Seminar Notes On ‘Harnessing Business Complexity through Agent-based Modelling’.

Abstract: Agent-based modelling (ABM) is a 'bottom up' approach to understanding business complexity. Unlike a mere process description, agent-based modelling captures emergent phenomena and provides a natural description which can include learning. It is scalable and flexible and will tell you where problems might occur and suggest ways of dealing with them when they do.

Dr Bonabeau discusses some basic principles of agent-based simulation and how these are applied to 'real world' situations.

(These notes constitute a synopsis of the presentation and discussion. Dr Bonabeau's slides are referred to in the text and a full set can be found at the end. Because of sound recording difficulties, only Dr Bonabeau is referred to by name).

Director of Complexity Research : Eve Mitleton- Kelly
London School of Economics
Houghton Street
London WC2A 2A

Presenter : Dr Eric Bonabeau, Icosystem Corporation, Boston.

Compiled For The L.S.E.
by Geoffrey J.C. Higgs 25/2/03

First Session

Some events in business organisations are a bit like 'sunspots', we know there's a probability of them happening but we just can't say exactly when or how big they will be. Sunspots have undesirable effects on telecommunications and scientists have been trying to predict them for many years. Ten years ago I was in a team at France Telecom, looking at different ways of tackling this problem. Since sunspots affect the electronic composition of the ionosphere, we were trying to find an optimum range of frequencies for Hertzian transmission.

Traditional techniques for predicting sunspots have relied on statistics; numerical data collected over a time period. Scientists then look for patterns and try to extrapolate them into the future. Other techniques, including the ones we were using, were based on neural networks which were non-linear, but their drawback was that they were what I call 'statistical black box' models in which the dynamics of the system remained opaque. The real problem with this approach is that such models cannot take account of non-stationarity (stationarity is where a process can be determined by a given set of numbers). The models rely on capturing the physics of the system at the macro level, and where a relatively small amount of past statistical data can lead to a prediction of the future, the models are mechanistic STATISTICAL and relatively cheap. However the question we are really asking MAY WANT TO ASK in the case of the sun is basically, "How does the sun work?" which is a much more fundamental question than the frequency of sunspots next week or next month. The question of how sunspots are produced is about the mechanisms behind the production of sunspots and a simulation providing an answer would be a lot more expensive than a statistical model BUT WOULD ALSO BE A LOT MORE POWERFUL TO DEAL WITH NON-STATIONARITY AND NEW EVENTS IMPACTING THE SYSTEM’S DYNAMICS.

Slide 2 - 'Why agent-based modelling?' 1.

Agent-based modelling (ABM) attempts to model things from the 'bottom up', and involves simulating the interactions between a number of basic elements or constituents. It can be very gratifying in terms of the value it creates, but it's often hard building models of systems in the real world so the client has to understand the cost/benefit ratio of the work. Nevertheless if we go that route it has a number of very big advantages. First we can often get leverage from experts leading to a mass of data on which to base our model, secondly it is less stringent on initial data requirements and in the second part of my presentation I will show you why. Building a model can point to data requirements and help us to determine which data is important and which is not so important. A mechanistic model based on statistics may capture the deep intrinsic properties of a system but if the environment changes the model will have to be completely rebuilt. Agent based models enable us to capture the new dynamics resulting from environment change by simply feeding in the new information.

Slide 3 - 'Why agent-based modelling?' 2.

If we transfer this approach to a business organisation then again we can draw a comparison between a statistical model of a process and a bottom up understanding of how a business works in terms of the human agents that make it up. Modelling a business in this way is especially valuable if we want to play 'what if' scenarios which can either mean "What happens if I change some of the stimuli from the external environment" or "What happens if I change the way the organisation itself works?" How, for example, could I reorganise the supply chain to avoid bottlenecks? There are many interesting problems that can be solved concerning the dynamics of a system but again we might contrast an 'aggregate' level approach to one which is bottom up and deals with the constituent units of the system. Take the way people shop in a supermarket. In an aggregate approach we might use fluid dynamics equations to model the density of shoppers. So for example, we might have an average density of 0.3 people per square metre round the fruit and veg aisle, but this doesn't tell us why people came to form that configuration and it won't capture any emergent properties. And, importantly, the clients eyes are sure to glaze over if you show them several pages of maths equations. Aggregate level modelling is cheap and it may capture the physics or the interaction dynamics but we lose a lot of information by using it.

A good example which makes the point here is Per Bak's 'sand pile'. Imagine that there is a steady stream of sand running out of a container to form an increasing pile. The avalanches that are created on its sides are not regular, there are some big and some small. Now we could try to model the sand pile using hydrodynamic equations, but we wouldn't capture the interactions between the grains of sand and we would lose highly non-linear and significant events in the life of the sand pile, namely big avalanches. But if we model the sand pile starting at the level of the sand grains we can capture some very complex features. Per Bak's model was appealing because it provided a description of the avalanches though it wasn't in the traditional physics mode. And there are important analogies with business here. Small avalanches in the sand pile are like small losses of money in business, but the past is not very good for predicting the future. Barings never made any big mistakes before they disappeared down the plug hole. If a CEO makes a policy decision he or she may feel that it is a small one but it can have ripple effect through the interaction of different parts of the business so the resulting effect is very large.

I'm not going to talk about policy decision-making here though agent-based modelling is a powerful tool for it and we have to remember that in the case of human organisations constituent units learn and adapt. I'm often asked to build models for people who don't understand what a model does and what the model can teach you is crucially dependant on the amount of work that goes into it and the quality of the data. It's a metaphor and since I think an awful lot of policy decision is based on metaphor. If you don't have the right metaphor it can be dangerous.

Agent-based modelling is not really a technology it's a mindset (learning approach?) and I have tried to give you a bit of the philosophy behind it. Modelling is not an end in itself though it may be a means to an end. And there's no such thing as a general purpose model. We build a model to address a specific issue or set of issues. But computing power today even compared with what we had ten years ago is enormous. My company's distributive computing system means that I have access to 50,000 machines and can run many programs as screen savers. I have more computing power now than the Pentagon had five years ago.

<u>Slides 4 and 5 - 'Agent -based Modelling'</u>

Questioner 1 : What do you mean by emergence effect? I think there are two points of view on this subject: that some new quality emerges when some parameters are not additive. For example in a simple situation it is an emergent fact that (in the logical system of maths) 2 plus 2 = 4, but in a new emergent effect 2 plus 2 might equal 6.

Eric : OK I'm talking about the occasion when 2 plus 2 equals 6. So the simple definition of emergence is almost circular. It's behaviour that happens at the collective level that cannot be easily connected to the level of individual behaviour. You need to capture the interaction between the basic constituent units to get access to this property. I've been involved in a lot of debates about emergence and I'd rather stick with something that is boring but doesn't raise a big issue. So my view of agent-based modelling is that it is the 'bottom up' mindset, and I'm not wedded to any specific technique when it comes to modelling. A lot of people contrast agent-based modelling with equation-based modelling though if I model a system from the constituent units using equations, that's still a bottom up approach.

Questioner 2 : I very much believe in using that approach and I'd like to ask how you actually convince policy makers that they must use a bottom up approach when they think about and operationalise matters. Its very difficult to get top level management to consider (the consequences) of what the person in the office is doing.

Eric : Well I haven't found the right language to convince people who don't know what I'm talking about before I go in and fortunately all my clients have been ready to go ahead with this approach before I started. I think the answer to your question is in the presentation and if it doesn't work then you don't have any more weapons.

Questioner 3 : May I ask a more specific question? If the older techniques, using finite elements and linear programming with or without integer solutions is limited, is there a sort of phase transition as you go from those models to agent-based modelling. Of if you talk about agent-based models are you really talking about anything that in some sense can be quantified.

Eric : Well I think there is continuum in terms of the transition of models and some people argue that finite element methods are a precursive basic minimum. I have no problem with that except that I think in agent-based modelling you also have the idea that you are modelling constituent elements of a system and that implies that such units make sense in that they actually do things. You are modelling a system by looking at the smallest possible things of the system that are relevant to what the system does. It could be a white cell in the immune system, an ant in a colony, a car in a traffic jam, a person in an organisation or even a grain of sand in a sand pile. I'll argue that the sand pile is actually an agent-based model because the grains of sand do things. But yes, it is a continuum. If I use equations to model an aspect of behaviour then it is still an agent based model.

Another thing which is very appealing to a number of people is that it is often technically very easy to program an agent based model. The universal appeal of complexity science is that you get order or pattern from simple rules and you only have to write a few lines of code. So people tend to think it's very easy to do agent-based modelling and at some levels that's true. But good modelling is often hard and takes experience and agent-based modelling is even harder because you are trying to manifest properties at the aggregate level and you don't know what the mapping is between the lower level and the aggregate one. So I think it's harder than traditional modelling though I would prefer not to be challenged to define what traditional modelling is.

Questioner 4 : How do you convince the client that it it's still not black box modelling?

Eric : That's a very good point and will come back to that later.

One of the reasons for using agent-based models is precisely the fact that it's a very natural description and therefore if a manager or a top executive wants to make the effort to understand how the model works it's orders of magnitude easier for him or her to do so, rather than with a bunch of hydrodynamic equations for example. If you talk to the CEO of Macey's or Sainsbury's and you say this is the equation that covers the dynamics of the density of shoppers at the fruit and vegetable isle in your supermarkets he's going to have a hard time. On the other hand if you say "Here is the model, this is what it does. This is the profile of a typical shopper - 45 year old male shopping for alcohol. This is the beer that he didn't intend to buy but due to the promotion he has", he understands. So one of the winning points of agent-based modelling is that the data is easily translatable.

I tinker with models and tweak this or that parameter and when it replicates the dynamics of the real world system it must count as a good explanation of what's going on. One of the most important aspects of this work is estimating the relevance of the model . That's what econometricians do if that word is not sacred. About 30% of our company resources is dedicated to building an econometrics framework for agent-based models. How you connect real world information to the model is an important issue if agent-based modelling is to make a significant contribution.

Agent-based modelling is flexible and scalable. Flexible in the sense that you can tune the level of complexity in terms of heterogeneous agents and the ways in which they interact, and you can add learning and adaptation. You can do this simply by adding the extra features to the agent classes if you do your basic modelling well from the beginning. And it's scalable in the sense that you can see what happens if you increase the size of the system: increase the number of shoppers for example.

Because an ABM captures emergent phenomena it enables the client to build a mapping from policy space to performance space. In other words we can say, 'Here is a set of interventions that will enable you to influence the system and produce desired effects'. You want some system out there in the real world to behave in a way that is consistent with your desires. The military in the US is very open-minded in terms of looking for and applying new techniques in this direction and one of the big buzz terms is 'effects based operations' . You don't measure the performance of an operation against the number of bombs that you've dropped because there is no linear relationship between that and the outcome. You have to measure the performance of an intervention by how well it achieves it's mission in terms of the disruption of say the information or communication infrastructure, the energy delivery network, the oil production and so on. Basically the idea is that you want to map from all the 'levers' you have available such as the air strikes, intelligence, ground forces etc. and the strategy. This can be at a very tactical low level or at a policy level. Agent-based modelling allows you to do this and make 'effects based operations' work. It's the only way of actually capturing the network effects that you want to exploit. Because for example, you only need to strike two or three nodes in a communications network and the whole thing goes down. But you have to strike the right nodes and you can't do that if you don't have good mapping between the 'levers' at your disposal and the aggregate level effect on the system. This is a military example but obviously in business you have these situations as well. One way of looking at strategy is as the allocation of scarce resources. How do you map between your allocation of resources and the performance? You might as a manager ask 'What is the impact of putting a coffee machine in a certain place in the organisation?' Will it maximise the spread of innovative ideas?" You need to be able to match the intervention to the performance in some way

Agent-based modelling goes beyond intuition and emergent properties can be quite counterintuitive. One example that sticks with people is the traffic jam. A traffic jam is an entity that is autonomous from the cars that form it. In other words it seems to have a life of its own in going backwards whilst every car is going forward. I was on route to the airport from Boston when they were revising the number of people who died in that awful night club fire. It's awful because 97 people died and I know that the place had a maximum limit of three hundred people and there were about 250 people in it on the night. There were four large emergency exits and it took several minutes for the fire to spread so you would have thought that there was time to evacuate the people. The sad truth is that everybody swarmed to one of the emergency exits and a third of all the bodies were stuck in the emergency exit. Almost all of the people that escaped used the other emergency exis. The outcome was counterintuitive and people afterwards said, 'What can we do? We had big exits and less than the permitted number of people' Some work done two or three years ago on fire escape modelling showed that if you put a pillar about a metre in front of the emergency exit slightly off-centre it optimises the outflow of people. That's a counterintuitive property. It's not the first idea you think of when designing a public space. I mean if the authorities didn't know it you could get sued for doing that. However it is the optimal way of getting people out. It's counterintuitive but to regulate the flow you have to prevent people from yielding to their primal fear of death and herd instinct. By modelling you can capture the emergent property which is the collective dynamics of the crowd.

Before the stampede people lost two minutes because they looked at other people and saw they were calm. They thought it was part of the show. Very dangerous. It's sometimes difficult to sell the result of an agent-based model though the counterintuitive aspect has an appeal. A lot of our clients who I might characterise as complexity geeks with a budget, are waiting for counterintuitive things to happen though they don't always.

But the other advantage of an ABM is that it's very easy to transfer the deep knowledge that experts have about their organisation or their customers, into the model and this kind of participation leads to a feeling of ownership in the client. Because it's a natural description it's easier to calibrate and validate qualitatively. You go to the experts in the organisation and ask "Does this model make sense?" and they will be able to answer because the model reflects how they do things and that's the important point. It's a model of agents activities.   So you can ask: "Is this what you do" or "Is this the way your customers behave?" All business organisations have deep untapped expertise and agent-based modelling enables you to tap into that potential. Moreover the modelling process is incremental. I can start with something simple both in terms of the number of agents and their complexity and I can raise the complexity as I dig deeper into the system.

We can play a game to demonstrate an emergent property. I would like you people here to pick two other people A and B. A is going to be the aggressor and B your protector. So B is protecting you from A. Then I will ask you to move so that B is always between you and A....... That's fine everyone is moving around and around and you successfully keep B between yourself and A. Now we're going to play a second version of the game where you are the protector and you must always protect B from A by standing between them...... You see this simple change in interaction means we all collapse into a ball and if we hadn't been solid physical entities we would have merged.

<u>Slide 7 and 8 - 'Simple rules and Outcome'</u>

This consequence tells us a number of things. First that you can have aggregate level behaviour that has characteristics and properties that de-coupled from the rules that I gave you. Simple rules can generate complexity. Secondly it's very difficult to predict what will happen at the aggregate level by looking at the simple rules. Thirdly if you change the rules a bit you get a very different outcome. In the first case everybody was moving around and moving nothing happened. In the second case we collapsed into a ball. Policy making is difficult because it's hard to predict how a system will behave at the aggregate level by knowing how it behaves at the individual level. When I present this game to the military I also say "Imagine that you have a decentralised organisation that is the enemy. When you interrogate one of the individual in the organisation he can't tell you anything about the goal of the whole even if you can get all the information that he knows. The best he can do is to give you the rules he has been asked to follow". Using agent-based simulation you can predict the dynamics that you have observed in this room and this kind of game, which appears to be very simple is a very rich source of collective patterns which we have been testing with real people.

Questioner 5 : We were playing a game and you gave us a simple rule to which we all agreed, but life isn't like that because we have the person who says "I will not play the game and I will change the rules and play by my own rules".

Eric : Well I asked if everyone was willing to play the game (which is what the CEO of a company would do). Still you are right, people change the rules. Let me tell you a funny story. We played the game at Icosystem when there were only ten of us and a single female who was young, British and attractive. And in the second version of the game everyone wanted to protect her and it created a completely different dynamic which actually prompted us to explore what this game can do. We've played the game when some people wanted to play with different rules, when some people didn't want to play the game, or played with a single defender or aggressor. We've played many versions of it.

Questioner 6 : Do you actually map the pathways that people follow over a period of time. Is it a symmetrical fractal pattern or completely chaotic or what?

Eric : I have no idea about the paths. We have just been looking at the stationary patterns; what happens when the system has been running for some time and asking if it settles into some kind of dynamic pattern or continue to change all the time?

Questioner 7 : Yes I was going to ask that question. What can you milk or extract from the agents themselves as opposed to the collective picture? In the science of networks you can have different kinds of network. For example where the connection between nodes is normally distributed or where you might have a power law distribution of the number of links per node. It strikes me that these agents are a bit like the nodes in that there is a lot of richness in the agents.

Eric : Right. Actually there is enough richness in the agents for the rules to be reconstructed. So now we have to ask the question "Can we reconstruct the rules that individuals are following by observation of what they are doing?" And if we just look at the pattern we ask whether we can we infer what everybody is doing by just looking at the fact that they collapse into a single cluster. Probably we can't because there must be 20,000 different ways of collapsing into the same cluster. Of course if you actually had access to the entire information about the pathways of people you could reconstruct the rules. One of the 'holy grails' of our work is to be able to infer individual behaviour from aggregate data. And even in this simple example, you would have to go very deep into the aggregate level data to infer individual behaviour. It depends on what you know about the strategy space at the innovative level. If you have some idea about what people are doing and can test hypotheses between different options then you have a chance.

Question 7: If you have more than one rule and each agent needed to choose between say two rules it will add to the stochastic content. Could your model do this?

Eric : Yes, it is very easy to do. And I will show what happens when you have a certain stochasticity and other complex parameters. I have two people working extending the game in all kinds of directions which will probably lead to another book.

Questioner 8 : I want to ask a question about the design for individual agent behaviour. If we're talking about a military agent-based model in which say a rifle is worth one ground weight and a tank a hundred ground weights, then you set up your agents with those brownie points and let the model run to see what happens. But the problem with that is that the real world doesn't behave in that way because the situation is always changing. So for example, if you took an aircraft carrier and you were using its brownie points as a basis for your parameter design there would be a great difference between it being on the ocean and stuck down the river. So my question is "How would you model capability when the circumstances can change so much?"

Eric : Well I think your question is about design and this is important because in designing things with an objective function in mind the results are crucially dependant on whether the brownie points apply to the real world and whether the system is going to have to learn and adapt to achieve its mission. For me it comes down to knowing when you can mathematically formulate an objective function that is relevant and makes sense and when you can't. And if you can't you have to resort to alternatives and one of those is the human brain. I have a big project with the Office of Naval Research in the U.S. which is dealing with the design space in which ships can be built. The design space is enormous and the objective function cannot be formulated. But we have people with brains and twenty years experience in ship design. What we would like to do is leverage the expertise without being constrained by the conservatism that goes with it. Experts are often very good but completely biased by their experience. What we do is to use them to evaluate solutions. You want to deconstruct their expertise into building blocks and then reconstruct different options by assembling the blocks in ways that they would never have thought of. But because you have no objective function there is no way you can evaluate the new assemblages and you have to go back to the experts and say "How does that look?" He or she might initially say "This is bullshit" but on looking more carefully admit there is something interesting there.

This is what I'm going to show you later with a game. It's called 'interactive evolution' and it's using the search power of computers with the evaluation power of human beings. Problem solving and decision making has two dimensions : one is search; how do you explore a space of solutions and the other is evaluation; how do you value the potential performance of an option or a solution? People have not in the past thought how it might be possible to de couple the search from the evaluation. and interactive evolution makes that possible. This is out of the scope of this talk but it has for example been used by Pepsi Cola to design new packaging because formulating the objective function for a bottle with all kinds of aesthetic factors is hard so they use people to judge whether it's a nice looking bottle. Honda has been using it to come up with new car designs. So on the one hand you have aesthetic concepts that are hard to formulate and you have ten thousand constraints that need to be satisfied. So you need both the power of the computer search and the human brain.

Questioner 9 : You were talking earlier about mapping from policy space to performance space in a non-intuitive manner. Should the model explain why that non-intuitive result occurs.

Eric : That's a very good point. The explanatory power of a model in general and an agent-based model in particular does not guarantee an explanation because the phenomenon is not explainable in the sense in which we require a 'cause'. We can't identify a 'cause' because in reality it may be two causes or five million causes that contribute a small amount to the behaviour that I am observing. I'm sorry I cannot reduce my explanation to a few simple words. Which is why the kind of intervention that I will design for you is hard to understand in terms of what it does and the problem in business is that managers would rather have a problem they can't solve than a problem they can't understand. In that respect the military is much more open minded and willing to experiment with things, which can be scary in some ways but as far as using models is concerned it's a pleasure to work with them. To be honest I think that in 80% of situations the reasons for what you observe may be obvious and in 19% it may not obvious but you can explain it by doing some analysis of what's going on so that you <u>can</u> reduce it to a simple set of causes and for the remaining 1% it's impossible to explain with a simple set. So it's an exception. But it happened to us with a software company for which we designed a solution for a technical issue using the power of a distributive storage network . It was hard to understand what the system was doing because it was taking advantage of small pieces of available storage here and there in the network to form a coherent storage picture that no engineer with twenty years experience and a centralised storage mindset could understand. So sometimes you cannot understand what you're seeing.

<u>Second session</u>

In this second session I'm going to say some more about 'flows' and something about the logistics of controlling such complex systems. I'll then say something about markets and risk analysis and fraud and terrorist behaviour and finish with the econometrics of agent-based modelling in terms of what has been done and what needs to be done.

Agent-based modelling is ideal for a system where the constituents are people who do things and their behaviour can be characterised at the individual level. So for example, modelling the way that people move in a public space such as a supermarket or theme park or in the case of the fire I mentioned earlier.

Questioner 10 : What kind of rules do you use at the individual level?

Eric : Well this is not my example but it's basically rules that govern the way people move and what their motivations are, in other words their mindset. Are they panicking? Are they trying to follow other people? Are they attempting to move towards the emergency exits? Obviously you have to test the number of hypotheses that you bring to say a fire event, but there is a lot of literature on that. There are several journals dedicated to fires in public spaces. One of your fellow citizens, Giff Steele, is very well known for modelling pedestrian behaviour in public spaces and spent many years studying the Wembley Stadium.

An ABM enables you first to replicate what you observe, but you can then play with a range of interventions to change how the system behaves collectively and test the impact of say changing the configuration of the public space or adding more emergency exits. If you did a simulation of the fire situation you would be able to see that there will be a symmetry breaking event in which people swarm towards one of the exits and forget the others. Some people will use the others but the consequences of a herd instinct will be observed. Now in the Wembley Stadium model putting a pillar in front of the emergency exit enabled 72 people to escape in 45 seconds without injury. The sad thing about the fire that happened near Boston a short while ago was that people fell and others tried to walk on top of them, fell and became obstacles themselves. If you just do an aggregate level model and fix the spatial layout you don't explore changing it. So we have to ask "How did they come up with the idea of a pillar in the first place?" And I'll come back to that.

Questioner 11 : I can see how you can play with counterintuitive physical designs, but isn't the other potentially interesting thing how you play with the assumptions or deliberately change the rules? So for example you could have a rule that says 'people much prefer to leave by the way they came in' which apparently is true when under stress. So as well as the herding instinct it could be that people were worried about going out some other way because they don't know how far they have to go and they might be worried the exit could be locked. Is this what also happens?

Eric : Well I don't think that was done in this case. They looked at the body of knowledge in the literature and on videos of how people react to a fire event and they did interviews. They built a model of the actual dynamics and people did a lot of tinkering to be sure it simulated what people actually did under those circumstances, but yes, you can do that and other people have done it. And it's a very good point in the sense that you may not know exactly people will behave. You have a pool of rules that they might be using but you don't know which ones and you could use the ABM with econometric techniques to determine which existing behaviour patterns were activated during the event.

Josh Epstein and Rob Axtel of the Brooklyn Institution did a piece of work with Ernst and Young six or seven years ago with an ABM to answer the question of the optimum spatial layout for a theme park. And this kind of question can be extended to many other business situations in which there are spaces in which people buy things. Disney had this problem and Disney is a very sophisticated company in terms of customer knowledge. They know your age, how many children you have. They have video data about how you go about the park. They give you tags and $5 coupons so that you can spend $20 in the park. They know exactly what you do during the day by tracking the tags so you wonder why can't they answer the question? The logic seems simple. Suppose they know 10000 people visit the park. They know that 5000 people visited the new Indiana Jones attraction. The park is open ten hours a day and Indiana Jones has a capacity of 600 people per hour. 5000 people over 10 hours is 500 people per hour visiting Indiana Jones so it seems that the capacity is right and there won't be a queue. But there is a queue which has a 2 hour waiting time because when the park opened there were 2000 people and 1200 of those wanted to visit Indiana Jones.

The truth of the matter is that the dynamics of the park unfolds in a way that it is impossible to capture without an ABM. And there are all kinds of reasons for that. Some people will say "I came to visit Indiana Jones and I don't care if I don't see anything else so I'm going to stick to the queue". Some people will be unwilling to wait and will visit other attractions and come back later. Other people will stay in the queue for a while and then get fed up and go and visit other attractions where similar dynamics may happen. The only way of capturing the emergent properties is by modelling the behaviour of each entity in the park and seeing what the collective patterns are. You can then use some kind of optimisation algorithm to search for the optimal layout but you can also use your eyes to see that there is an emergent flow of customers between Indiana Jones and Space Mountain. Now you don't want to people to just go from Indiana Jones to Space Mountain. You want them to spend money on the way so you put a few souvenir booths and a couple of restaurants where they can use their coupons.

Here's a video of the Disney model Axtel and Epstein did.... What you see are blue dots which are customers though there are a few bugs and some of the customers appear to eat the walls!.......... Each blue dot has a profile attached it which is the shopping list, the priorities of the attractions. That provides the elasticity with respect to waiting time and how they are going to respond to various events. They also know how much they are willing to spend an various categories of items. These are things that Disney as a company knows about it's customers. They go through the park from attraction to attraction depending on their shopping list of attractions with priorities. That's how they move according to certain rules. In this simulation you only have a few hundred customers in the park which is why appears Indiana Jones is never fully utilised. However the important point is that an agent-based model provides Disney with a way of leveraging the knowledge that they have about the customers and making it operational to answer the question: "How do I organise my park so that I extract the most money from the customers?" A customer that is waiting in a queue is one that is not spending any money and it is a customer that may not return to Disney.

Slide 12 - Theme Park

Questioner 12 : Is there any way to test the robustness of the various parameters in that program? For example somewhere there may be an implicit preference of the average agent for going from one attraction to another. The question is if you then move those attractions to a different geographical space and the parameters are somehow coded into the rules (change) can you test for robustness?

Eric : Yes and it is a very good question. We're talking about using the model outside of its specific domain of validity because its data has been collected in specific situations and we want to use that data to populate a model and test how it works in different conditions. So the robustness of the model is key. How far can we extrapolate and when does it fall apart?

What we have is a description of customer behaviour in the park that is very natural yet captures the emergent properties. Demographic data can be quantitatively plugged into the model directly along with data from videotapes and tracking tags. Visual displays enable easy interpretation and models are scalable and flexible. In other words the model can be run with 1000 or 2000 or 3000 customers and the level of complexity in terms of agent characteristics can be tuned as more data becomes available.

Los Alamos National Laboratory have done some amazing agent-based models of Dallas, Texas and Portland, Oregon, that have been built over eight or ten years, where 25 million people have been modelled to a pretty low level of granularity. The original Dallas model was to see how air pollution built up in order to put together an infrastructure plan that would reduce it and it shows what people do during the day and what routes they use to go to work. I actually have three different routes that I can use to go to my office. I can switch roads on the way and I usually adapt according to my perception of traffic. Sometimes of course I end up in even more congestion. The model shows the aggregate level properties of the traffic system that result from how people will react to their perceived traffic situation and enables mapping from policy space to performance space in terms of infra-structure planning. In traffic modelling there is the famous paradox where adding a lane to a highway actually increases traffic problems. This is another example of a counterintuitive emergent property which helps us understand the impact of adding another lane to the highway versus a side street. Trying to predict how traffic is going to be impacted by changing infrastructures is impossible if you don't have access to this individual level behaviour and if you can't capture the emergent properties.

Slide 15 - 'Hydrocarbon Emissions in Dallas, Texas'

Questioner 13 : Can I ask a question about the rules behind this? Do you assume that there are types of people who have different rules, so there are some people like you who find themselves in a traffic jam and always turn left and there are some people who will stick with it and so on?. Presumably you end up with stereo types.

Eric : There's a lot of data on that. In a simulation of a supermarket there are things that are well known because people have been studying shopping behaviour in supermarkets for twenty years. Paco Underhill wrote a book called *Why We Buy* which should be called 'How We Buy'. He discovered for example that 73% of all shoppers turn right just after entering the supermarket and 67% of men will make a U turn before the middle of an aisle if they see it's too crowded whereas 89% of females will do so before reaching 20% of the aisle and so on. And of course there's a lot of data about what people buy in supermarkets. But the problem is that the data bases are separate and not necessarily consistent across all fields. What you have to do as in the traffic model is reconcile sources of data to a single model and leverage that to draw conclusions.

We have possibly the simplest model for useful infrastructure planning. It's a model of a roundabout with a number of different types of agency; cars, buses, pedestrians and so on and we want to optimise the light system so that the flow of entities is maximised. What it enables us to do is to test different light synchronisation strategies because when you do something here you get a ripple effect throughout the entire system. Again it seems a simple system but it's impossible to predict what will happen.

Another example of a traffic model is the train station Rennes in Brittany . Again we have a similar number of different types of entity and what we want to do is optimise the flow while satisfying a number of safety constraints. It's a complicated system and you want the solution to be robust to variations in traffic and to work at any time of day and maybe adapt to time of day using a light system that responds to its perception of traffic and so on. Flows are probably the area in which the most significant results have been achieved by agent-based models because we're dealing with physical systems even though they involve human behaviour. A human being on a highway is highly constrained. OK I can do all kinds of things and my driving behaviour can be eccentric, but behaviour in general is still highly constrained. And so these models can make quantitative predictions much better than if you start adding soft factors into the picture such as the psychological characteristics of human beings

Now here's an example of the kind of problem that is common in business organisations. Area branch managers of a company need to satisfy customer demands at marketing units whilst minimising stock-keeping costs in a highly seasonal and weather sensitive business.

Slide 27

Current incentives for each branch were based on short storage of product because cost is high, but this resulted in product shortage at the retail outlet. What happened as a result was that a branch manager would order twice as much as his average forecasts were telling him for the year because he was scared of running short. The cause of the problem was that the demand for the product was very weather sensitive so the month to month forecasts didn't mean much. If, for example, the weather was bad in Italy huge shipments would have to be made to Denmark where the weather was gorgeous. and transport costs were high. The task was how to encourage area managers to buy the right amounts of stock so that globally the system was optimal? So, we have a trade off between how much it costs to store the product and customer service levels and we have to understand what drives area managers and how that maps onto their purchasing practices. The solution has to be simple because the incentive has to make sense to the managers and you want to make them responsible for and sensible to storage and shipping costs. But the mapping between area manager behaviour and overall behaviour is not obvious and we had some surprises there because ordering more than you need is not necessarily a bad thing if you do it with moderation because it buffers the business. If the weather is great and production capacity is saturated then you're going to have shorts. So there's a number of trade-offs and you don't want to drive the organisation with specific directives.

Another example of designing a human organisation with a particular goal in mind was ensuring that innovation was treated properly at Du Pont Capital Ventures. One of the problems there was that a lot of good projects would not be selected and would fail too early. A corollary or connected situation was that many of the backed projects would never fail. I don't know whether you've seen the current Harvard Business Review, but one of the articles is called 'Why backed projects are so hard to kill' and it happens in any human organisation. We had to use a very soft agent-based model , because it's difficult to manage your client's expectations. A soft model brings insight but not quantitative predictions and one of the things that we tested was the incentive structure for the different stages of the' innovation funnel'. The innovation funnel' is where you have a large number of ideas that get a small amount of funding for the first stage of testing, and ideas which make it to the second stage by satisfying certain requirements get more money and so on, until the product can be marketed. The process involves a transition from 'truth seeking incentives or behaviour' to 'success seeking behaviour'. What this means is that in the early phases you don't want people who will cling to the project no matter what; because for example, they think their career is going to be destroyed if they fail. If the project is going to be a 'dog' you need to kill it quickly. The current mentality at Du Pont and almost the entire pharmaceutical industry is often: "We've got to make this thing work" even at the early stages. That's why when everything has failed people either go for male erectile dysfunction or anti smoking drugs. Anyway we had to build a model that would tell us how to build an incentive structure that favoured full exploration of concept space so that good ideas had a fair chance of getting to market.

<u>Slide 18</u>

Questioner 14 : I assume that in order to do this you need quite a lot of understanding of behaviour. Do you work with psychologists for example?

Eric : No, we use tools from behavioural economics studies but we don't work with psychologists.

Questioner 15 : As a psychologist can I just say that one of the ironies that's come over to me about your opening remarks about being happy with models and not wanting to get involved with people is that hese models are absolutely full of psychological and sociological data and assumptions.

Eric : Yes, my discomfort comes from this Du Pont example. I'm very uncomfortable doing this kind of work for a client because I have to manage my client's expectations very tightly. I don't want them to think that I'm going to deliver any kind of predictive power here. The more 'soft' factors or psychology I add to my models the more uncomfortable I become. Not because I don't like psychologists but as a scientist I think we need to make more progress there. And I actually think that a meeting between agent-based modelling, behavioural economics and cognitive psychology would help a great deal. Yes my models are full of psychology but I'm scared when I sell them to my clients. What I'm saying is that the more soft factors you add to the models the more cautious you have to be with respect to what you can get out of it.

An interesting problem in logistics which Bios handled, was offered by South West Airlines. They needed to optimise their freighting activities but they didn't have a fully connected network of airlines. If, for example, you sent a package from Seattle to Miami it had to go though intermediate airports. To deal with this issue Bios had to build an agent-based model in which the agents were not human beings though human beings came into the operation. But basically they were looking at planes and packages and attempting to find rules which would optimise the system. The symptom of the problem was that overall cargo capacity was only 25 % taken up yet saturated in particular places. What they were looking for was a routing system that would enable them to re-distribute the load so that they could increase their activity level. Freighting is potentially a very profitable business but issues such as the weather and strikes can play an important part. The model threw up quite a lot of emergent properties and handlers had to be given very simple rules that they could remember and apply without thinking. Bios came up with simple rules that led to 71% improvement on weight transfer by delaying package transference from one flight to another. It might take slightly longer for a particular package to reach its destination, but it completely removed the saturation points from the system and enabled South West Airlines to save $2 million dollars a year in labour costs and increase their business by 30 or 40%.

OK, one of the interesting things that the (Barras?) group did was to model the stock market (Nasdaq). The market index follows a random but Poisson type distribution curve. Investors receive 'noisy' information about this value and decide to trade by comparing it with the available price of stock, seeing what the market trends are and using other technical devices. Market makers receive buy and sell orders and have to learn how to set their quotes profitably. The upshot is that the rules of the Nasdaq stock market impact on market characteristics such as spread etc. The simulation had to take into account a large number of factors (slide 27). When large institutional investors found doing 'limit orders' (large deals) affected the 'spread' (buying to selling price) and therefore the extra margin they had to pay they switched to watching the market and doing smaller deals which led to more volatility. Increase in spread was actually explained by looking at how agents reacted to it in the model.

<u>Slides 27 and 28</u>

I do want to talk about risk in financial institutions. By that I mean the risk of losses due to operational issues such as fraud or malfunctioning of the information system or human error. At present the Basle Committee, the regulatory body is putting a lot of pressure on banks to come up with a way to measure operational risk and to put aside money to cover losses. The problem with that is that it's a lousy investment, since the money is not working and if I'm managing $500 billion and I have to put aside 3% that's a lot of fairly idle capital. I need a way to quantify what the operational risk really is and the problem is that there is no data to tell me at what point I'll lose so much money that I'm going to go bankrupt, like Barings. There's not much data because banks don't go bankrupt every day and they usually go bankrupt for different reasons so the past is not a predictor of the future. The fact that I've had a number of small losses because of everyday mistakes doesn't mean I won't have a big mistake tomorrow which ripples through the organisation and takes it to the verge of bankruptcy. What do I do if I have no data to build a statistical model?

Well, I have a lot of expertise about how the organisation works, about what people do, how money flows and I can build a model from the bottom up that will capture big losses as emergent properties. And I may be able to quantify to some extent how risky my organisation is. What we found in the actual model we built was that big losses arose from senior management practices. For example, an employee came to work for the bank in August when practically all the fund managers were on vacation and there was one fund manager who had to deal with 10 funds instead of 1 or 2. So he was overworked and then something like the Russian crisis of 4 or 5 years ago occurred where the activity of the market increased dramatically and the employee starts getting orders from his clients who want to buy and sell and it's Friday afternoon and there's a guy who says "I want you to use my roubles to buy 500 million dollars". What happens is that the employee takes the order, makes a couple of mistakes by writing down billion instead of million and that he has to buy roubles instead of selling them. The order goes into the system and during the weekend the entire Russian economy melts down, the rouble loses 99.9% and bingo you have a 500 billion dollar loss. So it's not a single event that leads to disaster but a dysfunction in the system amplified by external shocks. What a model shows is that although in normal day to day business there are ups and down, some greater than others, the variance of the loss distribution increases dramatically when external conditions change quickly and you need careful monitoring of the internal and external environment. So this is an important factor to take into consideration, just as a lot of volatility in the stock market dramatically increases risk to investors.

But the important conclusion to reach is that risk is carried by the agents that do things in organisations. Risk is not something that can be picked up by process analysis, an ABM is the only way of dealing with it. So here's the epiphany that I present to accountants who try to map risk onto process by saying well this bit has a certain risk and this bit and then they add the risk up. This is useless. OK you need the process picture but what you really need is a picture of the activities in terms of what people do as agents. It's a different picture from the usual process picture. And remember you can go back to each agent and say does this model make sense? In my model I said you do this and this in this situation and so on. And the person might say "Well that's not exactly what I do" and you say "OK, we'll change it . So you have feedback. If you have a gigantic set of 2000 equations to describe the aggregate level properties nobody will want to talk to you. This is my epiphany.

<u>Slide 35</u>

Now my last application is using agent-based modelling to find a needle in a haystack i.e. the detection of abnormal behaviour. You have a huge volume of normal transactions in the system and you're looking for the abnormal transaction which may be indicative of fraud or sabotage. And you can't use a statistical model because you have too few examples of abnormal behaviour and that behaviour in the past may not be relevant to the future. ABM as behaviour from the bottom up enables you to create a haystack full of needles. You model the normal behaviour and you think about (make assumptions about) what the bad guy may want to do with the system. And you run it many times and then you have statistics about what could be the signature of abnormal behaviour. We've done that for the US Army CCIU

(Computer Crime Investigation Unit) and we came up with a model of hacker behaviour to see what kind of evidence would be left behind in the log files and errors that they would have made during the intrusion which would help the forensic investigators and the incident handlers all over the world. So if you're dealing with an intrusion here's what you should look at in the log files and then the investigator can say "This is what I find and I can tell you what is the most likely motive for intrusion" So you're creating a haystack full of needles to then generate rival statistics by running billions of simulations to create a very reliable statistical signature and then you can design the detectors that you want.

<u>Slides 40,41 and 42</u>

I'm going to skip the econometrics which is really about the need to develop techniques to calibrate and validate agent-based models but I just want you to go away with the knowledge that a lot of people have been tinkering with the technique to find a frame work that goes beyond and deals seriously with the calibration issue. We have a methodology and my slides 44 to 47 illustrate some of the principles involved but I do not have time to go into it here.

In 2003 we have about 20 times the number of applications examples that I have presented today that have been real successes but calibration issues have to be taken seriously in evaluation. I can reproduce the finance models for example, but I need to have a rigid framework for value estimating my models in terms of objectives. We should all be scared of GIGOT (garbage in garbage out) and be aware that agent-based models do not solve all problems effectively and it can be too expensive even if it is the best approach.

I want to ask you a question which goes back to the game in order to demonstrate what it can do in creating interesting patterns. I'd like to ask you where you would take the game from here in terms of human organisation and what results you would like to see. For example you talked about injecting stochasticity into it to see what happens. I'd like to hear suggestions relevant to situations that people are facing.

Questioner 16 : I'd like to see more examples of models where people operate with different rules rather than all the same and also where people change them.. Because sometimes people are very good organisational citizens and communicate and sometimes they deliberately withhold information. So people change the rules during the process.

Eric : Why, to get nearer to reality?

Questioner 16 : Yes, I'd be interested to see what would happen if you move from a simple linear axis. So for example, you might say you want the protector to be between yourself and two other agents instead of one. Because at the moment everybody is in a straight line.

Eric : OK, what you see on the screen here is what we've played with simple extensions in some of the directions you're suggesting but simpler. Playing with two roles. So for example 30% of the audience would use rule A and 70% would use rule B (the rules I gave you earlier). With the right proportions and the right interconnections between people . Here's an amazing pattern you get. This is completely unexpected

Questioner 17: Do you tell them who will be interacting with who?

Eric : Yes to get this pattern you have to have a specific topology of the interaction network. Which means I would have tell everyone here the names of two people that are going to be their aggressors and defenders. So it's a specific configuration but it seems to be very robust for that combination of aggressors and defenders. In the design space, i.e. all possible ways of doing the relationship between people for about 15% of the time I get this pattern which we call the 'Chinese streamer'. In other words if I give you the exact network that you should satisfy and a few people make mistakes it doesn't affect it. So in that respect it's stable and robust.

This one (another one),is not as clean as the previous one we played, because I say OK you can pick anyone you want. However I'm still using the two simple rules and asking you to have a defender and an aggressor. Red is using the second rule and blue is using the first rule. And of course in the simulation you can adjust the speed. We started playing this game as a toy and we said "Oliver you're going to be the aggressor for everyone" and we ended up creating an amazing pattern that was a line that was symmetric on Oliver and circled around Oliver like a propeller plane rotating around its axis. But then it would be like a Michael Crichton novel because some of the constituent units of the propeller would move between the two sides. But you had this rotating pattern which was a richness we didn't expect. This is actually pretty stable in the spatial. We get several patterns but after a while it very likely to form the Chinese streamer. I don't know how many of you have read Michael Crichton's novel called *Prey* but he's basically talking about swarms of inanimate entities gone wild. And the swarms can aggregate and disaggregate and form collective patterns that are reminiscent of a human body so my goal with this project is to produce my face.

Questioner 18 : I'd be interested in experimenting with the information flow, where you can see where the two people are. In other words you see the need and you respond to it. Suppose people either saw each of those two people as a delay or there was a delay throughout the whole population That might be interesting.

Eric : I agree completely with you.

Questioner 19: I'd like to know why you choose these particular rules about protector and aggressor.

Eric : Because there was an improvisational theatre troupe in Boston who ran a seminar for a consulting firm two years ago and they asked us to play the game. I asked myself "how I could go further by actually building a model to predict what happens?" I'm discovering the richness of this model every day. I don't even know all the new rules. people are trying. Someone sent these examples yesterday by e-mail.

Questioner 20 : Do you ever get stationary states where there's no movement?

Eric : We do.
Here's an example where we have 52 people and you can see it moving collectively as a straight line but none of the participants think they are collectively trying to form a straight line. They're only following the defender and aggressor rules.

Questioner 21 : Could you elaborate a little on what the relationship between business process and agent attributes or activities. How do you relate one to the other in this approach?

Eric : What as I said, process description is OK and I was involved in the bank project with a company called Ernst and Young whose expertise was process or business modelling. That's OK for the people actually interested in the management of the fund or the back office, but what a person actually does can span five processes and one process can have 25 activities in it at least. So process is a very abstract description of a clients organisation. You have to talk to people in the organisation about what they actually do.

Questioner 21 : But if you're implementing a process one of the things you would try to do for each actor is to present him or her with a role which would be the individual rules that they are going to follow?

Eric : I agree and you have to be able to reconcile the two perspectives. But I do believe especially for risk modelling it's very dangerous to be process focused.

Questioner 21 : But I got the impression that you had a clinical development model and I thought you were going to tell us that you defined some sort of process and then gave some sort of random behaviour to the agents going through it.

Questioner 22 : What you seem to be talking about is, when a process is implemented, people start off with an overall view and then develop their own behaviours and forget how it all hangs together.

Eric : I think the process description is a very valuable scaffolding or framework in which your agents perform their activities, but at some point, especially in risk analysis, I don't need that scaffolding anymore. But it is very useful to see how consistent an agent's activities are with the process.