

Policy Analysis from First Principles

**Scott Moss
Centre for Policy Modelling
The Business School
Manchester Metropolitan University
Manchester M1 3GH
United Kingdom**

s.moss@mmu.ac.uk

<http://www.cpm.mmu.ac.uk/~scott>

Abstract

The Sackler Colloquium on on Adaptive Agents, Intelligence and Emergent Human Organization is intended to force a re-examination of current social theory by means of an exploration of adaptive agent models. The primary issue addressed by this paper is why it should be assumed that adaptive agent models are suitable for that purpose. The argument is predicated on the view that social science should start with observation and the specification of a problem to be solved and then reason from that basis to define the appropriate properties and conditions of application of relevant tools of analysis. Accordingly, evidence is adduced from data for sales volumes and values of a disparate range of goods to show that, in all presented cases, frequency distributions are fat tailed. This result implies that, if there is a suitable and stable population distribution, it will generally have infinite variance and perhaps undefined mean. Agent based models with agents that reason about their behaviour and specifically do not conform to the rational choice model and are influenced by, but do not imitate, other agents known to them will typically generate fat tailed time series data. There is no reason to suppose that this data is drawn from any stable distribution. An agent based social simulation model of intermediated exchange is reported that has the same type fat tailed time series and cross sectional data that is found in data for fast moving consumer goods and for retail outlets. This result supports the proposition that adaptive agent models of markets with agents that reason and are socially embedded have the same statistical signatures as real markets. This statistical signature precludes any conventional hypothesis testing or forecasting – a point which pertains both to intermediated (e.g. retail or financial) markets in general and to the models. However, these agent based social simulation models offer unique opportunities for validation on the basis of domain expertise and qualitative data. While they cannot be used for prediction or forecasting the specific consequences of any policy or commercial action, they can be used both to identify differences in stakeholders' perceptions and to investigate the properties of effective responses to unpredictable events.

1 The issues

*The purpose of this colloquium is to explore adaptive agent models and, in so doing, to force reexamination of current social theory and encourage rethinking of the processes by which human organization emerges.**

The presumption in the meeting overview, from which the above passage is extracted, is that adaptive agent models are particularly well suited to capture the nature and consequences of social complexity whereas current social theory is not. Two aspects of current social theory are identified in the meeting overview: the dominance of social theory based on the assumption that economic actors maximise utility and the assumption that “social organization evolves from a top-down hierarchical system of culture and norms that serves to shape individual behavior.” It is then asserted that “Adaptive agents methods are likely to become the foundations for modeling and simulation that may help to resolve many of the problems of complexity and help in the development of policy tools that provide enhanced insight into the likely effects of policy action.”

Why should adaptive agent models be relevant to, much less force the reexamination of, current social theory? And why might adaptive agents models usefully inform the provision of enhanced insight into the likely effects of policy action? After all, agent based simulation is hardly the first analytical approach for which great promise has been claimed. Game theory and then dynamic game theory was going to provide powerful and relevant models of competition. Yet a recent survey [1] of game theory papers in a leading economic journal showed that game theoretic analyses of processes was limited to two or in one case three agents and all n-person game theoretic models were concerned with equilibrium outcomes rather than any process. Econometrics and Keynesian theory together were going to provide means of forecasting the effects of policy actions provided that the number of policy targets were the same as the number of instruments. I recently asked on the email discussion list of the International Institute of Forecasters whether there are any counter-examples to the following claim:

Since the invention of econometrics by Jan Tinbergen in the 1930s, there has not been a single correct econometric forecast of an extreme event such as a turning point in a trade cycle or a stock market crash. Every such forecast – without exception – has yielded either a type I or a type II error.

Apart from one undocumented claim, the strongest responses were that, when applied to past data, some new modelling techniques look better than most previous modelling techniques. No one was able to point to a correct forecast.

This experience, and many like it across the social sciences, is reason enough carefully to investigate the claim that agents and simulation are indeed promising elements in a new approach to policy analysis. A careful investigation specifies the problem to be addressed and does not alter the problem specification in order to conform to the requirements of any tool of analysis. The selection of any analytical tool to be applied to the problem, or the requirements analysis of the properties of any such tool is to be based on available and applicable empirical evidence. The tool of analysis considered in this paper is broadly in the class of agent based social simulation models.

* From the programme of the National Academy of Sciences Sackler Colloquium on Adaptive Agents, Intelligence and Emergent Human Organization.

The whole range of policy analysis in complex environments is much too broad to be the subject of the sort of careful investigation suggested here. It is however possible to address a class of policies: those that seek to use some market or competitive mechanism to manage resources. European privatisations of public utilities and transportation systems or the use of internal markets in the UK's National Health Service or proposals for carbon trading and carbon taxes to mitigate the extent of anthropogenic global climate change are examples of such policies. The choice of this class of public policies is motivated both by the importance of the policy goals and by the importance of representations of markets in the development of much of the current social theory to which the meeting overview refers.

In order to evaluate policies intended to use or create markets, we require analytical tools that capture the properties of markets that the policies are intended to exploit. For this reason, we begin in section 2 by a consideration of data describing the demand for and sales of a range of goods. The result is a demonstration that even most mundane goods are subject to the same sort of volatility and uncertainty that is found in financial market prices or earthquake magnitudes or any of a wide range of social and natural phenomena. This finding motivates the discussion in section 3 of the relevant statistical issues from the standpoint of both econometrics and statistical physics. Three mutually exclusive possibilities consistent with the available data are canvassed, one of which is consistent with current social theory and two of which are consistent with adaptive agent models. In section 4, a model is reported to support the investigation of the implications adaptive agent models for policy analysis. In section 5, we explore some wider issues concerning the ways in which it is justifiable to adaptive agent models for policy analysis. Perhaps the most striking conclusion is that neither current social theory nor any similar construct will ever support an effective policy analysis. However, adaptive agent modelling is an effective substitute.

2 The statistical signatures of competitive intermediated markets

If a competitive market is one where neither buyers nor sellers are able to set the prices in their transactions, then pricing is either the outcome of a process of negotiation or some third party must set the price. Such third parties are well known and include the market makers on the financial exchanges, retail shops and processors and sellers of information such as credit agencies. An important policy question is whether competitive intermediated exchange is appropriate for allocating public services or the services of privately owned public utilities and, if this is an appropriate arrangement, under what circumstances is it appropriate.

Economic propositions about the efficiency and social benefits of competition are based on equilibrium models. It is therefore worth asking whether equilibrium models actually provide the best available descriptions of the phenomena we observe. We begin with statistical observation.

Weekly scanner data from supermarkets show that sales of fast moving consumer goods such as alcoholic beverages in the US and the UK and tea, biscuits, shaving preparations and shampoo in the UK are marked by the kind of clustered volatility that we associate with asset prices in the financial markets.

Benoit Mandelbrot [2] first noted that log price changes in financial markets are commonly power law distributed. He pointed out that this phenomenon is consistent with a stable Paretian population distribution. The value of the stable Paretian distribution is that there is a known functional relationship between the moments of a distribution constructed by multiplying a constant by observations drawn from other

distributions and the moments of those other distributions. These functional relationships underlie such important statistical techniques as regression and correlation analysis.

The characteristic function of the stable Paretian distribution takes the logarithmic form

$$\log \int_{-\infty}^{\infty} \exp(iuz) d \Pr(U < u) = i\delta z - \gamma |z|^{\alpha} \left[1 + i\beta \left(\frac{z}{|z|} \right) \tan \left(\frac{\alpha\pi}{2} \right) \right]$$

where α in the interval $(0,2]$ is a “peakedness parameter” and β in the interval $[-1,1]$ determines skewness. Together with values of δ and γ , these parameters determine the first four moments (mean, variance, skewness and kurtosis) of the distribution.

The value of α is of the most interest here. When $\alpha=2$, this characteristic function reduces to that of the normal distribution. For all values of $\alpha < 2$, the variance of the distribution is infinite. Moreover, for values of $\alpha < 1$, the mean of the distribution is undefined. This means that, for all values of $\alpha < 2$, the law of large numbers does not apply to the variance and for all values of $\alpha < 1$ the law of large numbers does not apply to the mean. No laws or theorems of classical statistics or econometrics are applicable in these circumstances.

It is, however, important to note that the central limit theorem will typically apply to data with a stable Paretian distribution. Since aggregating time series data – say daily into weekly data points – is effectively to calculate the mean of daily data over seven data points and then multiply by seven, samples of these sample distributions will be approximately normally distributed even if the underlying daily data is not. Consequently, it is possible to generate data that appears to be normally distributed simply by taking data of sufficiently low frequency. However, the variance and possibly the mean of the distribution of sample distributions will not converge to stable values.

Distributions with infinite variances are easily distinguished visually from normal distributions because, for the same means and variances, they have fatter tails and therefore thinner peaks – a condition known as *leptokurtosis*. Figure 1 shows clear evidence of leptokurtosis in the weekly sales values of three brands of shampoo in UK supermarkets for the 65 weeks beginning 2nd January 2000. Similar results are found for virtually every one of the 120 or so brands of shampoo for which I have the data as well as every brand of tea, shaving preparations, biscuits and, in the US as well as the UK, every one of some 200 brands of spirituous alcoholic beverage and beers. The first row of Figure 1 shows weekly sales values. Brand A is a leading brand with no discernable sales trend while sales values of brand B are declining and sales values of brand C are increasing. Both of the latter have small market shares. The second row shows the time series of relative sales changes. Over the 65 weeks there obvious clusters of volatility and it is these clusters that generated the extreme values that cause the leptokurtosis evident in the third row showing the frequency histograms of the relative sales changes compared with the corresponding normal distribution.

These results are typical for all of the products considered as well as daily metered consumption of water in southern England. Leptokurtosis and clustered volatility is evidently far more general than has previously been recognised. Without doubt, the consequences for forecasting are enormous since leptokurtosis alone, independently of the clustering of volatility, implies the failure of the law of large numbers. That is, increasing sample size does not result in any convergence of any of the moments of the sample distributions – a result that itself renders parametric

statistical forecasting techniques wholly otiose (Mandelbrot, 1998). Clearly, if leptokurtosis undercuts the law of large numbers, then the clusters of extreme events cannot in principle be forecast by statistical means.

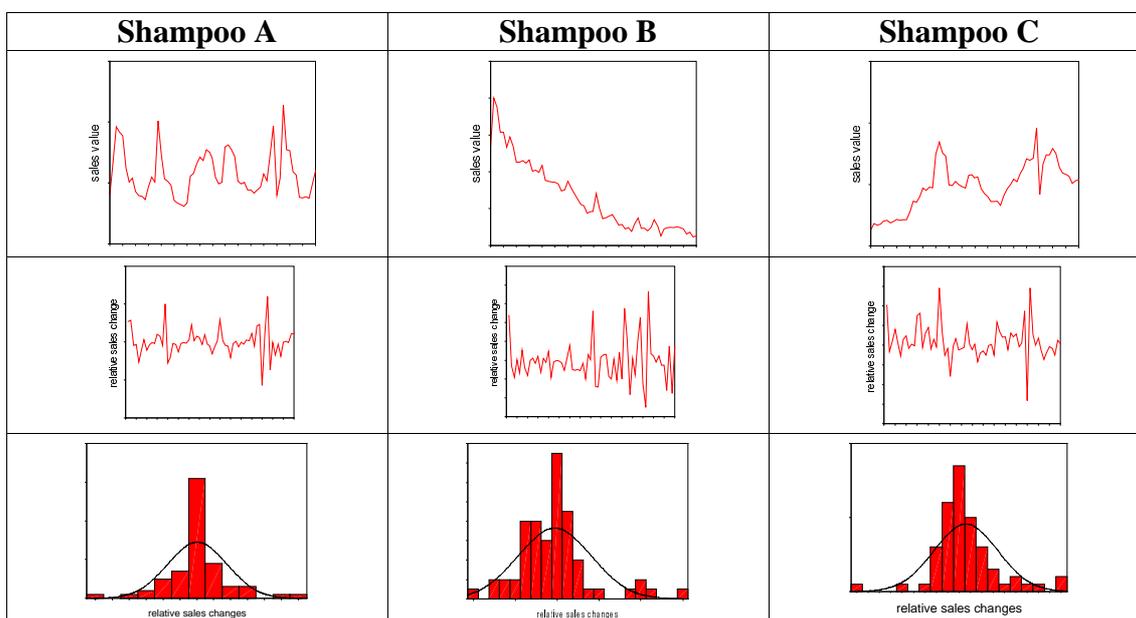


Figure 1: Weekly shampoo sales and relative sales change; 2 Jan 2000-25 March 2001
(Source: Information Resources International)

3 Three responses to leptokurtosis and clustered volatility

There have been two responses to this problem, both of them addressing the core issue of the failure of the law of large numbers in samples of the variance of time series data. One of these responses has been offered by econometricians and is in effect intended to preserve the law of large numbers and the other is from the physics community and is intended to bury the law of large numbers for good – at least when it comes to forecasting financial asset prices. A further response, drawing on agent based social simulation, is described below.

3.1 Time varying parameters (TVP)

The TVP approach is based on the assumption that the observed time series is drawn from a normally distributed population with constant mean and varying variance. The variance at the time of any observation is itself a function of previously observed errors so that standard regression techniques can be used to model the time series of variances. To capture the clustering of extreme events, the variance of a time series for any observation is treated as a function of previous error terms. Recent, large error terms tend to generate large variances. Because the observation is then drawn from a population with a larger variance, the probability of the observation being relatively far from the mean is greater than when the population variance is smaller.[†]

The motivations offered for particular TVP estimating methods are invariably related to rational expectations, the mean-variance representation of risk and risk

[†] Bollerslev [3] identifies the core econometric processes of relevance here to be the ARCH process [4], the GMM process [5] and GARCH [6].

aversion or some similar equilibrium notion from economic theory. There are, however, no microeconomic equilibrium models that generate both leptokurtosis and clustered volatility either analytically or by means of simulation.

3.2 Self organised criticality and econophysics

While there are no models demonstrating any microeconomic foundations of TVP-based leptokurtosis, there is an extensive class of models – both canonical and applied – that were developed and explored by statistical physicists because they do generate clustered volatility and hence leptokurtosis. The physical problem being addressed was the observation that an extraordinarily wide range of physical phenomena are power law distributed. The power law distribution is:

$$N(s) \sim s^{-\tau}$$

where N is the number of observations at scale s and $\tau > 0$ is a parameter. Mandelbrot (1963) pointed out that the power law distribution is a characteristic of the stable Paretian distribution.

The question of concern to statistical physicists starting with Per Bak and his colleagues [7] was to find the process that is both very general and yields power law distributed time series. The canonical model they developed was an idealisation of a sandpile with grains of sand being continually added. The sandpile model is closely related to a cellular automaton model in that it is located on a grid with non-periodic boundaries with grains of sand added to cells at each time step. Whenever the number of grains of sand in a cell reaches some specified critical level – say 4 – there is a “toppling” of the sand in that cell. This toppling takes the form of a redistribution of the grains of sand in the critical cell to other (not necessarily adjacent) cells in the grid. Not all of the grains are redistributed but the number of grains in the critical cell is nonetheless reduced to 0. That some grains are lost from the system in this way makes it dissipative.

Of course adding toppled sand to the grains at other cells increases the numbers in those cells until some of them become critical and topple and so increase the number of grains in yet other cells, and so on. The consequence is that, once the system reaches a critical state, there will be a sequence of topplings involving different numbers of cells in the grid. The time series of these topplings is power law distributed.

There is a growing family of such models that yield power law distributed time series and cross sectional data. The key feature of these models is that they do not require fine tuning of the parameters of interest in order to produce data with this statistical signature. In this sense, the models self-organise into the critical state and remain in that state thereby to produce such power law distributed data with clusters of extreme events.

Jensen [8] has summarised the conditions in which self organised criticality (SOC) emerges as those where:

- Model components (cells, agents, etc.) are metastable in the sense that they do not change their behaviour until some level of stimulus has been reached.
- Interaction among the model components is a dominant feature of the model dynamics.
- The model is a dissipative system.
- The system is slowly driven so that most components are below their threshold (or critical) states most of the time.

In social terms, agents and the individuals they represent are metastable if they do not respond to every minute stimulus they face. They would not, for example, reconfigure their desired shopping basket as a result of a penny rise in the price of a tin of tuna. A particular implication of metastability is that the behaviour of individuals cannot be represented by utility maximising software agents. The dominance of interaction among the agents amounts to social embeddedness in the sense of Granovetter [9] and Edmonds [10]: the behaviour of individuals cannot be explained except in terms of their interaction with other individuals known to them. Dissipation in a social system, analogous to the dissipation of grains of sand in the sandpile model, equates to individuals being influenced by other individuals without slavishly imitating them.

There is a literature on SOC models of financial markets though these appear almost entirely in journals such as *Nature*, *Physica A*, *Physical Review E* and *Physical Review Letters* and, apart perhaps from Lux [11], have had no noticeable effect on the economics literature.

There are several important differences between the self organised criticality and the TVP literatures. The former uses models to generate statistical signatures that do not replicate actual data series while the latter exploits the data to model changing values of the moments of a distribution function. This is part and parcel of a more fundamental difference: self organised criticality suggests that, without unlimited computational and information processing capacities, forecasting extreme events is inherently infeasible while the TVP literature is based on a faith (since there is no supporting evidence) that forecasting extreme events is feasible. A further difference is that SOC properties are, with a few special exceptions, known only from simulation experiments. There are hardly any analytical proofs [8]. The TVP literature, by contrast, is based on algorithms with properties that have been proved analytically.

3.3 SOC in social systems

One question that has not been considered in either the TVP or the SOC literatures is whether there is any stable population distribution underlying observed or simulated time series or cross sectional data. Since the TVP literature is concerned with the application of parametric statistical techniques, the assumption of a population distribution with stable characteristics is essential. The only question is where the stability lies. In the TVP literature there is assumed to be a stable function relating variance to previous deviations from the stable mean. In the physics literature on SOC, a measure of success is naturally taken to be the degree of agreement between the power law distribution parameter obtained by simulation and the parameter obtained from the corresponding real data. While this measure of success is appropriate in models of physical systems, it may not be an appropriate measure of success of models of social systems.

The key difference here turns on universality. The assumption that the laws of physics are always and everywhere the same has been enormously useful in the physical sciences. It is also a natural assumption to make since fundamental physical relations are plausibly unchanged by their own consequences. So the law of gravity does not vary because objects catastrophically collide, for example. It is much less plausible to argue that social relations are unchanged by their consequences. On a grand scale, it would be lunatic to suggest that social relations were unchanged by the French Revolution – and not just in France. More particularly, major financial panics frequently result in changes in the rules and practices of financial markets. Also,

institutional arrangements are altered by extreme natural events but the laws of nature are not affected by extreme social events.

If we consider the natural and social systems as two data generating mechanisms, then we observe that the data emerging from natural systems typically have stable statistical properties while data emerging from social systems typically do not. For example, earthquake magnitudes are power law distributed (the Gutenberg-Richter law) and the observed power law distribution does not change over time. So while it is not possible to predict the occurrence of specific earthquakes or earthquakes of specific magnitudes, it is possible to describe with considerable accuracy the distribution of earthquake magnitudes. However, in social systems experience of extreme events leads to a search for means of reducing their incidence if that is possible or their impact if it is not. The point is to change the observed distribution by, in effect, changing the data generating mechanism. To the extent that such social engineering is successful, the magnitude of the exponent of the power law distribution will be reduced. The goodness of fit of a power law distribution will be reduced although that will not eliminate the leptokurtosis and clustered volatility of the observed data.

4 Implications for scientific method

Three mutually exclusive explanations of observed leptokurtic data series were described above: a normal distribution with predictably time varying parameters, a stable Pareto distribution with infinite variance generated by a self organised critical social process and no stable distribution but also generated by a self organised critical social process. There are no tests on observed data that will distinguish between TVP, stable Pareto or unstable leptokurtic distributions. Clearly a systematic history of success at forecasting volatile episodes with TVP methods would give convincing support to the hypothesis that there is an underlying normal distribution with a predictably varying variance. However, there is no such history. The timing, magnitude and duration of volatile episodes remain in practice unpredictable. We must look to some other means of discriminating among – or rejecting – these possibilities.

In the natural sciences, the search for explanations of previously unexplained observations has taken the form of a search for a data generating mechanism that could be validated independently of the observations themselves. A classic example is the validation of general relativity theory by comparing observations of star positions during a total solar eclipse with the predictions of the theory. A more pertinent example is the development of the sandpile model to explain observed power law distributions and then experimental testing of the canonical model with sand, rice grains and the like [12].

Scientific methods that have proved to be successful in the natural sciences are not necessarily equally applicable in the social sciences. However, it is hard to see any objection to treating a social system as a data generating mechanism and devising a model to represent that mechanism. If the model captures self organised criticality, then it is not possible to validate it by statistical means for two reasons. One, of course, is that the generated data will be leptokurtic and therefore have infinite variance. There are no parametric hypothesis testing procedures for infinite variance distributions and non-parametric procedures provide information only about the data in hand. The second reason is that self organised criticality implies that the timing, magnitude and duration of clusters of extreme events are in practice unpredictable. There seem to be two, mutually exclusive ways forward.

One is to continue to develop the TVP approach in the hope that it will someday yield consistently accurate forecasts of extreme events. The other is to focus on the relevant system as a data generating mechanism and to devise means of modelling that mechanism *independently* of the statistical data itself.

4.1 From problem to approach

The issue of concern is policy analysis in conditions where the objective is to develop strategies to mitigate the impacts of clusters of extreme events the magnitude, duration and timing of which cannot be forecast. The alternative to forecasting suggested here is to try to understand the social processes generating the extreme event clusters in order to assess the effectiveness of different responses to their occurrence. The means to be chosen for understanding the underlying data generating process must obviously be able to capture a process yielding unpredictable extreme event clusters. This requirement filters out all equilibrium processes.

A second requirement is that the approach should be robust in explaining leptokurtosis and unpredictable clusters of extreme events. A mathematical model that generates data with the appropriate statistical signature only under a very narrow range of values of key parameters would not be appropriate unless there were independent evidence that those parameters and the particular values required robustly describe observed phenomena. While chaos and edge of chaos models meet the first requirement, there is some evidence that they do not meet the second. Differential or difference equation models with strange attractors have been known from the discovery of chaos to require parameters to be set to specific ranges. Kaufman (1993) has worked on models in which values of those parameters are driven to the chaotic range and are otherwise at “the edge of chaos”. However, Per Bak (1997, p. 127) reports that these results are not robust with respect to parameter settings and initial conditions. If Bak is wrong, then edge of chaos approaches will satisfy the second requirement. If he is right they will not.

The third requirement is that the approach must support independent validation. To date, there has been no independent validation of the edge of chaos models.

The approach investigated here starts expositionally from SOC. Historically, however, a set of models designed and implemented by researchers in the Centre for Policy Modelling (CPM) to analyse the effects of social embeddedness and to be open to validation by stakeholders turned out to yield leptokurtic data with clustered volatility. The research programme of the CPM that yielded these models and the models themselves were developed in ignorance of SOC.

Two examples of such models are Moss’ model of household water demand [14] and the effects of exhortation by government and other authorities during conditions of drought and Edmonds’ model of a financial market. Moss represented agent cognition by means of a combination of the problem space architecture of Soar [15] and ACT-R [16] together with an endorsements mechanism[17]. Edmonds represented agent cognition by means of an elaborated genetic programming algorithm. Both of these representations of cognition yield metastable agent behaviour in that some non-negligible weight of evidence and incentive is required to induce agents to change their behaviour. In both models, agents were socially embedded. In the Moss model, social embeddedness took the form of observation of neighbours’ public consumption activities such as garden watering and car washing as well as word of mouth communication. In the Edmonds model, there is word of mouth communication among agents. In both cases, agents were influenced in their behaviour by the behaviour of and communication with the subset of other agents

with which they had formed some relationship of trust and regard. This, of course, is the essence of social embeddedness.

In these social simulation models, the representations of agent cognition could hardly have been more different. Even the representations of the model spaces were different – Moss implemented a grid with periodic boundaries while Edmonds had no spatial location of agents. Indeed, apart from metastability of agent behaviour and social embeddedness, it is hard to identify common features in the design and implementation of the two models. The natural conclusion is that, on the basis of such experience, cognitive representations yielding metastable agent behaviour and social embeddedness of agents drive processes yielding leptokurtic distributions with clustered episodes of volatility.

It would be wrong to expect such social simulation models to replicate observed time series from their target systems. What is being sought in these model specifications is a shared statistical signature in the sense that the time series of both model and target systems are marked by unpredictable clustered volatility and therefore leptokurtosis. If volatility is unpredictable in both systems, we can hardly expect to be able to engineer the model system to replicate the time pattern of volatility of the target system since, to do so, would render the volatility of the target system predictable!

It is nonetheless possible to validate the goodness of the representation of the target system by the model system. The validation must be more direct and expressive than statistical validation techniques. The basis of the validation technique is the implementation of agents to represent specific observable social entities. Such entities could be individual persons or collections of persons as constituent components (departments, sections, or the like) of organisations, whole organisations, government agencies or any other recognisable entity. Validation must be undertaken in collaboration with (or actually by) domain experts who know the behaviour of the target entities and the they interact with other social entities. The key question here is whether such behaviour is plausible to the domain experts who may themselves be the target entities or members of them. In this case, the domain experts are participating stakeholders.

Because every model is sensitive to the values of some parameters, an essential element of model validation is that either the model behaviour is not sensitive to parameter values without unobservable target statistics or the model endogenously drives key values into the range that supports the replication of the behaviour of the target system.

4.2 An example system

In addition to the finding reported above that sales volumes and values by brand have the indicated statistical signature, Moss [18] has reported that market shares by retail outlet are power law distributed when the markets are competitive. The data for three branches of the retail trades is reproduced in Figure 4. The suggestion that market shares might not be found for the less competitive trades is due to the deviation of the market shares from the power law distribution for multiple UK grocers but not (or at least much less markedly) for all grocers. In light of this data, we would expect a model of a competitive market with intermediaries to yield power law distributed market shares.

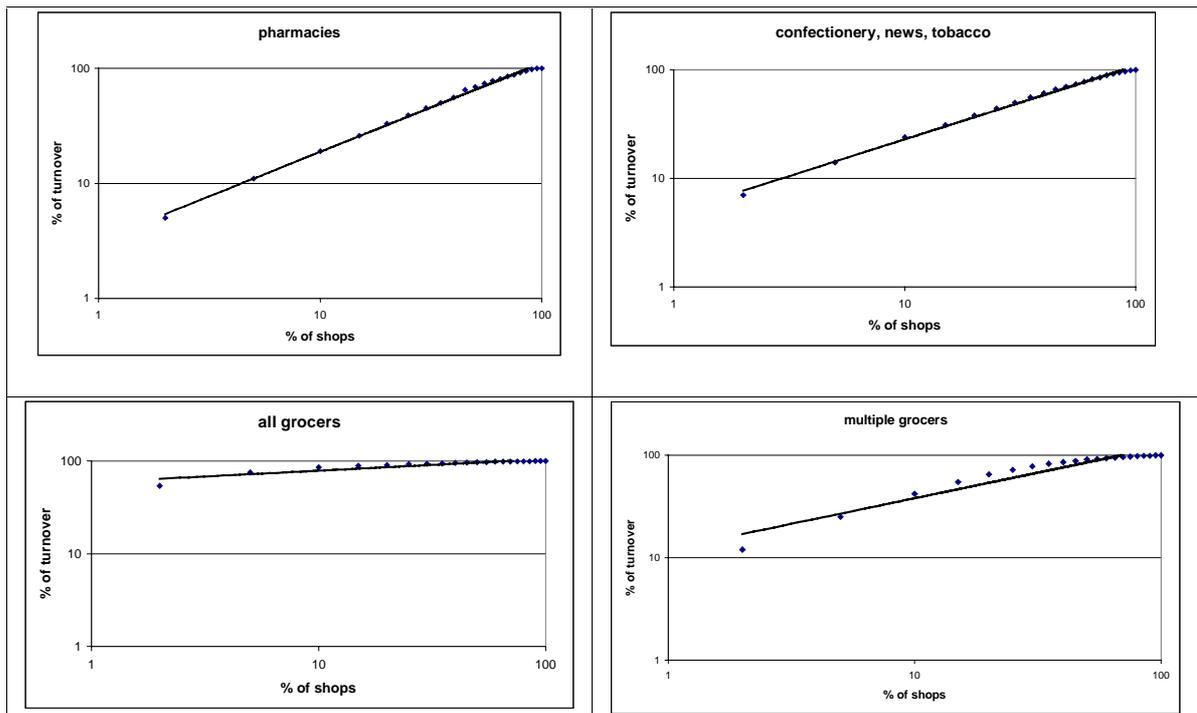


Figure 2: Market share distribution of UK retail outlets [19]

This proposition is tested with a model of a market was implemented in which there are adaptive agents representing customers and adaptive agents representing intermediaries. There are also product sources that are not given any representation of cognition. The social network in this model is represented by a grid with periodic boundaries in which agents can “see” a limited number of cells in each of the four cardinal directions.

Cognitive agents in the model buy and/or sell items represented by the values of digits in an ordered list – a digit string. The values of the digits in the string can be to any arbitrary base. At each trading cycle, an digit string generator produces a digit string. The length of the string is constant over each simulation run.

There is a user-determined number of product sources distributed at random on the grid. Each source holds the current values of digits at specified positions in the digit string. These values change as the system digit string changes.

The intermediaries acquire the values of digits from sources. These values can be acquired only as packets of all items held by a source. However, the intermediaries can sell items individually or in any combinations available to them, selling on to other agents only those items the other agents demand. Moreover, intermediaries can combine the items acquired from several sources. There is a flow of intermediaries chosen at random from the $[1, B]$ interval where B is the maximum number of intermediaries, set by the model operator, that can enter the market at each trading cycle. Each intermediary begins builds asset reserves from profits and leaves the market when its asset reserves are exhausted.

Each intermediary is initially allocated to an empty cell but can choose to move to some other cell if it is unoccupied and no other agent is seeking to move at the same time to the same cell. The motivation to change cells is the knowledge that there is a profitable intermediary in the neighbourhood of the destination cell.

Customers either acquire packets of items from sources in the same way as do the intermediaries or they buy just the items they want from the intermediaries. The customer agents each inhabit a cell during the whole of the simulation run. Although

the number of customer agents is determined at the start of each run by the model operator, their locations are determined at random.

At the start of each simulation run, customer agents are allocated demands for the values of digits at specified positions in the system digit string. The number of items demanded is determined at random in an interval set by the model operator at the start of the simulation run. Intermediaries demand only items for which they have previously received enquiries from customers or other intermediaries.

Intermediaries and customers are synchronous, parallel agents. To enable them to communicate with one another, a series of *communication cycles* is nested within each trading cycle.

In all of the simulation runs, the system digit string contained 40 digits; there were 15 sources and 100 customers. Each customer could demand up to 12 items and each source could hold up to 15 items. The maximum number of broker agents that could enter the market in any trading cycle was 15. Agents could identify the existence of sources or other agents within eight cells of their own position in the cardinal directions (up, down, right and left). The only parameter setting that was changed for the different simulation runs was the size of the grid. Three grid sizes were used: 50×50 (2500 cells), 30×30 (900 cells) and 25×25 (625 cells). A larger grid size implies a lower density of agents.

Experimentation confirmed that agent density is a critical factor in the viability of agent trading, that a high proportion of demands are satisfied only when virtually all trading is intermediated and market shares are leptokurtic.

A natural measure of the effectiveness of markets is the proportion of total customer demands that are satisfied through transactions. The time series of these proportions for three scales of grids are shown in Figure 3. The population density of customers and sources increases from panel (a) down to panel (c) with the corresponding proportion of satisfied demands rising from 3.2 to 14.6 to well over 90 per cent.

Demand satisfaction in all of the modelled markets was a result of intermediated transactions. In Figure 3, the lower line in each case represents acquisitions of items by customers directly from sources. Evidently, in all cases direct acquisition from sources was negligible.

The statistical signature identified for real intermediated markets is replicated by the simulation model. Figure 4 shows that the power law holds for cumulative sales volume against the rank of the broker (from lowest to highest sales) at the 50th trading cycle of the simulation of the 625-cell market. Although the power law distribution prevailed consistently during all trading cycles, the parameters of that distribution were changing over time. The frequency distribution in Figure 5 demonstrates that leptokurtosis of each intermediary's sales volume changes mirrors that of brand data as reported in above.

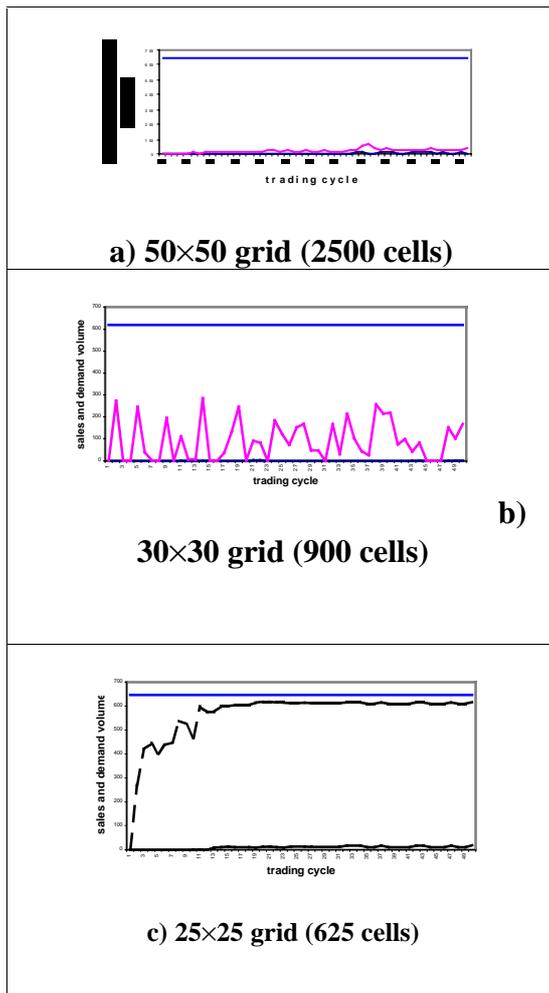


Figure 3: Sales volumes and demands at different agent densities

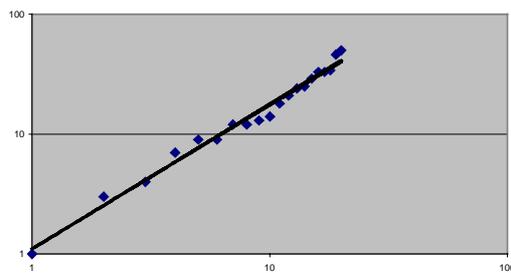


Figure 4: Intermediaries' market share distribution at the 50th trading cycle

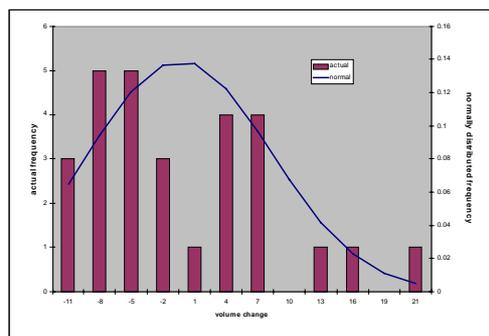


Figure 5: Frequency distribution of simulated sales volume changes – actual and normal

5 Adaptive agent models and the process of policy analysis

The model reported above has been validated to the extent that its statistical signature shares leptokurtosis, clustered episodes of volatility and power law distributed market shares with high frequency data captured from supermarket and other retail sales outlets. This result coheres with the longstanding results on high frequency price and volume data from organised financial markets. The differences in the uses to which items traded in these markets are put suggests that the common factor among these markets considered as data generating mechanisms is that they are competitive and transactions take place through intermediaries – brokers, supermarkets, newsagents, *etc.* The need for competitiveness is consistent with the implication from the simulation experiments that some critical density of agents is required for exchange to be conducted efficiently.

While this result supports a general presumption about the requirements for efficient trading in intermediated markets, it would be rash to build any policy prescriptions thereupon. And one implication of the validation process reported here is that the models cannot in principle be used for purposes of the prediction of events defined by the time of their occurrence and their magnitude – including the prediction of specific outcomes from any policy actions. Indeed, if SOC models support accurate descriptions of social systems as data generating mechanisms, then the sort of prediction sought by physical scientists are *in principle* impossible to achieve in those social systems. Consequently, the use of adaptive agent models where the agents are metastable and their behaviour is influenced by the behaviour of other agents is strictly incompatible with the positivist approaches that justify current social theory.

In general terms, the approach of positivist social scientists is to start from a social theory, derive a specific model from that theory (usually a regression model) and then to apply that model to the data. The approach taken in this paper has been to identify the statistical signature of the data and then to consider alternative means of capturing that data. The advantage of the adaptive agent models is that they can (and the TVP models cannot) be used to describe components of the data generating mechanism in arbitrary detail. Consequently, these models can be treated as tools for describing the data generating mechanism. The validation of these models takes the form of assessing the accuracy of the descriptions they entail.

The result of this process is a model that captures the leptokurtosis and clustered volatility of the relevant empirical data *and* captures the behaviour and interactions of the relevant social actors. This leave two questions: (i) what does it mean to “capture” individual behaviour and interactions? and (ii) how can the models validated in this way be used for policy analysis?

Policy analysis by its nature targets existing social systems. Consequently, there must be stakeholders who are sources of expert information concerning the particular social domains of concern to the policy analysts. Stakeholders and independent domain experts can provide descriptions of the goals and actions of the relevant actors as well as the patterns and modes of interaction among them. They can also evaluate the plausibility of the models designed to incorporate those descriptions in the software code that constitutes each agent. A good agent based model for these purposes will provide information about the agents' goals and behaviour in a form that will enable stakeholders and independent domain experts to evaluate that behaviour as descriptions of actual social entities. The "capture" of individual behaviour and interaction is the design of agents and interaction mechanisms that define software systems (models) that generate system data with the appropriate statistical signatures and produce data about the agents and mechanisms that are validated as accurate or plausible by domain experts.

Stakeholder participation entails not only validation by domain experts but also a more organic process of development of the models in which the stakeholders both explicate and refine their understanding of the target systems and use the models to investigate alternative policy or other strategic options. In the latter case, the models are being used in the same sense as flight simulators to develop responses and abilities to identify the emergence of critical events and a relatively early stage. There is of course an important difference from flight simulators. Flight simulators are based on systems that are well understood and based on clear principals of good physical science and engineering practice that support clear predictions of the outcomes of actions affecting the system. We have seen that systems well described by SOC models are not well understood in this sense. Consequently, stakeholder participation in the modelling process requires that be models are based on good science where that is useful but that stakeholder perceptions take precedence in model design over any conceptual frameworks or system representations developed by independent (*e.g.* academic) observers.

An important prospect here, currently being explored in the EU funded project on Freshwater Integrated Resource Management with Agents (FIRMA) is the development of models by different stakeholders for use in, as one example, negotiation regarding measures for flood safety, water quality, environmental preservation and development, navigation and economic exploitation of the Limberg basin of the River Meuse in the Netherlands. The stakeholders include a ministry of central government, the provincial government of Limberg, NGOs concerned with environmental issues, farmers, community groups, commercial companies and an organisation established to coordinate these various, conflicting interests. The process of participatory agent based social simulation modelling is used to identify conflicts in goals and in perceptions of existing conditions and the consequences of alternative courses of action. It is intended to graft segments of models developed with one set of stakeholders on models representing the understanding of other stakeholders in order to clarify differences and to provide each stakeholder with a greater understanding of the interests and concerns of the other stakeholders.

This sort of process is very different from conventional investigation in the social sciences. Instead of developing a particular model based on some more general construct (theory), the models are devised on the basis of observation and developed by means of a process of empirical validation. Particularly for models incorporating SOC, no hypothesis testing procedures from classical statistics are appropriate and no predictions of particular events are supported. It may be that model development with

stakeholder participation will lead to some more general propositions that can inform social or physical or biological theory. However, the usefulness of agent based social simulation models developed with stakeholder participation is that they support the development of a social process of policy and strategic analysis when forecasting and prediction is infeasible with respect to the relevant natural and social systems.

References

- [1] Moss, S. (2001), "Game Theory: Limitations and an Alternative". *Journal of Artificial Societies and Social Simulation*, <http://jasss.soc.surrey.ac.uk/4/2/2.html>.
- [2] Mandelbrot, B. (1963), "The Variation of Certain Speculative Prices", *Journal of Business*, **36**(4), 394-419.
- [3] Bollerslev, T. (2001), "Financial Econometrics: Past Developments and Future Challenges", *Journal of Econometrics*, **100**(1), 41-51.
- [4] Engle, R.F. (1982), "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation", *Econometrica* **50**, 987-1007.
- [5] Hansen, L.P. (1982), "Large Sample Properties of Generalized Method of Moments Estimators", *Econometrica* **50**, 1029-1054.
- [6] Bollerslev, T. (1986), "Generalized Autoregressive Conditional Heteroskedasticity", *Journal of Econometrics*, **31**, 307-327.
- [7] Bak, P., C. Tang and K. Wiesenfeld (1987), "Self Organized Criticality: An Explanation of 1/f Noise", *Physical Review Letters* (1986), **59** (4), 381-384.
- [8] Jensen, H. (1998), *Self-Organized Criticality: Emergent Complex Behavior in Physical and Biological Systems* (Cambridge: Cambridge University Press).
- [9] Granovetter, M. (1985), "Economic-Action and Social-Structure - The Problem of Embeddedness", *American Journal Of Sociology*, **91**(3), pp. 481-510
- [10] Edmonds, B. (1999), "Capturing Social Embeddedness: A Constructivist Approach", *Adaptive Behavior* **7** (3-4), 323-347.
- [11] Lux, T. (1998), "The Socio-economic Dynamics of Speculative Markets: Interacting Agents, Chaos and the Fat Tails of Return Distribution", *Journal of Economic Behaviour and Organization*, **33**(2), pp. 143-165.
- [12] Bak, P. (1997), *How Nature Works: The Science of Self Organized Criticality* (Oxford, Oxford University Press).
- [13] Kaufman, S.A. (1993), *The Origins of Order* (New York: Oxford University Press).
- [14] Downing, T.E., S. Moss and C. Pahl-Wostl (2000), "Understanding Climate Policy Using Participatory Agent Based Social Simulation" in S. Moss and P. Davidsson (eds), *Multi Agent Based Social Simulation* (Lecture Notes in Artificial Intelligence 1979), 198-213.
- [15] Laird, J.E., A. Newell and P.S. Rosenbloom (1987), "Soar: An Architecture for General Intelligence", *Artificial Intelligence*, **33**(1), 1-64.
- [16] Anderson, J.R. (1993), *Rules of the Mind* (Hillsdale NJ: Lawrence Erlbaum Associates).
- [17] Cohen, P.R. (1985), *Heuristic Reasoning: An Artificial Intelligence Approach* (Boston: Pitman Advanced Publishing Program).
- [18] Moss, S. (2001), "Scalability, Social Simulation and Self Organising Software", in Y. Demazeau (ed.), *Multi-Agent System Organisations* (Lecture Notes in Artificial Intelligence
- [19] Neilsen (1992), *The Retail Pocket Book 1993* (Henley-on-Thames: NTC Publications Ltd.).