PART I

The development of complexity perspectives

1. An Interview with W. Brian Arthur

Interviewed by Robert Delorme and Geoffrey Hodgson of EAEPE on 8 November 2002 in Aix-en-Provence

EAEPE Could you briefly outline how you developed your ideas on path dependence, positive feedbacks and complexity, and what problems you encountered in convincing others of their importance?

WBA As an undergraduate I had been trained as an electrical engineer so I knew quite a bit about positive feedback. Also, when I started to study economics in the early 1970s at Berkeley my key interest was in economic development and so I was exposed to the ideas of cumulative causation by Gunnar Myrdal, the forward and backward linkages of Hirshman – the sorts of things we would now describe as positive feedbacks. I took plenty of courses in neoclassical economics but they didn't speak to me as much as the dynamics of development did. My PhD was in operations research and that was heavily dynamical (I worked on control theory under Stuart Dreyfus). So my whole background was to look at dynamics. I was attracted to dynamics with cumulative causation positive feedbacks.

Around 1979 I read *The Eighth Day of Creation* by Horace Freeland Judson [1979], which was a rather thick book of the coming into being of molecular biology – a wonderful book. As a result I started to read about molecular biology and eventually about enzyme reactions. I was visiting the University of Hawaii at the time and had eight weeks to do as I pleased. I read Jacques Monod's book *Chance and Necessity* [1972], in English. I started to encounter chemical reactions, where you could get one of two products A or B coming out of the same reaction. They were self-catalysing in the sense that the more A was present the more A would be produced, and the more B was present the more B would be produced. Some reactions seemed to show a tipping possibility – they could go either way and there were two possible outcomes: A dominates, or B does. Such examples were not that well known at the time. This led me to read the work of Ilya Prigogine and eventually of Hermann Haken in physics.

I realized this was a methodology for looking at positive feedbacks. And my key realization was that all this applied to increasing returns in economics. Increasing returns had been looked at since the time of Marshall, but statically. In the static neoclassical theory people simply said that you can get multiple equilibria, some of which could be better than others. That was essentially the whole story. Nobody knew how one equilibrium out of the several might be selected. I realized I could solve this selection problem by re-expressing economic problems as stochastic processes. One set of small historical events would get magnified by increasing returns and would lead to one solution; if you reran history a second time, so to speak, different small events could lead to a different solution.

So, to be specific, what I did in increasing returns was to solve the selection problem – how an equilibrium is selected. I did this by making the problem stochastic and dynamic and watching how small events accumulated and became magnified by increasing returns to arrive at an outcome. I also introduced the theory of how solutions locked in and how non-ergodicity (or historical path dependence) worked.

Around 1981 I started to search for examples and realized that the best ones were in technology. If you had steam cars, petrol cars and electric cars one hundred years ago, then if one of them got more adopted it would get more developed, and possibly lock in. I began to see many examples of technologies that were not necessarily superior but were locked-in. Another thing I looked at was economic geography. It seemed clear that Silicon Valley was around Stanford because of small events in Palo Alto and its environs in the early 1900s. (Federal Telegraph had started as an offshoot of Stanford as early as 1905.) I used to wonder why Silicon Valley wasn't around Berkeley, which was my alma mater. If some key personages had done different things at Stanford or Berkeley in the first decades of the twentieth century, Silicon Valley might have been elsewhere. High-tech companies want to be where other high-tech companies are already and where you have some seed effects or founder effects, little clusters grow like in petri dishes. But these were just examples and I was interested in the theory.

When I brought these ideas out – I wrote them up in 1983 in a working paper at IIASA [Arthur 1983], then produced them in Stanford – there was an enormous amount of enthusiasm on the part of economic historians. Paul David was an early supporter and so was Nathan Rosenberg. So were some theorists such as Kenneth Arrow. But overall there was bitter opposition within economics. Many people found the ideas repugnant – lock-in to inferior outcomes, dependence on history, no unique solution. I submitted my 1983 paper to three economic journals in succession, including the *American Economic Review*, the *Quarterly Journal of Economics* and the *Economic Journal*. It was turned down in each case and for six years I couldn't get it published. In the meantime, Paul David [1985] pushed ahead with his version of the ideas, which became the QWERTY paper, and the ideas started to circulate widely. People knew about my ideas but there was no official version in journal form. After a second submission the *Economic Journal* finally published my paper in 1989 [Arthur 1989]. The hold up, I believe, was largely because economists were unfamiliar with the notion that different random events could lead you to different solutions. This seemed strange – and un-economic. And particularly in the USA we were still in the middle of the cold war and nobody wanted to hear that perfect capitalism could lock you into an inferior solution. Europeans, I've found, accepted that very naturally. But Americans ideologically found it difficult to accept and I got a huge amount of flak in the seminars. Then there was the standard line of resistance; 'If this was such a good idea somebody else would have thought of it already'.

In general there was a lot of puzzlement and I don't think these ideas began to be accepted until I went to the Santa Fe Institute in 1987. At a conference there the physicists turned around and explained to the economists, including Kenneth Arrow, that my approach was absolutely standard in physics, and there was nothing to worry about. If this wasn't regarded as proper economics it was certainly regarded as good physics. After that I could see the economists in the room relax. I started to have a lot of support from first-rate economists. By 1992 the ideas were so well accepted that the fights were about who had thought of what first. So those ideas went in the standard sequence: from ridicule to puzzlement, then to reluctant acceptance, then to total acceptance where everybody claimed that they known that all along. One challenge I had in the early 1980s was to find examples. I was told that this was like the theory of black holes – possible theoretically but no example existed of lock-in to an inferior solution. I used to say Microsoft DOS was totally inferior and totally locked in.

EAEPE Was there a turning point in 1992?

WBA Yes – actually in 1991. It was abrupt. In 1990 I published an article called 'Positive feedbacks in the economy' in *Scientific American* [Arthur 1990]. The physicists helped me publish it and after that those ideas became totally accepted in the scientific community, and in economics. As I said, after that the scramble in economics was to say who had thought of it first.

I surely believe that no idea in science is ever completely original. You can always find wisps of thinking that precede it. I have no doubt Marshall and other people thought along the same lines about increasing returns. What I did was to find a way to solve the selection problem – the dynamics of how an outcome is 'selected' and to make that formal. For Marshall and others the apparatus of non-linear stochastic processes wasn't available to them, so they couldn't do that. I had to do a lot of work in probability theory and to collaborate with good probability theorists like Yuri Ermoliev and Yuri Kaniovski to nail these ideas down.

The last part of your question was about complexity. I confess I hadn't much thought about complexity until I went to Santa Fe. But it turns out that non-linear thinking or positive feedbacks are very much a property of complex systems. You need a mixture of positive and negative feedbacks to make any system complex – to get the possibility of multiple solutions or multiple metastable states. I can't say I developed ideas on complexity. It was more that, *à la* Moliere, I was speaking prose in that I was doing complexity without realizing it. When I got to Santa Fe Institute everybody there in physics instinctively understood what I was doing.

EAEPE You describe complexity as a movement and as a popular issue. It has been criticized as being more of a fad than a profound perspective. How do you respond to this?

WBA It is not a fad. It is actually a long-term movement in science. For about 300 years we have been looking in science at objects from the top down. It is as if we are looking at a Swiss watch under a microscope looking at the mechanisms of how things work. This is very much a Cartesian perspective, looking for the mechanics – the causal mechanics – of how something works. This reductionism has been extremely successful and I certainly don't think it is going to go away. It constitutes virtually all that we know of science to date. But there is another way to do science – a growing movement – and that is to look from the bottom up at how interactions form structures and patterns. Certainly there are many precedents for that, Poincaré's dynamics for example. But looking at the world this way became a serious proposition only when we all got desktop computers. So from around 1980 on, you see in science the notion of looking at elements interacting to form patterns. How do elements interact to form structures? How do structures emerge out of interactions? How do you get anthills? How do you get intelligence among ant colonies out of very simple rules? What happens when billions of stars in two galaxies collide? How do spiral arms form in a new galaxy? These are all questions of how individual simple interactions lead to structures. So the movement is towards looking for causal explanations that work algorithmically and looking much more at process and simple interactive rules leading to dynamical processes that lead to structures.

I'm not sure 'complexity' is a very good word for this movement. But we are stuck with the label. Complex studies look at interacting elements producing aggregate patterns that those elements in turn react to. This seems to be the theme of all studies I've seen on complexity. The elements could be ions in a spin glass, or they could be cells in the immune system, or they could be cars in traffic. But each element is behaviourally reacting in some way to the overall pattern, or a local pattern, that those elements are co-creating. It is very much a self-referential system. And for me that's what gives the juice to complexity. Elements are creating that which they react to. The outcome could be incredibly simple. But often, with the right combination of positive and negative feedbacks in the interaction, the outcome is quite complex with unexpected structures forming. The idea of elements forming a field that they in turn react to has been looked at for a long time in equilibrium terms both in physics and economics. But looking at this dynamically and heterogeneously, element by element, is new. It is hard to do this analytically, because patterns form from heterogeneous interactions. Once computation came along in the sciences, we could recreate such systems on our machines and it became natural to look at how elements interacted to form patterns.

If you call this approach complexity, it is a long-term movement in the sciences. The sciences underwent a huge switch in the middle of the 1600s, about the time of Descartes, Galileo, and Newton, when explanations went from being geometric to being equation based. Now we are making the switch from explanations being equation-based to algorithm-based. This change going to be slow; it will build on top of what is already there; and it will not replace reductionism. And it will not go away.

EAEPE Isn't the concept of emergent properties central to this discourse?

WBA Yes, emergent properties are absolutely essential. Emergent properties are difficult to define because when you see an emergent property you immediately think of a simple explanation for it and then that property seems obvious. Take our stock market model. An emergent property is that you get periods of quiescence where prices don't vary much, followed by random periods where there is a lot of volatility. An explanation is that changes in forecasting behaviour can cause the overall price system to change, which causes other people to have to change their forecasting. So you get small and large avalanches of change cascading in such a system. Quiescence (a period of small change) begets quiescence; and volatility (a period of major change) begets volatility. That's an emergent property and it arises from simple rules. The problem is, that when you look at any emergent property in hindsight, it seems to be trivially explainable. Yet you may not have been able to deduce it unless you looked via a computer into pattern formation. Generally, emergent behaviour is something that's 'emergent' or surprising only when it emerges. After it emerges, it seems pretty obvious.

EAEPE Coming back to the complexity approach as a movement, is it a collection of varied insights, or does it – in actuality or in potential – represent a coherent theory? In a different way, is there or can there be, a universal theory of complex systems?

WBA Is there is a universal theory of complex systems? I don't think one has been reached yet. But it often takes many decades for theory to emerge in many new fields. Think of thermodynamics or even mechanics in the 1600s. We're often asked at Santa Fe, 'Why isn't there a coherent theory of complexity yet?' Well, it'll take time. And I am not sure there is a theory – we may be looking down rabbit holes where there is no rabbit. We can say this: if you are looking at systems whose elements create the overall field those elements are reacting to, we do see many common properties. Sometimes these self-reactive systems will reduce in simplicity and lockin to a static and unchanging pattern early on. At other times they become chaotic. But there is an area in between chaos and order, the so-called 'edge of chaos', where certain structures start to appear. There you find similar structures at many scales, no stasis or equilibrium, patterns that lock in for a long time then disappear, and long correlation links. If there is a coherent theory of complex behaviour, it will be found there. That is not yet a universal theory, but commonalities are starting to emerge in this area. Until we find such a theory, complexity will simply be a common approach looking at elements reacting to the patterns those elements create.

EAEPE When we look at the literature on complexity, it seems divided into several camps, with minimal communication between one another, even between those groups writing in English. For example, there is the von Foerster–Maturana–Varela autopoiesis tradition, Prigogine on dissipative structures, the Checkland approach in Lancaster (UK), the (US) Santa Fe school, and so on, without mention of the traditions writing in French and German. To what extent does the compartmentalization of complexity perspectives constitute a problem, and how could it be resolved?

WBA There is compartmentalization, you are absolutely right. At this stage this is actually a good thing. It would be a great mistake if this emergent approach to science got narrowed down too early. I think it is great there are so many different approaches. That means that the subject isn't locking in and narrowing too early. But there is an unfortunate tendency in each of these areas – Brussels school, Strasbourg school, Santa Fe school – to believe that they are the only place thinking these thoughts, which is of course nonsense and there is not enough citation of other schools. That's just pure ignorance. It's not I think lack of goodwill. What we do need is more information and interaction among these different centres, but I certainly think it is too early to narrow things down. There should be lots of variety to apply selection on. There is another sort of compartmentalisation too. Some of this work is coming from computer science and some from physics, some from evolutionary computation, and still other work coming out of economics. So besides these various geographically-based schools, the disciplinary approaches differ. But all this variety is wonderful.

EAEPE You have argued that agents in the economy are in a 'Magritte world', where the distinction between subject and object 'often blurs'. Can you elaborate on this idea? How have others reacted to this insight?

WBA This is slightly difficult to explain. What particularly fascinates me is that in the economy there is no action taken without some expectation of the result. A farmer, acting alone, has some expectation about what the weather will be like in the growing season and may plant soy rather than corn. That's simple, and there is a simple dichotomy here between subject (the farmer) and object (the decision environment). In most of the economy, however, we are dependent upon *other* people's actions to form the future outcome. Let me give an example. I live in the Silicon Valley, and if I am an entrepreneur there, I have to decide now on whether to go ahead with a high-tech product based on what I think the market will be like in two years' time. I am trying to forecast a market that will consist of my product and other firms' products competing in two years' time. But those products – that situation – in two years' time will be the result of other people sitting down today and also trying to forecast what the situation will be. In Keynes's phrase, I (and others) are trying to form an opinion about [something that depends on] what the average opinion will be. In other words I am trying to form a forecast or expectation about an outcome (and everybody else is too) that is a function of my and others' expectations. This blurs the distinction between subject and object. I and others are the subject, and the object we are trying to forecast is a function of ourselves – the subject. You cannot separate subject from object in this case.

This leads to a fundamental indeterminacy that economics hasn't addressed properly. You can't act without forming some expectations of the future. But your and others' choice of expectation depends on your and others' choice of expectation. The situation is self-referential. There is no logical or deductive way to settle this.

There *is* an analytical dodge you can take. You can ask what forecast, if everyone adopted it, would lead to actions that would (on average) validate that forecast. That is the rational expectations approach. It might work in trivial cases. But in most situations rational expectations are highly singular, like a pencil balanced on its point – logically possible but unlikely in reality. The reason is that if anyone deviates (from rational expectations), or anyone thinks anyone will deviate, or anyone is not capable of calculating the rational expectations solution, or if anyone thinks someone is not capable of calculating the rational expectations solution, you are back to subjects trying to figure out how subjects will figure out. There is no deductively logical way to do this. This situation does not represent a set of 'measure zero' in economics, a trivial special case. It pervades economics – financial markets, producers' markets, foreign trade. Most agents' decisions in the economy involve other agents, and all must involve future value. So this is a fundamental indeterminacy in the field.

EAEPE What you are saying, I think, is a proposition – which is quite well known in social theory – that agents and structure are mutually constituted of one another. But at the same time there are agents and there are structures. Like in your El Farol bar example, which illustrates presumably the points of view you are making. There is still an El Farol bar, and there are still individuals deciding whether to go to the bar or not. But would you agree that there is still a distinction between agents and structure?

WBA I constructed the El Farol bar problem to illustrate to myself, at least, that there was such a problem. In El Farol, agents are primed to forecast attendance at the bar, which is a function of other agents' forecasts. (In El Farol, agents show up only if they think less than a certain percentage of others – 60 per cent say – will show up. Otherwise they stay home.) But any common rational expectations forecast is immediately negated. If we all say shared a rational expectations forecast that predicted that next week 74 per cent would attend, then nobody would show up, negating that forecast. I constructed El Farol to close off the rational expectations loophole. Agents can't act deductively here. So what do they do in reality? In such situations I believe people act *inductively*. That is, they try different hypotheses of what's going on, different forecasting strategies – much like statisticians do on a new data set – and act on the most accurate of these. If their hypotheses don't work well, they generate new ones. What forecasting method they choose alters attendance. So the forecasting methods are trying to be 'fit' in a mini-ecology of forecasting methods. This mini-ecology changes over time. It's fun to watch on the computer.

The reaction to this problem so far in economics has been mild interest – and some bemusement. I don't think economists quite know what to make of this yet. The physicists picked up on El Farol and turned it into the minority game. I'm told there are now around 150 papers on this version of the problem. The game version is valid, and of course it has a Nash solution. But that misses my original point about El Farol. The problem is structured so that if we postulate a 'correct' bar-attendance forecasting machine that everyone coordinated on, their actions would collectively invalidate it. Therefore there could be no such correct machine. The indeterminacy here is not just a behavioral failure of human agents to live up to an economic 'theory'. It is unavoidable.

EAEPE What do you think are the main challenges and obstacles to complexity research today? What research strategy would you advocate for further advance in this area?

WBA Well, I think we have just gone through an early exploratory period. As with the Lewis and Clark expedition, we know what the continent looks like – or so we think. There are several areas I'd like to see research in. We are moving into a period where what we need to see is whether there is any commonality of solutions. Is there any universal theory of complex systems? Are there universal properties we can attribute to complex systems? What structures tend to emerge that we might see as general in complex systems? Is there any commonality of phenomena? These are early days as yet. There are also many exploratory questions. We are shifting from an equation-based to algorithm-based approach in many fields, as I said earlier. And algorithms are able to introduce more realistic assumptions into theory. One of these is heterogeneity. So in economics and in physics, this should open the way to exploring the consequences of heterogeneity. This, I believe, will not just be a trivial addendum to standard homogeneous theory. Heterogeneity in any dynamical situation leads to an evolutionary approach. Heterogeneous actions or strategies or expectations try to compete in a overall situation (or ecology) they co-create. When you introduce heterogeneity into economics, it immediately becomes more like modern biology and less like nineteenth century physics. This is not a trivial extension of standard theory. It's a different way of thinking – one that Marshall would have approved of, I am sure.

EAEPE What suggestions would you make to help the further diffusion of complexity ideas among academics, especially those ideas that have no obvious expression in mathematical form?

WBA I've been arguing earlier that complexity isn't so much a theory as a point of view. Given that it is a point of view, you may suddenly click into it and your way of thinking about everything shifts. There's a before and an after. Before, I thought of economics in the conventional neoclassical view. After, I thought of it in ecological or process terms, and everything appeared very different. I don't know that you can convince people that such changes in view are necessary. It's not so much a lot of theory as a change in point of view: from looking at stasis and equilibrium and looking in a reductionist manner, to looking at process and emergence of formation and looking in a holistic manner. That comes naturally to some people. Others find it difficult. It's a sort of gestalt switch.

What convinces people is seeing novel phenomena in standard problems that can't be explained in the standard way. So we need to ask: are there phenomena that these new studies are showing? If there are, then people will have to sit down and master this point of view. I would claim that in many of the studies I have seen there are. I am thinking of Kristian Lindgren's game theoretic paper, where he shows the phenomenon of endless novelty in prisoner's dilemma tournaments [1992]. Lindgren's system never settles down. In the artificial stock market model I did with John Holland, Blake LeBaron, and Richard Palmer, we found a phenomenon of change begetting change – that is, changes in expectations causing the market to change, thereby causing others to have to change *their* expectations. Avalanches of change at all scales that ripple through the system. And we found these produced GARCH-like statistics. Other work I like is Tom Sargent's inductive approach to decision making. And the work of Per Bak and various economists, showing such phenomena as self-ordered criticality in markets. And there is much work showing how standard equilibrium solutions emerge from 'ecologies' of individual actions.

There is a shift in view in all the sciences, a Kuhnian [1970] transition, if you like, from a static to a pattern-formation approach. The old theories will not be invalidated. Rather, the new theory will come up with certain things like perpetual novelty of outcome that otherwise cannot be extracted from the old theory and new people coming along will have to study this and become familiar with it. So there will be a time lag in the viewpoint in economics shifting from an equilibrium one. It will be resisted, but that's true in every shift.

I am not keen on the label 'complexity' for this. What is really happening is the birth of an Out-of-Equilibrium Economics. This new approach specifies what happens out of equilibrium. For that you need to include heterogeneity of agents, allow agent-based behaviour, and allow patterns to form – possibly in the computer. In this larger economics, equilibrium becomes a special case. This is an inevitable widening of economics, and it will not go away.

EAEPE You once declared in a talk that the problem with economics was not physics-envy, but mathematics-envy. Could you comment further on the state of modern economics and its preoccupation with mathematical formalism?

WBA From the time of von Neumann onwards (that is from the 1930s and 1940s), mathematicians coming into economics found they could clean up a lot of problems that economics faced. Before that concepts were sloppy and whole subjects such as general equilibrium theory were a mess. People like von Neumann, Wald, Arrow and Debreu – the great mathematical economists – came in and cleared up a great deal. We owe them a lot. Arguments are now clear in economics. This sort of work, much of which was done in the 1950s and 1960s by the Cowles Foundation was controversial at the time – Kenneth Arrow himself told me that. But it was a badly needed cleaning of the stables in economics. The whole subject needed to be logically tightened up.

All this was fine, but in due course economic theory began to be confused with mathematics in economics. Theories are arguments about how the world works, and they derive from observation and insight. Mathematics is a means of expression for theory, just as language is a means of expression for thought. But it is not theory itself. Moreover, as with all fields of enquiry about behaviour in the wild, the economy refuses to be captured by a single set of axioms from which all truths can be deduced. So the mathematization of economics succeeds as a badly-needed clarification of thought, but it has not generated new theory in the sense of fundamental insights. In this it has failed. Mathematics is a good language for theory, but it is not theory itself. If you want theory, look to people like Kaldor, or Schumpeter, or Allyn Young – economists writing before 1950. These were people who were interested in the actual economy.

EAEPE Or Marshall?

WBA Or Marshall, indeed. I think Marshall would have regretted what has happened in economic theory. There's clearly high prestige to doing mathematics in economics as in any other field. But overall I do not see what novel theories or novel economic insights mathematics has delivered. What it *has* delivered is the fine print under which certain phenomena or certain insights hold. And it has done this well. But all fields go through certain approaches from time to time that are held to be fundamental. Philosophy attempted to base itself on logic in the early twentieth century, then on linguistics – and both these failed. These are necessary stages of adolescence that a field passes through. Originally they clarify much and help discard dross, then they become increasingly baroque, then they become sources of orthodoxy and of power over careers. Few people then dare to oppose it. In due course, mercifully, the field passes on. English departments are going through this these days with the dominance of postmodern critical theory. Don't get me wrong. I love mathematics. And I admit I have spent a fair amount of my career proving theorems. But regrettably, dressing up an argument in economics as a theorem is not mathematics. Nor is it good theory. Theory is basically a set of explanations for certain phenomena. It is a re-construction of phenomena from simpler ideas; and normally it can be expressed in English (or in simple mathematics). This construction of course needs to be logical and to fit properly together, so some mathematics can keep us honest in this. But the idea that theory equals mathematics is not just regrettable in economics. It is wrong. I am sure that it comes from a sense of inferiority to the so-called 'hard sciences'. We shouldn't feel any inferiority. Economics is at least as difficult as physics and it doesn't need to dress itself up artificially to prove this.

Let me put this more positively. I believe that economics is a science. It is a science in the sense that it is well-organized body of knowledge, it is logically sophisticated, and it provides a deep understanding of what goes on in the economy. But I believe that economics needs to become more seriously scientific. By this I mean that it needs to take itself seriously as an explanatory body of theories about how the economic world works, and it needs to look more deeply at actual behaviour and actual phenomena rather than try to shoehorn these into some fixed format dictated by 'theory'. Further, the structures that define the economy continually change and reshape themselves – and therefore constantly need reinterpretation. Therefore over the long term economics continually needs to shift its understandings. This means that our interpretation of the economy is forced to keep changing. To make economics seriously into a science means taking the world of phenomena seriously. This means moving away from shoehorning behaviour into the strictures required by abstract mathematics.

EAEPE What practical or career advice would you give to young researchers, hoping to engage in complexity research, especially those forms of research that are not centred on computer simulations or non-linear mathematics?

WBA It's a good question and a difficult one to answer. The career advice I would give to young researchers is simply to explore. This is a young enough area that it's still in a time of exploration. It's an extremely exciting time. Everything, almost everything, remains to be done. There have been a lot of new ideas in the last ten or 15 years, in the work of Chris Langton, Stuart Kaufmann, John Holland, and others – artificial life, genetic algorithms, all sorts of interesting stuff. So this is a new area and is totally open to exploration. It's like where radio was at the time of Marconi or where molecular biology was in the 1950s. There's always an illusion that a half dozen people have cleaned up the entire area. But that's not true. I think almost everything remains to be done. There have been a few initial insights, but that's all. Fearlessly explore, is the advice I would give.

EAEPE What is the place of the computer in complexity research, according to you? Is it one hundred per cent? If it is less, what room would you leave to other techniques?

WBA For this sort of exploration computer simulation is extremely useful. Backed up of course by a grounding in non-linear mathematics. In complexity research, the desktop computer is our lab. It's where we see new things, and where we try ideas out. It's where we look at how things unfold and evolve. It's very difficult to do that without being able to program a computer. Mathematics is useful (as long as we don't mistake it for theory). Especially, non-linear mathematics can be useful for analysis. But we usually wheel that out to tidy up our insights and not to use it as much for exploration. So my advice to people who are not centred on computer simulations and non-linear mathematics would be to learn some. The idea is to look at miniature 'situations' or thought experiments, not just mentally but on the computer, to sharpen intuition and to look for phenomena. Lots and lots of small 'back-of-the-envelope' experiments. It's terribly important to have questions to ask. To have hypothesis, to look at things, to get insights. Above all the knack in this area is to cook up clever experiments. Those don't come from computers. Blind research on a computer is not useful.

Someone asked me once how I would define complexity and I said: 'Complexity is what happens when you give Charles Darwin a computer'. I meant it seriously. Imagine what Poincaré could have done if he could have seen his dynamical systems unfold. Darwin had to restrict himself to talking to pigeon breeders and cow breeders to find the results of natural experiments, or look through palaeontological evidence. But imagine; we can actually physically allow really simple systems to literally evolve on our machines. Virtually everybody I know who has done original work in this area has sat up at night, often far into the night, staring at computer screens wondering, 'What if I change this, what if I change that?' This is true of Wolfram or Chris Langton. In each science today, there is a theoretical arm and a experimental arm, say in biology or physics. But there is also a computer-based experimental arm. That is true even in mathematics.

We have a good idea of what proper theory consists of, and of how experimental work should operate. But so far we haven't paid attention to any rules for computer experimentation. Should experiments be reproducible? Should they be available on the Web? How can they be verified? There hasn't been much thought yet on this.

EAEPE Isn't your last point an argument for collaboration between complexity workers and philosophers of science, as a live development in the future?

WBA What do you mean by that?

EAEPE I mean, the kind of epistemological questions you're asking about the status of explanation in simulation.

WBA Yes. Absolutely. I think there's an enormous amount of work by the way for philosophers in this area. Basically, if it's true that science is becoming algorithmic and process oriented, and looking at the results of simple rules interacting in a holistic way, then I think that this is an enormous field for entry by philosophers of science, to make sense of that. The people who are doing the actual work on computers are either not interested in the meta-level, or they are too busy trying to make further discoveries to think much about it. So there is a role for re-interpreting what science means in this algorithmic era. I think science has been interpreted in the past to be what it was in the twentieth century, which is rather top-down and reductionist, static equilibrium-based, and so on. We have very little idea yet as to what science is going to mean in this new way of thinking. I would like to see some good philosophers of science interpret this new way of doing things. We need a new Popper, or a Lakatos, or a Thomas Kuhn, to make sense – Or a Geoffrey Hodgson to help us here!

EAEPE Is there a question we didn't ask you and that you would have liked to be asked?

WBA I think maybe the question closest to my heart is, 'What might the future of economics be, based upon this different approach?'

I think the economy is a natural area for looking from a complexity point of view. In any area – trade theory, development theory, consumer theory, or production theory – the economy consists of agents that are reacting to the overall pattern those agents bring about. So the economy is naturally a complex system and is best viewed as in process. Everything in the economy is naturally unfolding, and strategies may never come to an equilibrium, and one structure forms on top of another. But to apply this process approach to economics is more demanding than in physics because the elements in the economy – unlike ions – have both strategy and foresight. So there is challenge here. I've often asked myself if Adam Smith or Alfred Marshall came back to life and were parachuted into the twenty-first century, what would they be interested in? I believe they would be interested this new movement: doing economics from this out-of-equilibrium process viewpoint.

Of course the standard neoclassical stuff will not go away. It is classic in the sense that Shakespeare is. It's beautiful and it contains deep, deep truths of the economy works. But on all questions of how things unfold, such as what effect does technology have on the economy, how do economies develop, how structures form, how do patterns and institutions come about, we do not yet know much. This new way of looking at the economy will add considerably to what we know. It is not a fad. Complexity-based economics means looking at the economy out of equilibrium. This is not an adjunct to equilibrium economics; rather equilibrium economics is a special case of this new economics. So as we enter this new century, we are beginning to look at the economy as something that is structure building, perpetually novel and exploratory rather than static, in equilibrium and based on deductive rationality. We are beginning to see the economy as organic and coming alive. It's an enormous shift and it's exciting to glimpse some of what might come.

EAEPE Finally, can you same something about the implications of your complexity ideas for policy?

WBA Well, two points here. One is that this point of view makes you aware that there are natural structures that want to form in the economy. If you allow a process to take place, there are natural patterns for it to fall into. You may prefer one pattern over another and you may want to nudge the system into that pattern rather than another. For example in the 1960s Britain planned new cities. There would have been organically natural places for these to form. But the new cities like Milton Keynes weren't put in natural places by planners and didn't work very well as a result. In cases like this you want to look for the natural outcomes the system would fall into and allow one of these to happen. Not an invisible hand, not a heavy central-planning hand, but a nudging hand. This means being aware of the natural unfolding of a process and influencing the system early on to go on a preferred direction.

The other point is that when you start to think in terms of process, you realize that process is inherited history. A history that can influence outcomes. If you suddenly release a system, as happened when Soviet Russia became Russia in the early 1900s, you release this inherited history. Static equilibrium thinking might lead you to think that a big bang – a sudden freeing of markets – will bring about a perfect capitalist solution. But from a process point of view, you realize the system inherits procedures, institutions, ways of thinking, and behaviours that are extremely well adapted to the previous system. As a result you get a continuation of the previous system working in a new environment. An evolutionary organism doesn't change itself suddenly when the climate suddenly changes. It slowly adapts its previous structure to the new circumstance. This seems a trivial insight. But it's important. An emphasis on process and adaptation would have been very valuable in Russia or Eastern Europe in the early 1990s, where the welfare of hundreds of millions of people was affected.

REFERENCES

Arthur, W.B. (1983), 'Competing technologies, increasing returns and lock-in by historical events', IIASA Paper WP-83–90, September.

Arthur, W.B. (1989), 'Competing technologies, increasing returns and lock-in by historical events', *Economic Journal*, **99**, 116–31.

Arthur, W.B. (1990), 'Positive feedbacks in the economy', *Scientific American*, **262**, 92–9.

Arthur, W.B. (1994), 'Inductive reasoning and bounded rationality', *American Economic Review Papers and Proceedings*, **84**, 406–11.

David, P.A. (1985), 'Clio and the economics of QWERTY', *American Economic Review*, **75**, 332–37.

Judson, H.F. (1979), *The Eighth Day of Creation. Makers of the Revolution in Biology*, London: Cape.

Kuhn, T.S. (1970), *The Structure of Scientific Revolutions*, 2nd edn, Chicago: University of Chicago Press.

Lindgren, K. (1992), 'Evolutionary phenomena in simple dynamics,' in C.G. Langton, C. Taylor, J.D. Farmer and S. Rasmussen (eds), *Artificial Life II*, Reading, MA: Addison-Wesley, pp. 295–312.

Monod, J. (1972), *Chance and Necessity. An Essay on the Natural Philosophy of Modern Biology*, London: Collins.