Complexity Economics
Application and Relevance to Actuarial Work

A Report from the Agent-Based Modelling Working Party

Jon Palin
Nick Silver
Andrew Slater
Andrew D Smith

Introduction

The purpose of this paper is to introduce the concepts of agent based modelling to the actuarial community. The paper will provide a brief introduction to the subject, will look at examples of agent based models, and will identify where they might be used in future and what the implications might be for actuarial practice and theory.

Agent based modelling (ABM) was motivated by the observation that there appears to be a commonality between different systems in the world – including ecosystems, financial markets and the weather. Their development was facilitated by the increased availability of computing power and the development of computational models (Miller and Page, 2007).

ABM was not developed specifically for finance, but there is a growing interest in its use in the field. We believe that the insights from ABM could have profound consequences for the risk management of financial institutions, and could also be fruitfully utilised by actuaries. However, in saying that ABM is currently in an early stage of development – whilst current models produce outputs which look like the real markets – they are a long way from producing reliable calibrated models with predictive power.

Disclaimer: We have drafted this paper to be informative and provocative. There are many views on agent based models, and the working party did not always reach unanimous views. Therefore, the reader should not assume that all the authors fully endorse every statement in this paper.

1. What is agent-based Modelling ?

1.1. The World as an Adaptive System

Proponents of agent based modelling views the economy (and indeed the world) as a complex adaptive system, the structure of the markets, the interplay between agents and time lags are the cause of much of the complexity and interesting behaviour we see in the real world. Much of science (including economics) is reductionist, an attempt to reduce the world to its basic elements. However, the interaction of these elements causes behaviour – often described as emergence - which cannot be predicted by studying the elements themselves – think of ant colonies, the brain and economies (Miller and Page (2007)).

The agents of a complex system, often by following simple rules, form a system that behaves qualitatively differently from the individual agents themselves. A good example of this sort of system is an ant colony – the colony as a whole can perform complex tasks – such as defending the colony from an aggressor by way of the individual ants processing simple pieces of information.

1.2. Typical Attributes of Agent-based Models

ABMs attempts to capture this and typically have the following features:

1. Heterogeneous agents: there are a finite number of heterogeneous agents. These agents are allowed to follow rules; these rules can be very simple (for example if stock goes up buy) to highly complicated rules. Heterogeneity arises as groups of agents are allowed different rules. This contrasts with neo-classical economic models which tend to assume there is only 1 or an infinite number of homogeneously rational agents.
2. Adaptation: agents are often allowed to adapt; i.e. individual agents can learn during the running of the model, or are allowed to combine with other agents ("sexual reproduction") or evolve as the less successful die off and the more successful agents flourish. Again, this contrasts with neo-classical economic models where agents are the same through time.
3. Feedback loops: agent behaviour can have positive or negative feedback – for example if agents buy when a stock goes up this will cause a positive feedback loop which would result in an asset bubble. Again, this contrasts with neo-classical economic models which assume that systems return to equilibrium.
4. Local interactions: neo-classical economic models assume that all agents can trade with each other. ABMS can allow for markets to have structure, for example traders can only deal with brokers, there are time lags and possible geographic affects (although the latter is not so important in financial markets).
5. Externalities: agents interact with their environment, which is in turn affected by agent's behaviour.

1.3. Parsimony and Realism

As can be imagined from the above discussion, the skill of ABM involves a playoff between parsimony and realism. On the one extreme, many neo-classical financial models reduce the workings of a market to a few easily solved equations. Whilst these equations or soluble and rigorous and useful in many circumstances, they lose much of the interesting features of real-world markets. On the other extreme, a 1-1 scale map of the world is almost completely useless (although for some circumstances, e.g. predicting weather it would be useful, provided you had a big enough computer).

The early results from ABM show that a complex emergent behaviour can result from relatively simple rules amongst agents, increasing the complexity of the agents often does not result in altering the system's behaviour. Many of the interesting "life-like" results from ABMs, for example asset bubbles and crashes, income distributions, are generated endogenously from the models due to the interaction of the 5 features described above – these are often assumed away of treated as exogenous by ne-classical economic models.

2. Agent-Based Modelling – A Brief History

2.1. *Origins*

The idea of agent-based modelling was developed as a relatively simple concept in the late 1940s. Since it requires computation-intensive procedures, it did not become widespread until the 1990s. The history of the agent-based model can be traced back to the Von Neumann machine, a theoretical machine capable of reproduction. The device von Neumann proposed would follow precisely detailed instructions to fashion a copy of itself. The concept was then improved by von Neumann's friend Stanisław Ulam, also a mathematician; Ulam suggested that the machine be built on paper, as a collection of cells on a grid. The idea intrigued von Neumann, who drew it up—creating the first of the devices later termed cellular automata.

Another improvement was introduced by the mathematician John Conway. He constructed the well-known Game of Life. Unlike von Neumann's machine, Conway's Game of Life operated by tremendously simple rules in a virtual world in the form of a 2-dimensional checkerboard. However, it is still not a game in a traditional sense – there is only one player, and no concept of strategy. The game simply evolves.

The creation of agent-based models of social systems is often credited to the computer scientist Craig Reynolds. He tried to model the reality of lively biological agents, known as artificial life, a term coined by Christopher Langton.

2.2. *Computational and Mathematical Organisation Theory*

At the same time, during the 1980s, social scientists, mathematicians, operations researchers, and a scattering of people from other disciplines developed Computational and Mathematical Organization Theory (CMOT). This field grew as a special interest group of The Institute of Management Sciences (TIMS) and its sister society, the Operations Research Society of America (ORSA). Through the mid-1990s, the field focused on such issues as designing effective teams, understanding the communication required for organizational effectiveness, and the behaviour of social networks. With the appearance of SWARM in the mid-1990s and RePast in 2000, as well as some custom-designed code, CMOT -- later renamed Computational Analysis of Social and Organizational Systems (CASOS) -- incorporated more and more agent-based modelling. Samuelson (2000) is a good brief overview of the early history, and Samuelson (2005) and Samuelson and Macal (2006) trace the more recent developments. Bonabeau (2002) is a good survey of the potential of agent-based modelling as of the time that its modelling software became widely available.

Joshua M. Epstein and Robert Axtell developed the first large-scale ABM, the Sugarscape, to simulate and explore the role of social phenomenon such as seasonal migrations, pollution, sexual reproduction, combat, and transmission of disease and even culture.

Agent-based models consist of dynamically interacting rule based agents. The systems within which they interact can create real world-like complexity. These agents are:

- Intelligent and purposeful, but not so intelligent as to reach the cognitive closure implied by game theory.
- Situated in space and time. They reside in networks and in lattice-like neighbourhoods. The location of the agents and their responsive and purposeful behaviour are encoded in algorithmic form in computer programs. The modelling process is best described as inductive. The modeller makes those assumptions thought most relevant to the situation at hand and then watches phenomena emerge from the agents' interactions. Sometimes that result is an equilibrium. Sometimes it is an emergent pattern. Sometimes, however, it is an unintelligible mangle.

2.3. Agent Based Models as an Extension of Traditional Methods

In some ways, agent-based models complement traditional analytic methods. Where analytic methods enable humans to characterize the equilibria of a system, agent-based models allow the possibility of generating those equilibria. This generative contribution may be the most mainstream of the potential benefits of agent-based modelling. Agent-based models can explain the emergence of higher order patterns -- network structures of terrorist organizations and the Internet, power law distributions in the sizes of traffic jams, wars, and stock market crashes, and social segregation that persists despite populations of tolerant people. Agent-based models also can be used to identify lever points, defined as moments in time in which interventions have extreme consequences, and to distinguish among types of path dependency.

Rather than focusing on stable states, the models consider a system's robustness -- the ways that complex systems adapt to internal and external pressures so as to maintain their functionalities. The task of harnessing that complexity requires consideration of the agents themselves -- their diversity, connectedness, and level of interactions.

2.4. Modern Applications

Agent-based models have been used since the mid-1990s to solve a variety of business and technology problems. Examples of applications include supply chain optimization and logistics, modelling of consumer behaviour, including word of mouth, social network effects, distributed computing, workforce management, and portfolio management. They have also been used to analyze traffic congestion. In these and other applications, the system of interest is simulated by capturing the behaviour of individual agents and their interconnections. Agent-based modelling tools can be used to test how changes in individual behaviours will affect the system's emerging overall behaviour.

2.5. Online Resources

A simple and accessible program for creating agent-based models is NetLogo. NetLogo was originally designed for educational purposes but now numbers many thousands of research users as well. Many colleges have used this as a tool to teach their students about agent-based modelling. A similar program, StarLogo, has also been released with similar functionality. Swarm was one of the first general purpose ABM systems. Swarm, developed by the Swarm Development Group, uses the Objective C programming language, and is recommended for C programmers with little object-oriented programming experience. Swarm can also be implemented by Java programmers, as can Ascape. Both MASON and Repast are widely used, and EcoLab is suitable for C++ programmers. Cormas is another platform, focusing on natural resources management, rural development or ecology research, based on the SmallTalk language. All the toolkits described previously are based on serial von-Neumann computer architectures. This limits the speed and scalability of these systems. A recent development is the use of data-parallel algorithms on Graphics Processing Units (GPUs) for ABM simulation. The extreme memory bandwidth combined with the sheer number crunching power of multi-processor GPUs has enabled simulation of millions of agents at tens of frames per second.

3. Simple Environment Models: Cellular Automata

Most agent-based models can be thought of in two parts: an environment and a collection of agents. The environment sets ground rules but evolves only passively. Agents, that live on the environment

Cellular automata is a field of academic study in its own right, and forms the environment half of ABM (the other half being the agents that “live on the environment). While the environments used for illustration in this paper have a direct physical interpretation they can have a much more abstract construct. For example in a financial model the stocks that are traded are part of the environment.

What cellular automata demonstrate is how complexity emerges from very simple rules. Consider a one dimensional line of cells, and each cell can be coloured black or white. The colour of each cell in the next time period is determined by the colour of it and its immediate two neighbours (one on each side) in the current time period according to a rule that does not change across time. Between one time period and the next the new colour of each cell is determined independently from each other according to the rule then all cells change to their new colour.

3.1. The 256 Possible Evolutionary Rules

There are eight different combination of how three cells can be coloured either black or white. Each combination can specify whether the cell should be black or white at the next time period. So there are 256 distinct rules that are possible. Examples of six rules are given below, referred (0 to 255) using an obvious binary notation. For each rule the eight combinations are shown and below each the new cell colour that results.

Rule 0

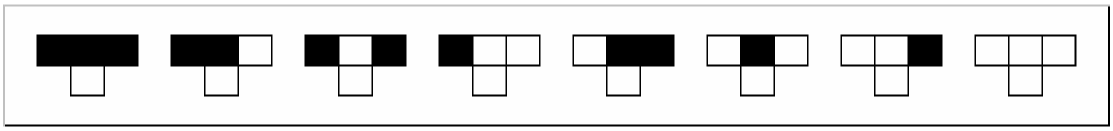

Rule 30

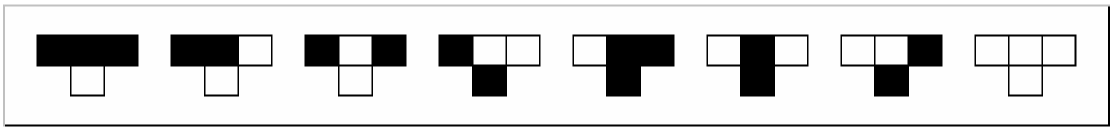

Rule 90

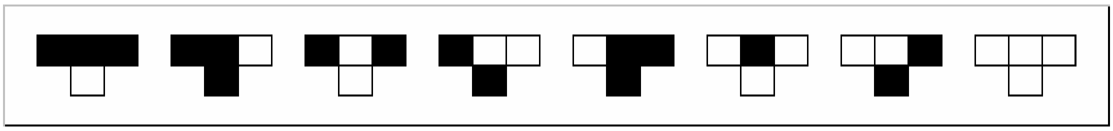

Rule 110

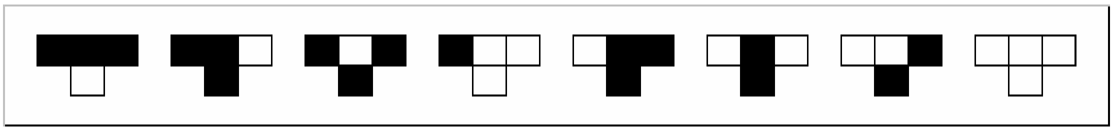

Rule 250

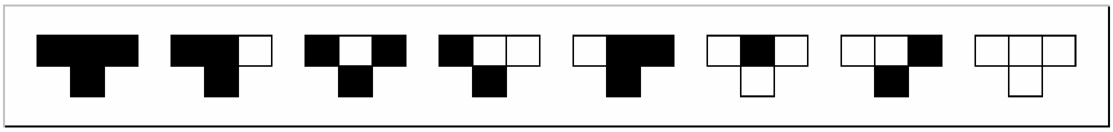

Rule 255

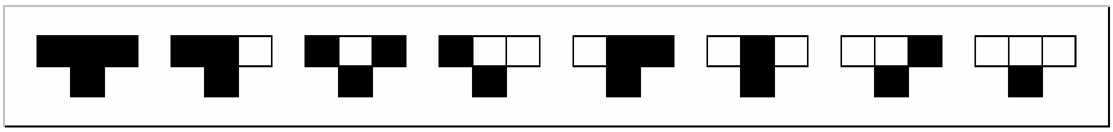

The time evolution, regardless of initial condition, of Rule 0 and Rule 255 should be clear: all cells become either white or black. To investigate the other four rules we consider an initial condition of one black cell with all other cells white. In the diagrams below each line as you move from top to bottom represents the next time period.

3.2. *Example Rules*

Rule 250 gives the simple behaviour of repetition.

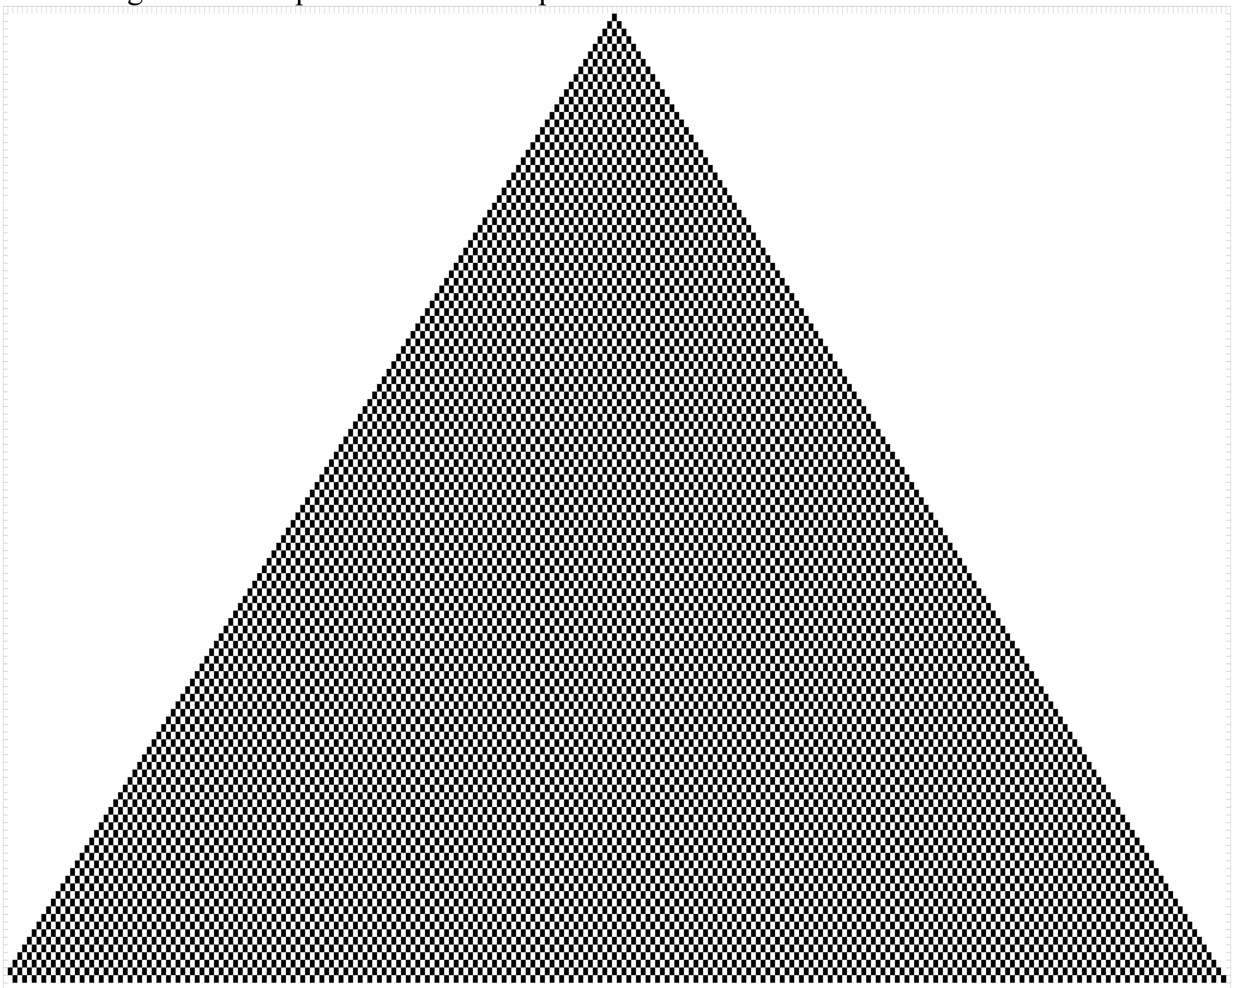

Rule 90 gives a more complex "nesting" behaviour. While the pattern may be reminiscent of a fractal remember that the diagram has been built up by applying the same rule repeatedly line by line from top to bottom. It is not the result of a global rule or equation, which is for example how the Mandelbrot set is constructed. And the structure of Rule 90 is identical to Rule 250 (or 0 or 255) with just the detail changed. The first indication that simple rules can create something interesting.

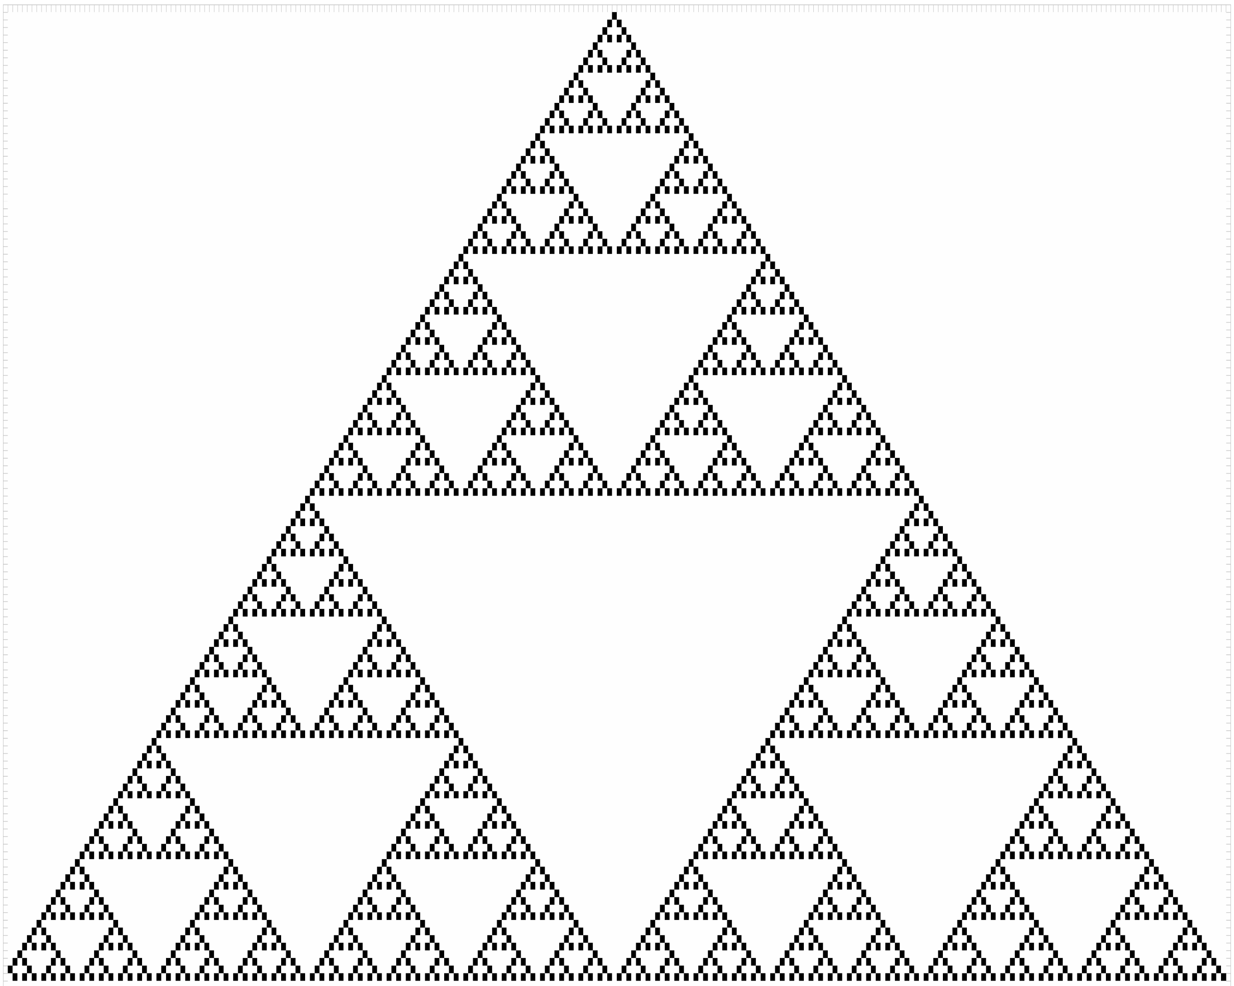

Rule 30 illustrates "from simple becomes complex" even further. The behaviour that results is random, showing no regularity, even though its construction is deterministic.

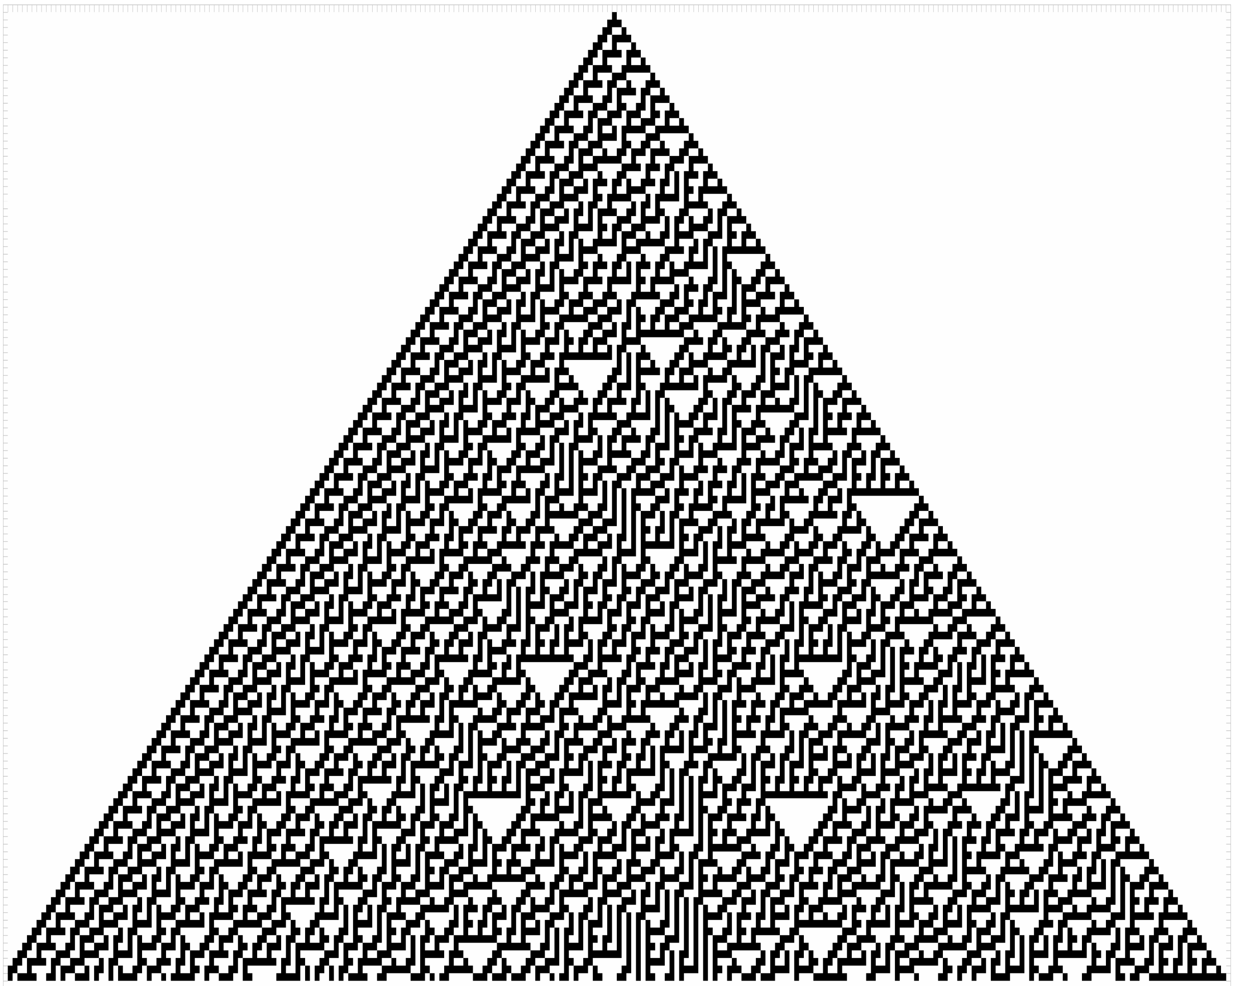

As the last we look at Rule 100, which exhibits localised structures. Periods of regularity punctuated by non-regularity; perhaps the basis for modelling the behaviour of a financial market.

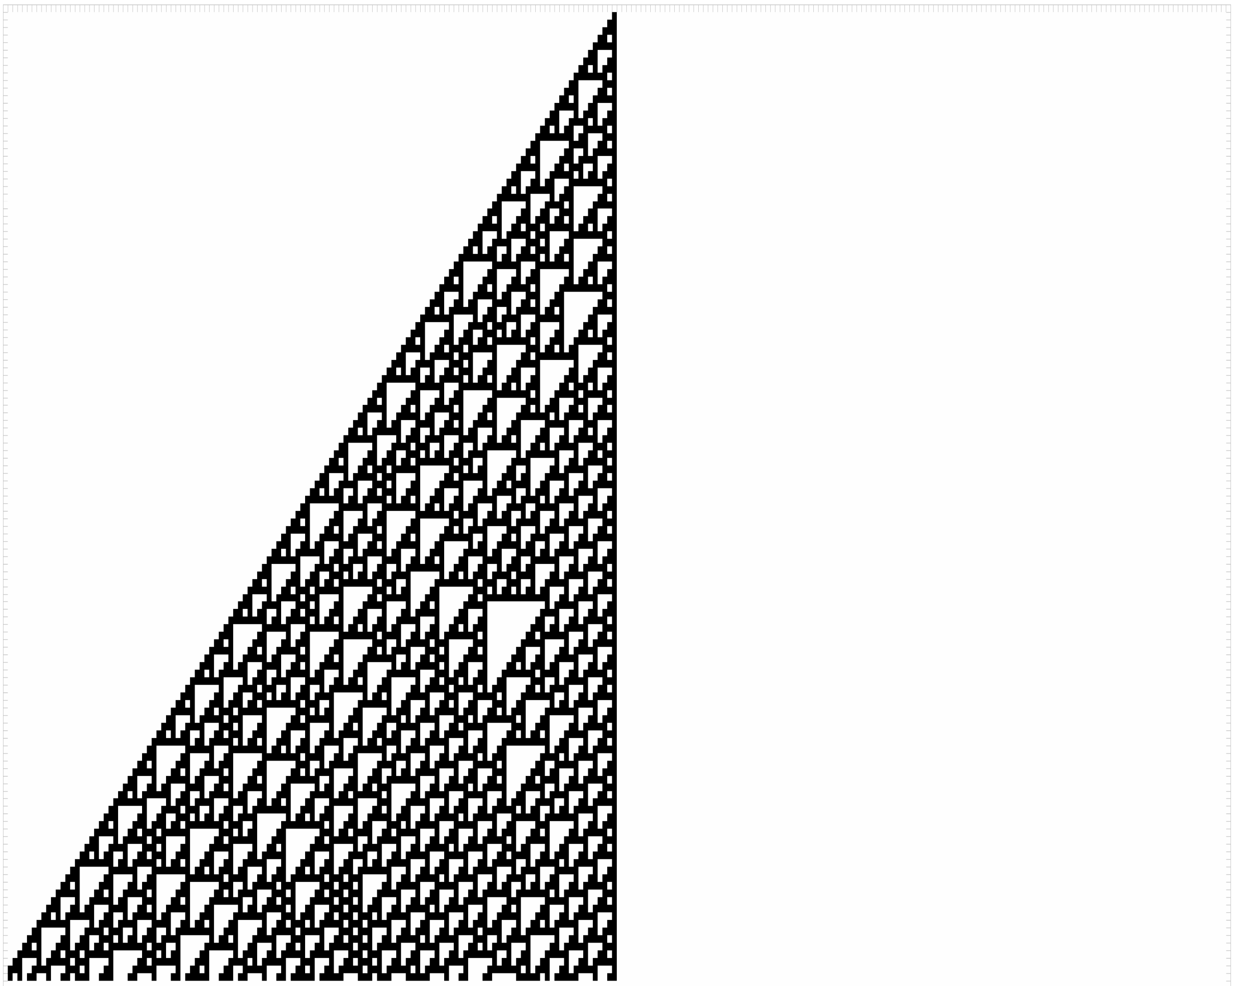

4. An Interactive Economic Model - SugarScape

Sugarscape is described in the book “Growing Artificial Societies” by Joshua Epstein and Robert Axtell, and in the book they deploy the model to investigate a wide range of social structures. Here we will look at just two. The first is growing a Pareto distribution of wealth which fits that which is observed. The second looks at trading between agents and the result that the actual price in a trade is often far from the equilibrium price determined from supply and demand curves of textbook economics.

Sugarscape consists of a 50 cell by 50 cell grid (actually a torus, if an agent goes off the left it reappears on the right, for example). For each point on the grid there is a maximum capacity of sugar that can be supported, but at a point in time the capacity may be less than the maximum. The initial distribution of sugar is such that the are two “mountains” of sugar. The landscape is subsequently extended to include a second commodity – spice – and there are two separate spice mountains. There is a rule on how sugar (or spice) grows back each time period. In this respect the landscape is a cellular automata. What makes it an ABM is the introduction of agents: *cellular automata + agents = Sugarscape*. So there are three types of rule:

- Environment – Environment (eg sugar grows back after cutting)
- Environment – Agent (eg harvest sugar, move)

- Agent – Agent (eg metabolise sugar, metabolise spice, trade)

In the simulations we will look at all agents are identical, though do show variation in their "genetic" makeup. (Other simulations, for example, differentiate between male and female agents and there is a rule for reproduction, and models can include different "species" of agent.) Each agent has a "vision" which determines how many grid cells (1 to 6) ahead (north, south, east and west, no diagonal) it can see the amount of sugar or spice. Each agent has a metabolism for sugar and spice, that is how many units of sugar or spice it consumes each time period (1 to 4). Each agent has a maximum age (which could be infinite) before it dies, if it has not already died from starvation (when its supply of either sugar or spice reaches zero).

4.1. Initial State

These genetic characteristics are randomly assigned for each agent initially (and similarly when a new agent is introduced when an existing one dies), together with a random position on sugarscape (no two agents can occupy the same position) and a random initial supply of sugar and spice between 5 and 25 units (there is no limit to the size of each agent's "wealth" of sugar/spice it can accumulate). Once a simulation has been set up there is no further random element, except when a choice needs to be made between two otherwise equal selections. So though not completely deterministic a simulation is definitely more deterministic than stochastic post an initial random set up.

A simulation consists of the initialisation and letting the environment and agents follow the rules that govern them. Basically the sugarscape just wants to grow back to maximum capacity. And basically the agents just want to move to where they can get the best amount of sugar and spice, harvest all of it, trade with a neighbour if that can further improve their situation, and metabolise. In each time period the agents take their turn to move etc in an order that is random each time period (agents go one at a time whereas cellular automata go all at once).

4.2. Model Evolution

With sugar and spice (and generally with two or more commodities) answering what is the "best" allocation of sugar and spice requires a quantitative specification of welfare to effect the move and/or trade rules. However, the qualitative description given in the text still holds and describes conceptually how Sugarscape operates. Readers are referred to "Growing Artificial Societies" for the detail which Epstein and Axtell implemented for their model, which they considered the simplest that could be constructed[1]. The detail is just that, a detail to make the model operate in practice and is not something that its design rests upon. Nor are the results presented sensitive to the formulation of the detail.

---

[1] Much of the book can be viewed at
http://books.google.co.uk/books?id=xXvelSs2caQC&printsec=frontcover&dq=animation+IV-2&source=gbs_summary_r&cad=0

The diagram below shows the initial distribution of sugar with each cell at its maximum capacity.

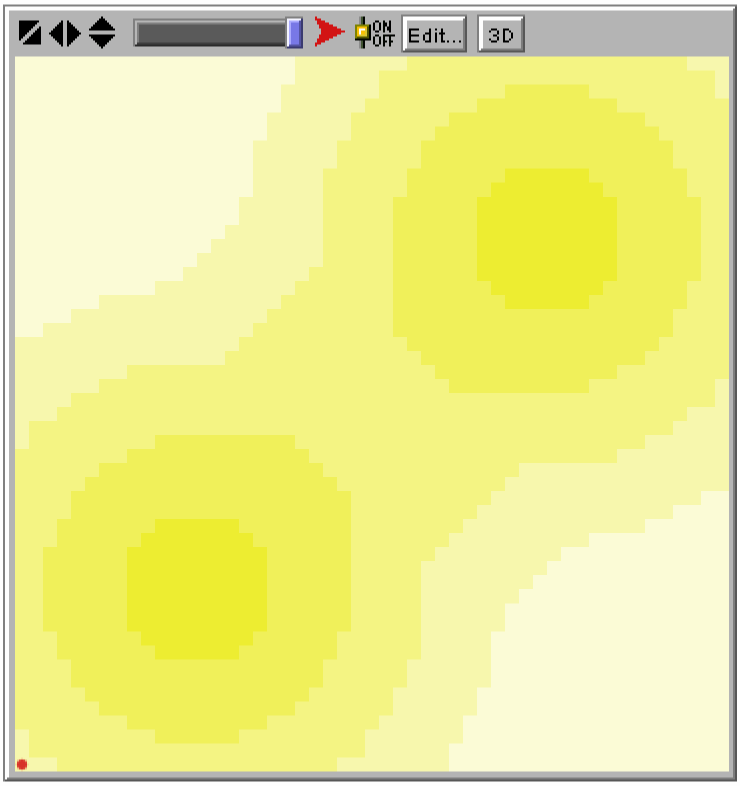

The shades of yellow represent 0, 1, 2, 3 and 4 units of sugar. (The red dot in the bottom left corner is a rather lonely agent.)

When spice is introduced the initial set up becomes as show in the next diagram, with sugar in yellow and spice in orange, the red dots being agents.

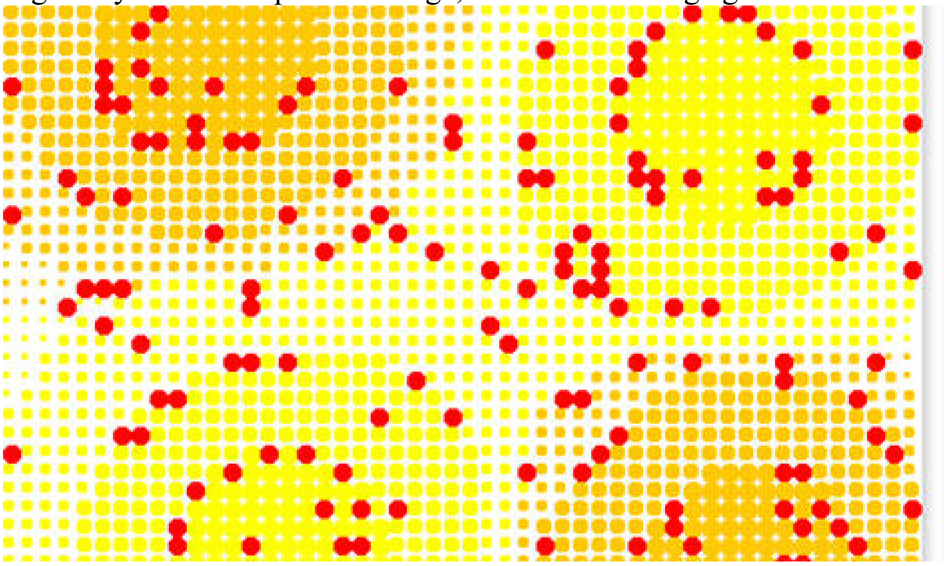

The agents obey the following rules. The rules are presented in generic form, a particular simulation will specify the necessary parameters. And the rules are presented in English rather than computer code, which will be dependent on the software being used.

Sugarscape grow back rule $\mathbf{G}_{\alpha}$

- At each lattice position, sugar grows back at a rate of α units per time interval up to the capacity at that position

Agent movement rule **M**

- Look out as far as vision permits in the four principal lattice directions and identify the unoccupied site(s) having the most sugar
- If the greatest sugar value appears on multiple sites then select the nearest one
  - (If it appears at multiple sites the same distance away, the first site encountered is selected, the site search order being random)
- Move to this site
- Collect all the sugar at this new position

Multi-commodity agent movement rule **M**

- Look out as far as vision permits in each of the four lattice direction
- Considering only unoccupied lattice positions, find the nearest position producing maximum welfare
- Move to the new position
- Collect all the resources at that location

Agent replacement rule $\mathbf{R}_{[a,b]}$

- When an agent dies it is replaced by an agent of age zero having random genetic attributes, random position in the sugarscape, random initial endowment, and a maximum age randomly selected in the range [a,b]

Agent trade rule **T**

- Agent and neighbour compute their marginal rate of substitution (MRS); if these are equal then end, else continue
- The geometric mean of the two MRS is calculated – this will serve as the price *p* [more elaborate bargaining could be constructed]
- The quantities to be exchanged are if p>1 then p units of spice for 1 unit of sugar; if p<1 then 1/p units of sugar for 1 unit of spice
- If this trade will (a) make both agents better off (increase their welfares) and (b) not cause the agents' MRS to cross over then the trade is made and return to start, else end

With these rules we first look at simulation ($\{\mathbf{G}_1\}$, $\{\mathbf{M}, \mathbf{R}_{[60,100]}\}$) applied when there is just sugar (no spice). The sugarscape grows back at 1 unit per time period (recall that agents harvest all sugar at a site when they move to it). Agents move as per the single-commodity rule and die when they reach their maximum age which is set at birth randomly between 60 and 100 time periods (or die of starvation).

Readers can run the simulation at by selecting experiment 3 at:

http://complexityworkshop.com/models/sugarscape.html

Agents swarm around the two sugar mountains, with some moving between the two. Those agents with metabolism equal to 1 can survive without moving on the "lowland" of sugarscape where the maximum capacity is 1 and regenerates each time period.

4.3. Wealth Distributions

The chart below is taken from the simulation after about 100 time units, it is an emergent structure, a stable macroscopic pattern, that statistical in nature. It shows the distribution of wealth: along the horizontal axis there is the range of wealth of all the agents split by decile; the vertical axis shows the number of agents in each decile. (The complexity workshop website has NetLogo embedded as the ABM software; NetLogo uses "turtle" as the label for "agent".)

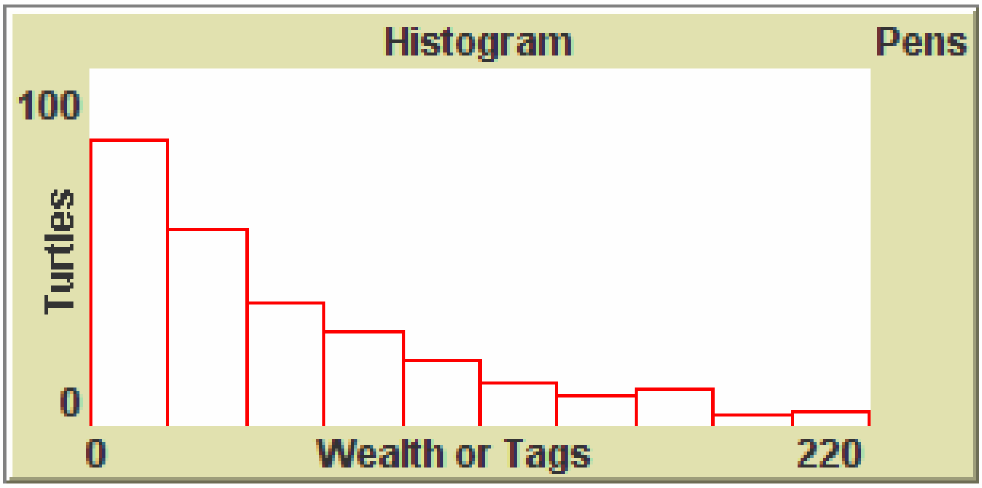

The distribution is Pareto: majority of agents have little wealth and a few have great riches. Recall that all agents are created equally with random variation of their genetic characteristics; no agents were favoured. It is the wealth distribution that is observed in the real world. Its growth in a simple model like Sugarscape is a powerful statement of what ABM can achieve.

4.4. Sugarscape with trade

The second simulation is ({$\mathbf{G}_1$}, {**M**, **T**}) in the sugarscape with both and spice. The movement rule is modified as above and a trade rule has been added. The death by old age rule has been dropped for simplicity as the focus is on trade as agents seek to better their welfare. We can observe each trade (volume transacted and price) in a time period and hence calculate averages for what actually happened at that time period. We can also "ask" each agent its situation regarding supply and demand for sugar or spice and corresponding price, eg how much sugar would it want to buy at what price if there was unlimited supply. This gives the aggregate supply and demand by price: the supply and demand curves of neo-classical economics with the intersection being the equilibrium volume of trade and price.

A snapshot of the supply/demand vs price, actual and equilibrium, that results is given below. We see the shape of the supply and demand is textbook, but remember that it is an emergent feature, not an assumed structure. We also see that equilibrium is never attained – we have a model to explore the dynamics of non-equilibrium.

Animation IV-2. Evolution of Supply and Demand under Rule System ({$\mathbf{G}_1$}, {**M**, **T**})

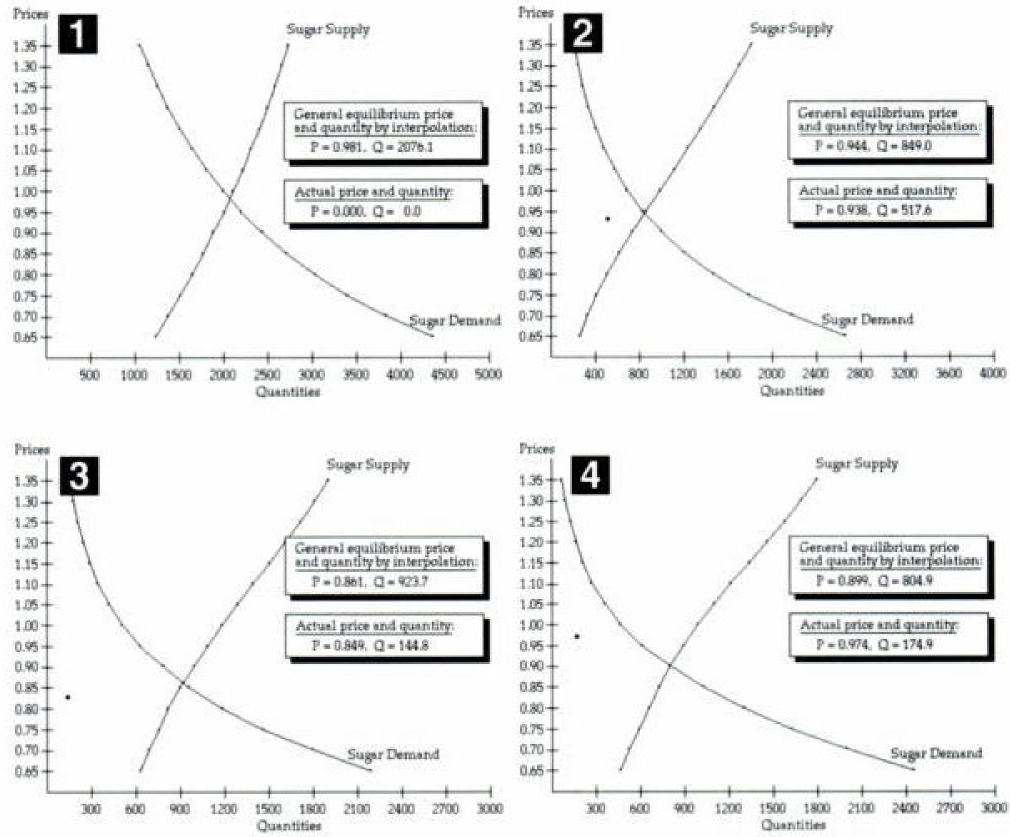

A movie showing a longer time evolution is available at
http://www.brook.edu/es/dynamics/sugarscape/animations/AnimationIV.mov

The charts below show the development of the equilibrium and actual price and quantity for the first 150 time units. The key observations are:

- The volume traded is always below equilibrium level (not surprising as agents can only trade with immediate neighbours and not on a wider scale basis)
- The average equilibrium quantity is 906.6, the average actual quantity is 212.5
- The equilibrium price or quantity is not stable
- Nor is the actual price or quantity
- The actual price can be either above or below the equilibrium price (it is above 56% of the time).
- The average prices are similar: for equilibrium it is 0.895 and for actual it is 0.911
- The actual price or quantity exhibits greater volatility than equilibrium
- The standard deviation for actual price is 0.0762 (8.4% of the mean) whereas the standard deviation for equilibrium price is 0.032 (3.6% of the mean)
- The standard deviation for actual quantity is 146.9 (69.2% of the mean) whereas the standard deviation for equilibrium quantity is 89.4 (9.9% of the mean)
- The distribution of price is visibly non-Gaussian with evidence of fat tails.

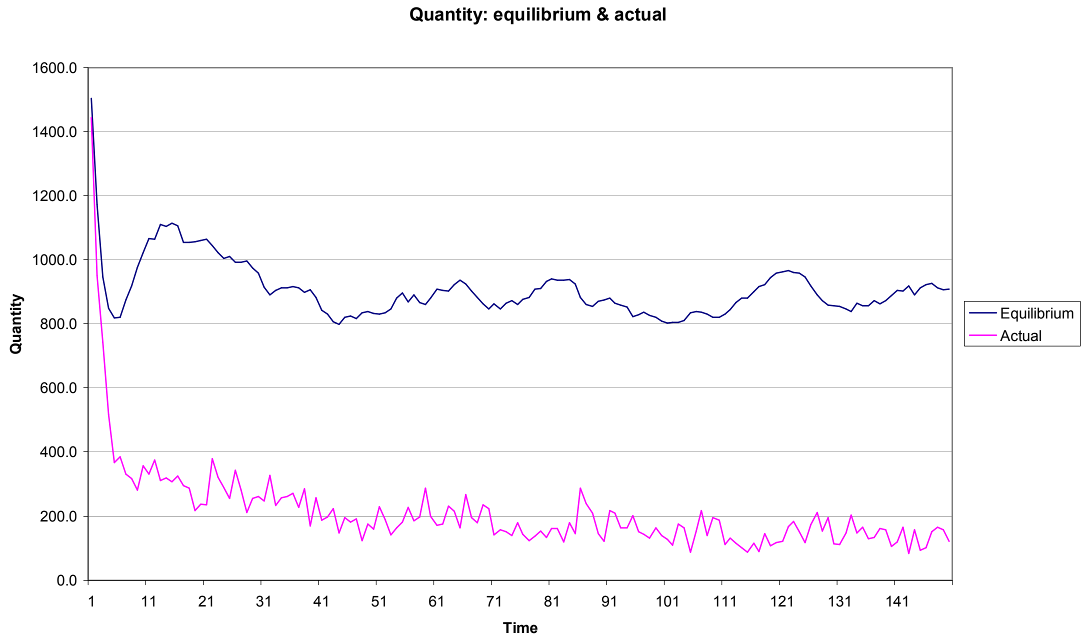

Quantity: equilibrium & actual
Quantity
1600.0
1400.0
1200.0
1000.0
800.0
600.0
400.0
200.0
0.0
1
11
21
31
41
51
61
71
81
91
101
111
121
131
141
Time
Equilibrium
Actual

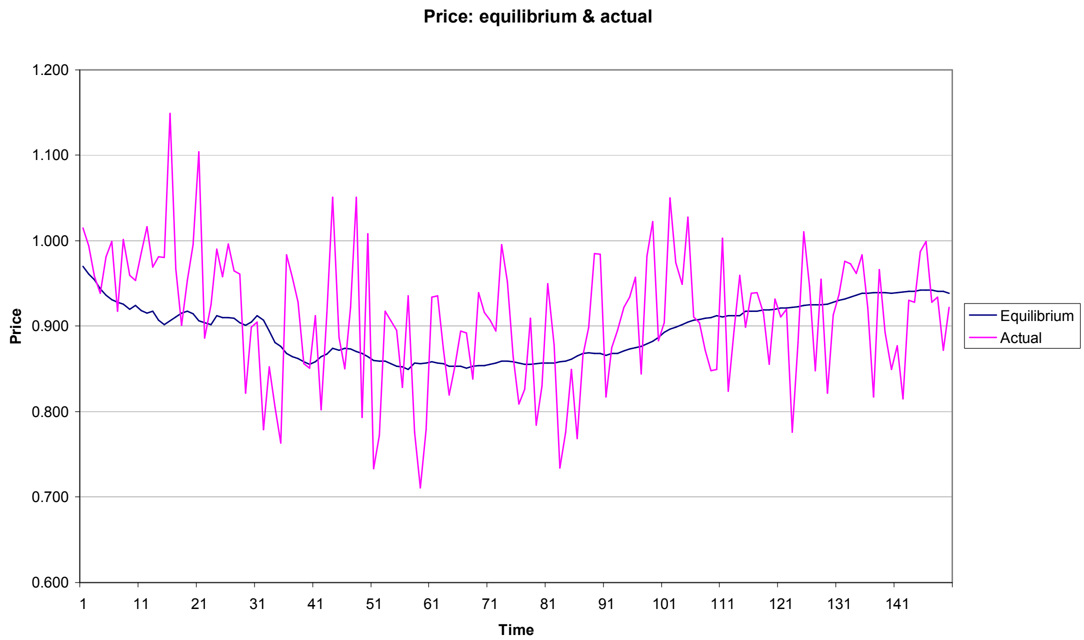

Price: equilibrium & actual
Price
1.200
1.100
1.000
0.900
0.800
0.700
0.600
1
11
21
31
41
51
61
71
81
91
101
111
121
131
141
Time
Equilibrium
Actual

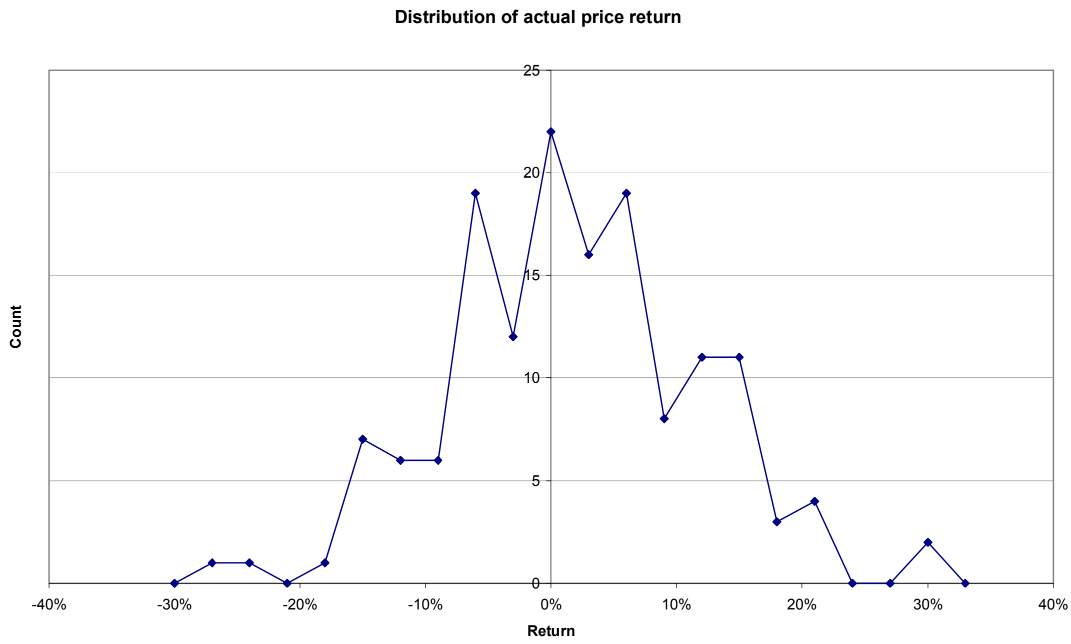

The movement of agents is also more complicated. Agents can no longer swarm on a single mountain top as agents need both sugar and spice to survive and these mountains are not coincident. Movements of agents which resemble trade routes emerge.

5. Agent-Based Stock Market Models

This section is largely a summary of an earlier and much more detailed paper (Palin, 2003, http://citeseer.ist.psu.edu/palin02agentbased.html) which contains a mathematical specification of the model and further results and commentary.

5.1. Motivation

Stockmarkets and other financial markets share many stylised features. These include a distribution of returns that is more peaked and "fat-tailed" than the Gaussian distribution; periods of persistent high volatility; periods of persistent high trading volume; and correlation between volatility and trading volume.

Traditional economic models have tended either to use a simple distribution of returns such as the Gaussian and treat extreme events as outliers, or to construct a statistical process which reproduces some of these features. But such an approach is purely descriptive and offers no understanding of why these characteristic properties should be seen and persist across so many markets.

To demonstrate a true understanding of the origins of these properties we would like to build a model which (i) produces these features without them being "hard-coded" into the model; and (ii) allows us to turn these features on and off by changing parameters of the model.

5.2. Description of Le Baron's model

We'll focus on a model by Blake Le Baron. While there are a number of models with broadly similar features, Le Baron's model is attractive because it is makes uses of many features of classical economics.

The model has just two assets: cash, with a constant rate of return; and an equity which pays a random dividend in each monthly time-step. For the equity only the dividend is specified (by a lognormal process); the price of the equity is an emergent property of the model which arises through the interactions of the agents.

The model also has a large number of agents. Unlike the earlier Sugarscape model there is no geographic distribution of agents: instead they all compete equally in the same market. Agents have an initial wealth and must decide at at each time-step how much wealth to consume, and how to invest their wealth. The agents make their decisions in order to maximise lifetime utility. This framework allows the agents to draw on standard results from lifecycle finance (due to Samuelson and Merton) in order to make their decisions; but in order to do so the agents must take a view on the distribution of equity returns over the next time-step.

The agents form their views on future equity returns by drawing on a pool of trading strategies. Each agent has several hundred rules which turn historic market information into a prediction. The rules are in the form of a neural network, which allows great flexibility in potential rules and can encompass rules which we might describe as "value", "growth", and "momentum", as well as complex combinations of these. The agents monitor the success of each of their rules over previous time-steps and act upon the rule which has proved most successful. Different agents have different lengths of "memory" over which they assess their rules. Some may only look at the last ten time-steps while others may use hundreds of time-steps.

The market price is determined using a "Walrasian auction" Under this method each agent says how many shares he would like to buy or sell at any given price, and the auctioneer chooses the price so that total supply and demand are equal. This method has been used in real stockmarkets in the past but is not in common use currently.

The model also contains elements of adaptation and evolution. If a rule has proved unsuccessful over a number of time-steps then it well be replaced by another rule. And agents which are unsuccessful in their trading will also be replaced.

5.3. Rational expectations price

A consequence of the grounding of Le Baron's model in economic theory is that we can calculate the "rational expectations price". This is the share price that would hold if all agents held the same rational views of the market. We find that the rational expectations price is equal to a constant multiple of the dividend; ie that the dividend yield is constant.

Pensions actuaries may recall that a decade ago it was common in a pension scheme valuation to use an actuarial value of assets rather than the market value. The "discounted dividend" method of assigning an actuarial value to equities is similar to using the rational expectations price. The agent-based model allows us to consider under which conditions the rational expectations price would or would not be seen in practice.

5.4. Results from the model

We'll look at a few different settings of the model. In each case the model is run for many time periods to avoid any artefacts caused by the initial conditions.

We first look at the case where all agents have a "long memory": they choose which rule to use based on performance over a period of around twenty years (240 monthly time-steps). In this case (Figure 1) we find that the share price does converge to the rational-expectations price, and the share price in the chart below is lognormal (since it just a multiple of the dividend which is itself lognormal). We also find in this case that there is no trading, since each agent is happy to maintain the same constant mix of cash and equity.

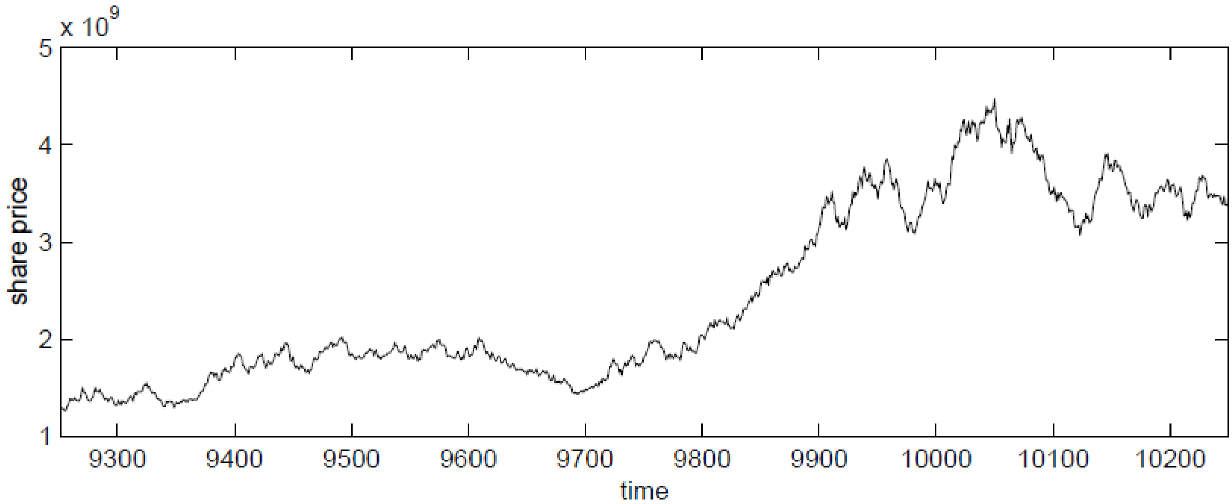

Figure 1: The "long memory" case. The rational expectations price is observed.

We can contrast this with the "all memory" case where agents' memories vary from six months to over twenty years. In this case (Figure 2) we see significant deviations from the rational expectations price with bubbles and crashes. We also find that in contrast to the "long memory" case there is significant trading volume as agents change their views in response to the changing market. The trading volume is correlated with changes in share price, with trading volume being particularly high during the instability around time 10000 and 10200.

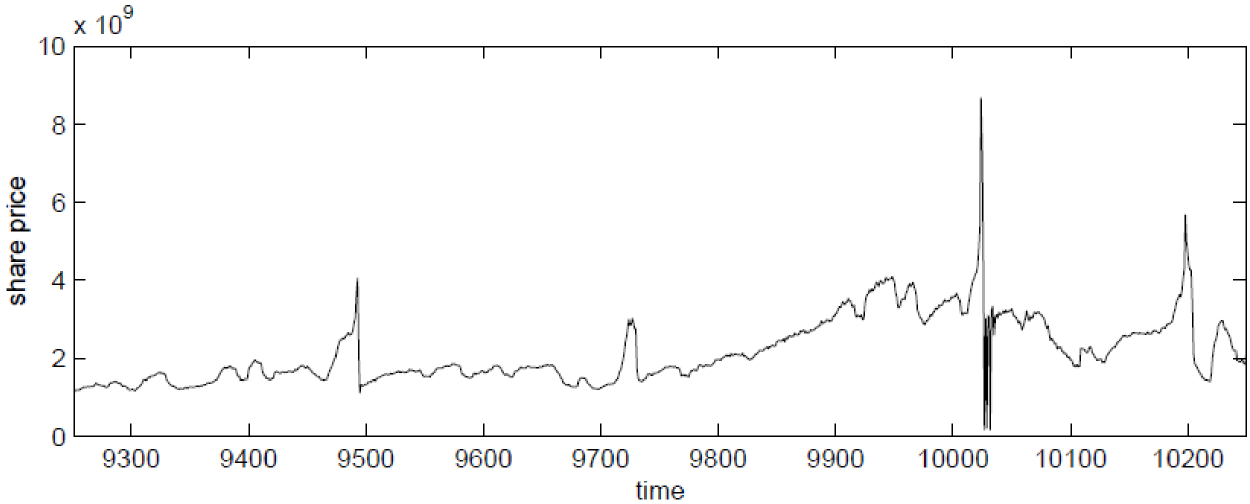

Figure 2: The “all memory” case. Significant deviation from the rational expectations price.

We would like to know whether the deviation from the rational expectations price in the “all memory” comes from the diversity of agents’ memory lengths or just from the presence of “short memory” agents. To address this we consider the “short memory” case where all agents have memories that range from six months to three years. This shows rapid cycles of bubbles and crashes.

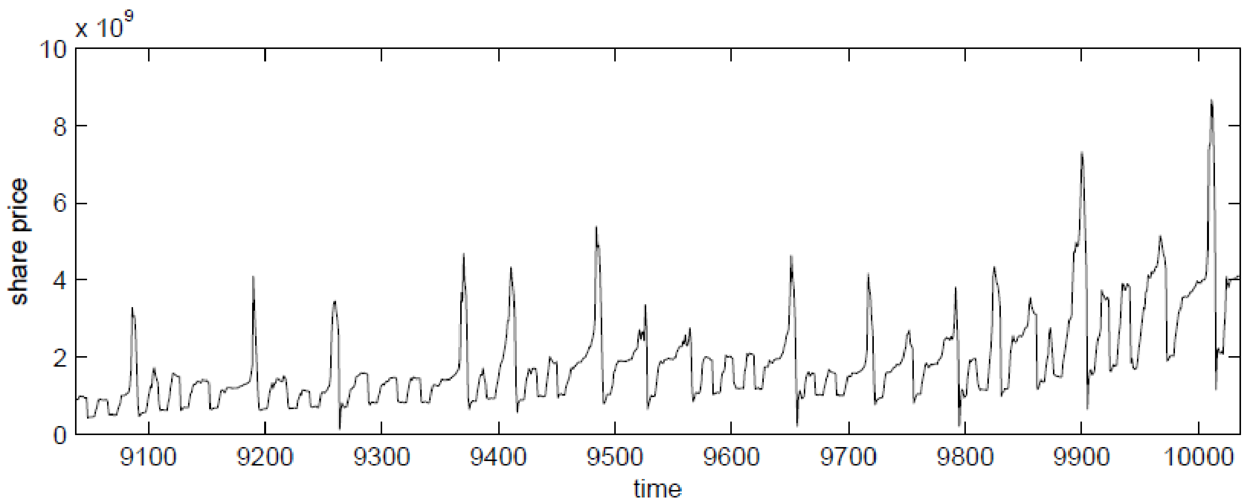

Figure 3: The “short memory” case. Very rapid cycles of bubbles and crashes.

5.5. Cautions and criticisms

Le Baron’s model is a partial success. It allows us to investigate how a share price can emerge from the actions of individual agents, and it shows how changing a parameter of the model (agents’ memory length) can move from lognormal share prices to stylised features of markets with fat tails, bubbles and crashes, and persistent volatility.

But while we can get a good qualitative result it is difficult to find a calibration which gives a good quantitative calibration. Changing memory length gives a sudden shift from rational expectations price to a wild caricature of the real world (the kurtosis of returns is an order of magnitude too high).

Also different models give different apparent reasons for features of real markets. One model claims that a calibration with fast evolution of agents and strategies is necessary to see fat tails; another shows fat tails without any evolution.

Although agent-based stockmarket models as they currently stand do not seem ready as an investment tool they remain a very useful tool for investigation. The processes which drive stock returns are not well understood with much commentary, particularly in the popular media, being glib and applied only after the event. The complex dynamics seen in LeBaron's model are valuable in showing that simple explanations of stock prices in terms of a single event are probably wrong.

6. Agent Based Models and Finance

When the authors of this paper were setting out on our careers, actuarial valuation methods were typically deterministic, based on expected outcomes minus implicit or explicit margins for prudence. On the pensions side, this is the "funding" approach to liability valuation (Exley et al (1997)), while in insurance liabilities were valued by reference to "reliable yields" (a calculation still required). The underlying methodology was challenged in the 1990s by a series of papers arguing in favour of market valuations[2]. These were backed up on strong modern financial economic foundations, a discipline which hitherto had had little impact on actuarial practice. These papers were highly influential in convincing the profession to alter practice, for example the adoption of market based valuations, which are now generally common practice in actuarial valuations and for accounting purposes in pensions and insurance. Alongside this move to market valuation was the growth of computing power which allowed the deterministic approach to be replaced with a stochastic one.

Agent-based modelling is not a mere tweak to neo-classical finance theory, but an alternative world view which matches many participants' intuition of what real markets look like – not necessarily a guarantee of their veracity. For example, agnet based models are inherently lumpy, while much of classical finance makes assumptions about infinitely divisible agents and investment markets in order to apply differntial calculus. Some agent based models consider market and intrinsic values separately, and explain how these two numbers diverge for significant periods of time. The agent-based model generates bubbles, crashes and high volatility endogenously, rather than exogenously as in the efficient market hypothesis (Farmer (2001)). For anyone who works with real markets, they are intuitively appealing.

6.1. Financial Uses of ABM

The use of agent based models so far has been dominated by aspects of explanation rather than prediction. Consideration of explanation and prediction need to differentiate between qualitative and quantitative, issues of calibration being more important for the latter. For example a model (agent based or otherwise) of financial

[2] For example Exley et al (1997) and Sheldon and Smith (2004)

markets needs to generate bubbles as they have been observed. A model which does not is clearly defective in either explanative or predictive capability. Having generates bubble behaviour it is for calibration techniques (which may not currently exist) to achieve a realist scaling between model and observation. Models which have passed the qualitative test may fail the subsequent quantitative test.

6.2. Modigliani, Miller and Agents

A series of papers 1990s laid the foundations of what we now call "market consistet modelling". These papers applied the emerging discipline of finance economics to actuarial theory. For example, a seminal paper was Exley et al (1997) which applies the Modigliani – Miller theorem (that the value of two firms is the same irrespective of their financial structures (Modigliani – Miller (1958))) to defined benefit pension scheme valuation and investment strategy.

The paper argues that a shareholder of a company with a defined benefit pension scheme has three equivalent ways of changing his asset mix:

1. Altering his directly held assets
2. Changing the balance sheet of the company
3. To modify pension fund asset strategy (assuming the shareholder has power to do this)

The implications of this equivalence are profound as they imply that it is impossible to achieve an optimal investment portfolio, as the shareholder can alter his own investments instead of the pension funds.

The paper is now going to look at how we might take an agent based approach to modelling a pension fund. The aim of this exercise is not to describe a realistic model, but to question whether ABM could give rise to a different outlook to Modigliani-Miller.

Doyne Farmer of the Sante Fe Institute constructs two kinds of traders, a seasonal trader and a technical trader. There is no movement in the underlying "value" of the commodity in the market. The seasonal traders buy and sell in a predictable pattern – they could represent, for example, farmers or even companies who need to meet emissions targets. The technical traders are purely in it for the money and are allowed to develop a variety of trading strategies; the more successful a strategy is, the more capital the trader gets and the more they influence the market. The result is shown in Figure 4. What happens is that the market becomes "efficient" after about 5,000 iterations when the technical traders make money off the seasonal traders and "iron" out the predictable seasonal fluctuations. But after that, as you see, the model suddenly goes mad. This is because the technical traders, who have now acquired practically all the capital (as shown in graph B), start trading against each other, devising ever more sophisticated trading strategies which work for a while until another trader develops a better strategy (Farmer (2001)).

6.3. Pension fund as agent?

To answer this question let us imagine a scenario in which an equity and bond portfolio both gave the same performance over a 5 year period, but equities increased by 50% over the 1st 3 years and then declined to meet the level of the gilts after 5 years – i.e. one not unlike our current experience.

At the end of this period, the shareholder is equally well off if he invested directly in bonds or equities. However, the outcome of shareholder value if the pension fund invested in equity or gilts might be quite different. If the pension fund invested in equity, at some point there would be an actuarial valuation, which would show a surplus – this could give rise to a reduced contribution rate and possible change to a more aggressive investment approach. The reduction in contributions would mean that the company has more cash available to invest in other activities. The positive balance sheet may boost the companies' share price, making capital cheaper which it then might invest favourably. The subsequent drop in equity value (exacerbated by the more aggressive investment policy) may put the scheme into deficit, causing the company to close the scheme, maybe boosting share value but also alienating the workforce.

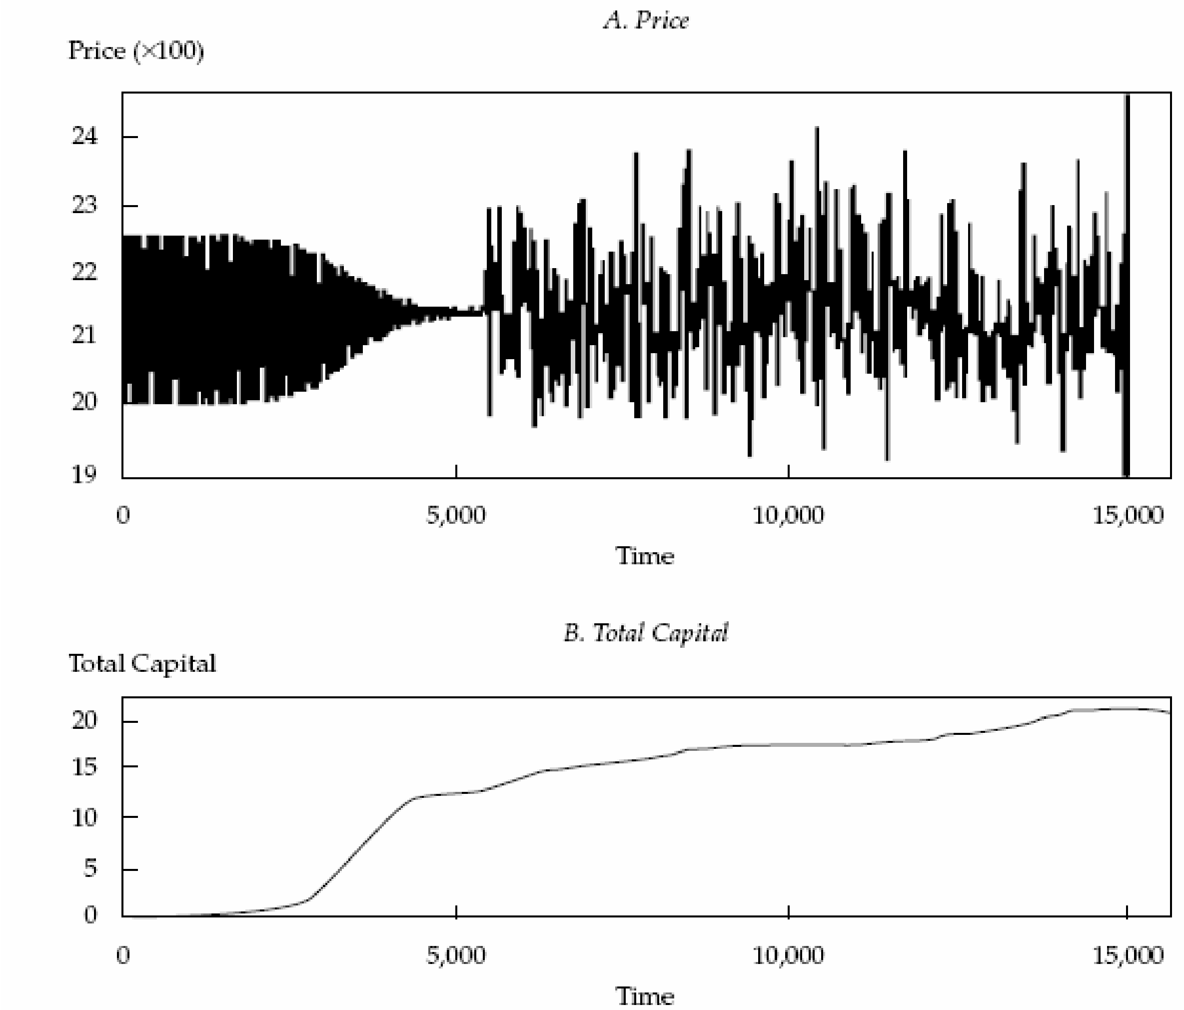

Figure 4 Behaviour of market with 2 types of trader (Farmer 2001)

This is the potential action of the pensions fund – we could in theory build up predictable investment rules for the pension fund, which could then be modelled as an agent. However, many pension funds tend to face the same issues at the same time. If we model pension funds as agents we could introduce them into the a Farmer or LeBaron type model, which are likely to challenge the Modigliani-Miller equivalence.

This scenario above is oversimplified– the aim is to demonstrates that the pension scheme's decision to invest in equity or gilts could have an effect on the pension scheme, the value of the company and shareholder value quite different to his decision to change his personal investments.

7. Implications for Actuarial Practice

7.1. Early Stages

Agent-Based Modelling is still at an early stage. What we have now is a collection of example computer models. The inputs are a collection of agents with assumed behaviour patterns, and ways of interacting. The outputs include simulated paths of financial variables, which in most cases share prices.

The models we have, and future we may build, allow us to conduct many experiments. With a combination of carefully controlled experiments and the growing mathematical insights from the field of dynamical systems, we can aspire to a better understanding of what in the model is most important at driving the outcomes. For example, is the absolute number of agents critical, or is there some point beyond which the model is "large" and adding extra agents does little to change behaviour. How could we measure the diversity of behaviour among agents? Is it the absolute level of irrationality or the difference between rationality of different agents that matters most? To what extent do different forms of irrational behaviour cancel out in observed market behaviour? Does the behaviour of all agents contribute equally to the modelled behaviour, or does a cadre of leaders emerge whose behaviour is more influential on the market as a whole?

7.2. The Need for Calibrations

We are far from understanding these issues at the current state of knowledge. But this is likely to improve as experimentation and analysis proceeds. An ability to rank inputs in order of importance is a pre-requisite for empirical calibration. The ranking tells us what aspects of real trader behaviour should be captured in order to build a useful model of the economy. Having determined the most important inputs, calibration might proceed by surveys, interviews and trading on market simulators. All this provides a way to calibrate the behavioural aspect of the models.

Without a calibration, the model tells us nothing concrete about the real world. Models of efficient markets have existed for some time. ABM provides us, in addition, with many alternative models of different ways markets could be inefficient.

Building an abstract computer model of inefficient markets cannot prove that markets are inefficient. Efficiency, or otherwise, is an empirical question for which disciplined observation of the real world is critical. That is why calibration is so important. A fully calibrated and tested ABM, with demonstrable predictive power, can add to our knowledge about real markets, not just to our knowledge of computer models. We are not there yet, so a view that ABM will one day provide insight into market efficiency remains, for now, a hunch, albeit a widely held hunch among ABM enthusiasts.

7.3. Rigour in testing Goodness of Fit

Informal model tests involve checking that the model accounts for observed data patterns. A more rigorous process takes account of the number of parameters estimated. The argument is that, with a large enough parameter count, a model can reproduce any data set. The trick in classical statistics is to find a parsimonious model than explains as much as possible of the observed data with a small number of parameters. Neo-classical optimising agents fit well into this framework; a parsimonious objective function is specified, and all aspects of behaviour are in single-minded pursuit of the chosen objective. Specifying the behaviour of a non-optimising agent requires an exhaustive list of responses to any eventuality. Without any over-arching optimisation at work, the potential number of rules is vast. As a result, ABM can generate many thousands or millions of parameters. In a statistical sense, such high parameter counts imply a high threshold for the model's ability to explain real world outcomes.

7.4. Forecasting Ability

Let us now conduct a thought experiment, and imagine we have a calibrated and tested ABM. What could we do with it?

The most obvious use is market prediction. A rejection of market efficiency is a rejection of the notion that returns are unpredictable. So an ABM that consistently predicts market prices (more accurately than a random walk) would pose a substantial challenge to theories of efficient markets. Orthodoxy might fight back, for example by rationalising observed biases as a reward for risk, but if the rewards are large enough then a latent risk is a less plausible explanation. The implications for portfolio construction and investment strategy are profound. We do not know how many of these are in use already; the owner of such money-making machine may be reluctant to publicise the fact. Conversely, if ABM's do not ultimately give better forecasts than a random walk, then the ABM has instead illustrated how irrational agents can nevertheless contribute to market pricing – a result that would also be of considerable interest albeit less lucrative.

It is less clear how other financial theories would need adapting. The Black-Scholes model for option pricing does not rely on market efficiency. It does rely on being able to trade continuously, without transaction costs, and on having accurate volatility forecasts. These conditions do not fully hold in real markets, yet the Black-Scholes and related models are still widely used. In our view, it is unlikely that the increasing use of ABM would dent the use of Black-Scholes model, but it could well provide new insights into the elements driving an appropriate choice of volatility parameter.

7.5. Adapting Modigliani and Miller

The work of Modigliani and Miller on capital structure has had a profound impact on recent actuarial thought, particularly in pensions. The underlying arguments are based on arbitrage and do not require markets to be efficient. However, to exploit the arbitrage of a violation of Modigliani and Miller, would require an arbitrageur to take offsetting long ands short positions in two shares with equal core businesses but different capital structures or different pension plan investments. Such trades are in practice impossible to execute, because companies do not come in such convenient pairings. We can instead turn to equilibrium arguments to imply the same (M&M) conclusions, but now we have introduced the additional hypothesis that economic agents are trying to optimise something. Take away optimising agents, and the arguments become much more delicate. Everything becomes very model dependent; for example under Modiglinai and Miller (and excluding tax, bankruptcy or benefit leaage) an extra €1000 pension fund contribution should have no effect on a company's share price. The usual lesson taken from this argument is that optimal contributions much therefore take account of tax, bankruptcy and leakage effects. However, in an agent based model, we cannot easily unpick the effect of contributions on share prices. ABM might suggest we need to know a great deal more in depth about the investor community in order to recommend pension contribution rates.

7.6. Feedback Loops

Some ideas of market inefficiency are already embedded in regulatory thinking. If markets are efficient, and prices reflect best information, then there can be little market benefit from regulatory interference in trading volumes. If a regulator prevents an institution from selling a risky asset, then the share price of the institution itself falls, so the fall in aggregate value of the economy's assets is unchanged by the regulation. One could even argue that a regime discouraging fire sale of assets expose policyholders or pension plan members to further risk, because in a default event it is policyholders or plan members who have to scrape around the remaining assets to recover some of their promised benefits.

Under agent based models, however, positive feedback loops may operate to exacerbate the effect of crashes. Relaxation in solvency regulation, and corresponding easing of the pressure on institutions to sell risky assets, could mitigate the market fall itself. This increases aggregate economic wealth and reduces some of the refinancing costs that otherwise might have applied. As calibration technology develops, we may well see a day when ABM provides meaningful support to regulators searching for the best way to protect plan members or policyholders in turbulent financial times.

References

Beinhocker, E (2007) *The origin of wealth* Business Books

Bonabeau, E (2002) *Agent-based modelling: Methods and techniques for simulating human systems* Icosystems Corporation

Brealy, R and Myers, S (2007) *Principles of corporate finance* McGraw-Hill Education

Copeland, T and Weston, F (2004) *Financial theory and corporate policy (4th Edition)* Addison-Wesley

Exley, C Mehta, S and Smith A (1997) *The financial theory of defined benefit pension schemes* British Actuarial Journal 3, IV 835-966

Fama, E and French, K (1992) *The cross section of expected stock returns* The journal of finance Vol XLVII, No 2.

Farmer, D (2001) *Toward agent-based models for investment* Association for Investment Management and Research

Holland, J (1998) *Emergence* Oxford University Press

Mandelbbrot, B and Hudson R (2006) *The (mis)behaviour of markets* Profile Books

Mandelbrot, B (1963) *The variation of certain speculative prices* The Journal of Business, 36(4):394–419

Mandelbrot, B (1967) *The variation of some other speculative prices* The Journal of Business, 40(4):393–413

Miller, J and Page, S (2007) *Complex adaptive systems* Princeton Paperbacks

Modigliani, F and Miller M (1958) *The cost of capital, corporation finance and the theory of investment* American Economic Review, **48**, 261-297

Palin, J (2002) *Agent-based stockmarket models: calibration issues and applications* University of Sussex

Ralfe, J Speed, C and Palin, J (2003) *Pension and capital structure: why hold equities in the pension fund?* Society of Actuaries

Sharpe, W Alexander, G and Bailey, J (1995) *Investments* Prentice Hall Fifth Edition

Sheldon, T and Smith, A (2004) *Market consistent valuation of life assurance business* Institute of Actuaries

Taleb, N (2007) *The black swan* Random House

Varian, H (1999) *Intermediate microeconomics* Norton 5th Edition