Self-Organisation and Innovation Policy, Edward Elgar, Cheltenham, 2001.

[6] P.M. Allen, The dynamics of knowledge and ignorance: learning the new systems science in: M. Matthies, H. Malchow, J. Kriz (Eds.),, Integrative Approaches to Natural and Social Dynamics, Berlin, Springer, 2001.

[7] R.D. Stacey, Complex Responsive Processes in Organizations, Routledge, London, 2001.