

Complexity and Agent Based Models in the Policy Process

Bridget Rosewell¹ and Paul Ormerod²

June 2012

DRAFT

1. Volterra Partners LLP, London browsewell@volterra.co.uk, corresponding author
2. Volterra Partners LLP and University of Durham, UK pormerod@volterra.co.uk

Abstract

Rosewell and Ormerod have collaborated for over a decade in Volterra Partners, building models using complexity principles for clients in both the public and private sectors. Both have extensive experience of policy making, especially Rosewell. For example, she worked directly with the UK finance minister in the 1990s and has been a board member of several major companies over the past decade. She has been the chief economic advisor to the Mayor of London for eight years.

This paper presents examples of agent based models which have been commissioned by policy makers and have been used as inputs into the decision making process.

The main focus of the paper is on complexity in practice. We discuss the problems which arise within both the public and private sectors when a complex modelling approach is proposed, the difficulties which are often encountered, and strategies for overcoming these.

For example, the dominant, standard approach to policy analysis, that of economics, has the advantage in that it often purports to give the answer. Frequently, the models are large black boxes which even the model operators themselves do not fully understand. The policy maker must take the results on trust – a point which actually often works in favour of these models – ‘if it is really difficult to understand, it must be true’

Complex systems modelling in contrast emphasises uncertainty of outcome. In addition, many complex systems models are much more transparent. A problem often encountered is with economists in the relevant institution, who cannot understand that a rule based system is a model at all.

The paper describes the way in which policy makers describe and identify problems, how they can be engaged, how they can ‘buy in’ to the results of a model, how to involve the decision maker in the validation of the model. These points are illustrated from actual models which have been built in practice.

1. Introduction

The policy process is a problem solving one. Policy makers need to produce convincing answers to a problem which seem to give reasonable certainty that a proposed policy will work. Such problem solving requires an intellectual framework which makes sense to both the policy analyst and the decision maker. Over the last half century or so, this framework has been based on a particular paradigm in which models are deterministic and marginal rather than large changes are the main focus of the analytical process

This framework is the framework of economics. The intellectual input into the process of both social and economic policy making in the West continues to be dominated by this discipline. This is not to say that other social sciences do not have an input, but the influence of economics is pervasive.

There are two aspects to this. First, the emphasis in policy is on the role of incentives, especially marginal ones, the fundamental building block of economics. Agents – decision makers – respond and alter their behaviour when the set of incentives which they face changes. So the emphasis is on taxes, subsidies, and the impact of regulation on incentives. Recently, for example, the Scottish executive decided to try and tackle the serious problem of binge drinking by imposing a minimum price at which alcohol can be sold. In transport analysis, time savings are additive and one minute has the same value whether it is the total time saved or part of a larger saving.

The second aspect is the equilibrium framework in which economists typically think. Even if a system is not at equilibrium at a given moment, the basic tendency of the economy, according to this view, is to move back towards equilibrium. So there is a great deal of policy which is based on the concept of ‘market failure’, that there are imperfections in information or other constraints on the behaviour of agents which prevent ‘the market’ from operating as it should and restoring equilibrium. In the context, say, of unemployment, for decades the idea has persisted within economics that many unemployed people are simply choosing leisure rather than work, and a potential solution is to cut the level of benefits to restore equilibrium in the labour market. This is a framework in which the world can be portrayed as deterministic.

The combination of marginal analysis and an equilibrium framework permits the derivation of optimising and optimal criteria. It allows policy makers to conclude that they can design a better world in which everyone can be better off either financially or in welfare terms. This is therefore a very comforting framework which will be hard to abandon unless it is either clearly shown to be misleading or can be replaced by something more effective. It explains why long established results can be set aside or ignored because they do not help decision makers come to a conclusion. For example, Lipsey and Lancaster (1956) showed over fifty years ago that a policy change that makes one market more efficient does not necessarily

improve the overall efficiency or welfare of the system. This result is both well known and ignored in practice.

In the real world decisions have to be made, by businesses, governments, and other organisations. They will rest on the available evidence and the available decision frameworks. Tacit and instinctive positions are just as likely to be important as a careful weighing up of all options¹. This paper looks at how in practice to cope with a world in which uncertainty has become much more endemic and apparent and in which structural change is taking place at a fast pace – but where decisions still need to be made. We consider some of the main sources of error that will lead to bad decisions and present examples of models and modelling framework which address them.

These examples are all variants of agent based models in the sense that agents may exhibit a variety of forms of behaviour and feedback mechanisms are explicitly identified. It is of course the case that the current policy paradigm is based on the simulation of agent behaviour. But this takes place under the restrictive assumptions that all agents are independent, that feedbacks produce equilibrium. The examples described here do not make these assumptions a priori but rather consider the extent to which a different set of assumptions might be relevant to the problem under consideration.

It is crucial in the decision process to remain relevant to the concerns of the policy makers. The extent of added complexity and variety of behaviour must be carefully considered. This leads to the second part of the paper which addresses the challenges of gaining acceptance for a less deterministic approach and of validating it.

2 Complexity, agent based models and the incumbent paradigm in the policy world

Agent based models will generate a range of potential outcomes for any policy which creates uncertainty for the policy maker. This applies across all ranges of policy problems – from the impact of interest rates changes to deciding taxation impacts.

This paper examines a range of policy problems and the kind of frameworks which can address the issues potentially more effectively than deterministic or marginal ones. While a less deterministic framework allows us to consider a wider range of outcomes and policy risks, it also leads to new questions of validation and acceptance. These are as important as establishing new frameworks in the first place. Although many existing models are ‘black boxes’ which are hard to explain or to use, they are nonetheless familiar boxes. A new approach will require new rules of engagement. We address these in the later parts of the paper.

¹ These are decisions based on what Daniel Kahneman (2012) has called System 1 thinking

In a practical context, it is essential for the agent-based, complex systems modelling community to appreciate the power of the incumbent discipline: economics. Keynes wrote as long ago as the 1930s that ‘The ideas of economists...., both when they are right and when they are wrong, are more powerful than is commonly understood. Indeed, the world is ruled by little else. Practical men, who believe themselves to be quite exempt from any intellectual influences, are usually the slaves of some defunct economist’. This remains true today.

In many ways, economics has authority simply *because* it has authority, and has had authority for decades. Decision makers, who in general are *not* experts in any particular discipline, are often predisposed to consider the recommendations based on economic models favourably. They are tried and trusted. Further, economists have the advantage in that they often purport to give *the* answer. Frequently, their models are large black boxes which even the model operators themselves do not fully understand. The policy maker must take the results on trust – a point which actually often works in favour of these models – ‘if it is really difficult to understand, it must be true’.

Economics, over the course of a century or more, has honed and refined its methodological and epistemological approaches to the analysis of practical questions. This does not mean that it correct. But because of this process within the subject, there is much more agreement amongst economists than is commonly supposed. So, for example, different teams of economists might disagree on the impact of, say, a road pricing proposal. But they will agree on how conceptually the problem should be approached.

The issues of methodology and epistemology are not mere ivory tower questions. Addressing such issues is a necessary condition for the use of agent-based, complex systems models to become more widespread in the policy process. Our models *are* different from those of economics. But we need to articulate and explain these principles more clearly, both to show the strengths of the approach and to make clear the disadvantages from the perspective of the decision maker.

The authors of this paper have collaborated for over a decade in Volterra Partners, building models using complexity principles for clients in both the public and private sectors. Both have direct experience of policy making, Rosewell in particular. For example, she worked directly with the UK finance minister in the 1990s and has been a board member of several major companies over the past decade. She has been the chief economic adviser to the Mayor of London for eight years.

We consider here some key methodological issues in agent based modelling, and illustrate them with examples of models which we have been paid to build. Models which have been used as inputs into the decision making process. We describe the way in which policy makers describe and identify problems, how they can be engaged, how they can ‘buy in’ to

the results of a model, how to involve the decision maker in the validation of the model. These points are illustrated from actual models which have been built in practice.

We also reflect on what we might think of as the 'supply-side' of complexity, in other words the experts who build such models. Relatively few experts in the Complexity community have sat down and listened to the challenges policy makers face. We suggest ways in which behaviour needs to be adapted from the supply side.

3 Extending the paradigm

This section looks at ways in which the decision frameworks need to be extended and examples in which agent based models have been used to do this.

3.1 Imperfect information and knowledge

There is a key distinction to be made between the concept of imperfect information and that of imperfect knowledge. Imperfect information is associated with the concept of bounded rationality, and developed within economics by Akerlof (1970) and Stiglitz (1977). Bounded rationality in this sense implies that the decision making process follows the same rules as when perfect information is available but includes an assessment of the costs of acquiring information. Information will exist but if its capacity to improve a decision is less than its cost it will be ignored. Equally some agents may have better access to information than others and are able to create false market signals.

Imperfect knowledge is different from this and can be regarded as being in the tradition of Hayek and Simon. In many situations, there are *inherent* limits to knowledge, which no amount of additional information can overcome (Hayek, 1974 for example). The same concept underpins Simon's seminal 1955 article on behavioural economics. Economists now interpret his phrase 'satisficing' to mean that agents look at a limited number of alternatives until they find one that is satisfactory, judging that a superior one might very well be found, but that the time and effort involved in the discovery of it would outweigh the additional benefits, in other words in the context of imperfect information. On the contrary, Simon argued that, in many situations, all that we are able to do is to find a satisfactory choice or make a satisfactory decision. Even *ex post*, the optimal choice or decision cannot be discovered.

It is not surprising that analysts shy away from the proposition that the right answer is unknowable. An analyst's instinct is that there is a right answer and more research is needed to acquire the right information to uncover it. The search for a right answer is congruent with the approach of economists that market solutions have optimal properties and that a world which fits the assumptions behind such solutions will also have optimal properties. Thus the policy world puts together a belief that somewhere full information

exists with a set of axioms that shows the optimal properties of a world with such information.

It is therefore very important to describe the consequences of imperfect – and inherently imperfect – knowledge. This is partly the difference between foresight and hindsight, and partly an essential randomness born of computational limitations. The Argentinian poet Jorge Luis Borges' short story, *The Garden of Forking Paths* (1941), is now widely discussed in academic literature as the seminal example of modern digital media and hypertext projects. Ostensibly a spy story, Borges conceives of 'a labyrinth that folds back upon itself in infinite regression', asking the reader to 'become aware of all the possible choices we might make'. It is impossible either completely to explore the garden or to map it. Making a choice of paths itself determines future options, while backtracking does not return you to the previous starting place.

Thus the story that describes the actual history which was experienced also provides the data on which a prescription for the future is based. Yet that history did not look the same at the outset, and there will be several possible futures that can be modelled even based on the available history. In the Borges story, there are innumerable possible futures, and by definition not all of these can be modelled or described in the time available to make a decision. In the physical world, the repetition of consequences from causes can make it possible to derive laws which enable hindsight – data – to create foresight. In the policy world such repetition is often lacking and each repetition of a policy takes effect in distinctively different circumstances. Such differences reduce the possibility of systematically accurate prediction and generate multiple possible futures.

Both companies and governments operate in situations which exhibit not only imperfect information but also imperfect knowledge. The problems associated with the former can be overcome in principle, while the latter cannot. They are an inherent feature of the real world. Essentially, if uncertainty did not exist (or if the environment was sufficiently stable for decision makers to be able to learn over time in ways which reduced uncertainty), running a company would be easy. But it is not.

We need both to understand the limits on our information and the limits on our knowledge. Often these are confused. In policy making, data can be assumed to be correct when there are uncertainties and significant scope for revisions over time. Data is collected for one purpose and used for another. For example, unemployment was for a long period reported based on administrative data on those in receipt of unemployment benefit. More recently, a survey has been designed to collect information on those who are looking for work, rather than just those in receipt of benefits. This is a more appropriate definition, but the result is based on a sample survey rather than a complete measurement. Which is the more accurate or appropriate depends on the purpose for which the information is needed. This in turn depends on other kinds of information. Statistical data is usually only part of

what is known and used in decision making. Other information, both explicit and tacit, is used to select the data to be collected, and to decide on how appropriate it is.

Data and its availability provide the building blocks of the modelling strategies that we might use. The data may itself be constructed within the framework of such a strategy. Much macro-economic policy relies on a measure of the output of the economy, often Gross Value Added (GVA), or Gross Domestic Product (GDP). These measures are widely used and reported, and policy is based on their growth and the difference between such measures and potential output – the output gap.

These concepts themselves rest on a framework of analysis which divides up the economy into different sources of income and output, and a statistical account of how the data are collected shows how great the uncertainty can be². One of the present authors, when advising government in the mid-1990s as a member of a Group of Independent Economic Forecasters otherwise known as the Seven Wise Men³, participated in a study⁴ to look at estimates of the output gap. She concluded that the data uncertainties were such that this was a wildly inaccurate guide to policy. It continues however to be used on the basis of the modelling framework which underpins it. It is of course true that we make decisions based on inadequate information and guesstimates at the real position. It is essential that we do not confuse inadequate data with an inadequate modelling framework or an inadequate application of our frameworks.

3.2 Stochastic and Deterministic Models

Imperfect knowledge can derive from both inadequate frameworks for models and from an inadequate application of models. One way to develop the application of our models is to ensure that stochastic elements are used to capture variations of behaviour outside the core model. The stochastic terms reflect the situation of imperfect knowledge in which we operate. We can identify at least four sources of imperfect knowledge, of uncertainty in the Knightian sense.

These are uncertainty about the starting point (the initial conditions), the order in which agents take decisions, the network which connects them and the behavioural rules that they follow. The first of these is not just about the accuracy of data. Depending on the starting point, a system can follow a very different path, either towards or away from an

² Maurice (1968) is the classic account of the compilation of national accounts data

³ The group only had male members at its inception, though latterly two women were members. The nickname however survived.

⁴ The Independent group of Forecasters, (1996).

equilibrium. In many cases, only a small change in the starting point can make a big difference⁵.

The other three elements all describe potential ways in which feedbacks can operate and behaviours diverge from the standard view of how decisions are taken. We look at examples where generating different modelling frameworks delivered more effective results and a different policy outcome.

3.2.1 Initial conditions

In 2011, we were commissioned by a financial institution to model whether the financial situation in Spain was stable or unstable. Kosko (1986) and Dickerson and Kosko (1994) develop the concept of a fuzzy cognitive map, or FCM. This is a concept which draws a causal picture between the key factors in any given situation. In particular, it allows for dynamic feedback between these factors. In essence, it is a system of differential equations similar to those used in the systems dynamics literature.

In conjunction with the financial sector institution, we identified a small number of key determinants of stability: unemployment; bad debt of savings banks; regional/local government bad debt; spread on government bonds; international sentiment; property prices; economic growth. We model these as a dynamic system with feedbacks between the various factors.

The client suggested the strength of the impacts of the various factors on each other, and after several iterations of discussion with them, a central view of their values was arrived at. The client recognised that uncertainty existed in the judgment which led to these values, and each parameter was drawn from a uniform distribution around its central value.

We then solved the model a large number of times, and investigated its stability properties with reference both to the numerical simulations and to the eigenvalues of the matrix of partial derivatives of the system. The eigenvalues are a mixture of positive and negative real parts and so the initial values of the variables are crucial to the stability or otherwise of the system. But there is considerable uncertainty about the initial conditions at any point in time. Some of these variables, in particular international sentiment, are hard to measure in any precise way.

3.2.2 The order in which agents take decisions

We can readily solve a model many times and create a large number of potential outcomes. But in the real world, there is only ever one solution. Historical counterfactuals, for

⁵ Another early and frequently ignored result illustrates this. It is known as the hog cycle or cobweb theorem, Kaldor (1938). This used to be in economics textbooks but has been purged in recent decades. It is a simple model of how the price mechanism can lead to *disequilibrium*

example, are both intellectually challenging and engaging, but as discussed earlier there has only ever been one actual path of history.

There is a considerable literature on this. In models in which choice, whether of a product or a technology, is governed not just by the set of incentives described by conventional consumer demand theory, but by the choices of others in which an individual's payoff is an explicit function of the actions of others, the alternative(s) which become successful often owe their success not to their inherent attributes but to the fact that they were selected early and at random. Examples of this principle are given, for example, by Arthur (1989), Brock and Durlauf (2001), De Vany and Walls (1996), Ormerod (1998) and Bentley et al. (2011).

A practical example was a model we built for British Telecom some ten years ago⁶, who were involved in discussions with the regulator on their market share of the landline market in the UK. BT had been a monopoly supplier, and the company continued to maintain a large market share despite the opening upon of the market to new entrants and hence to competition. On standard criteria, the market share was anti-competitive.

Our model focused on the process by which new entrants came into a market opened up to competition for the first time. The monopolist provides a combination of price and quality $[p, q]$, which are, respectively, high and poor. A firm is chosen to enter with a random combination of $[p, q]$, with the proviso that both are better than that initially provided by the monopolist. Not all consumers will be aware of this new entrant, but those who are have the opportunity to switch. There are rules specifying when they will do this, how the monopolist can respond to the newly offered $[p, q]$, and so on. For example, in subsequent periods, firms enter the market depending upon their assessment of potential profit, which in turn depends upon the $[p, q]$ combination which has become established at that point in time.

In general, the market eventually contains an indeterminate number of producers, each of which supplies the product either at or close to the minimum price which is both technologically feasible and enables a normal rate of profit to be made. In other words, we begin with a monopoly and end with the perfectly competitive market of standard economic theory.

However, judged on the conventional criterion of the distribution of market shares, the market structure which emerges is, in general, anti-competitive. As the outcome on market price shows, the model is highly competitive in any meaningful sense of the word. The market price which emerges, given the number of firms in the market in any particular solution of the model, is in general lower than that implied by the equilibrium Cournot

⁶ Ormerod and Rosewell (2001)

price, substantially so when the number of firms is small. The model further differs from conventional theory in that, across a large number of individual solutions of the model, there is no connection between the market price and the number of firms.

The potential range of outcomes for the model is wide, reflecting the importance of contingency in the process of the evolution of market structure. In determining the outcome in any particular solution, the $[p, q]$ offer of the early entrants, and the extent to which they generate awareness amongst consumers, is crucial. The order in which agents take decisions has a strong influence on the solution.

3.2.3 Structure of network

The differences between the percolation properties of different kinds of networks are well known (see for example, Newman, 2010). In practice, a good approximation can often be made to the type of network which connects the relevant agents by combining straightforward survey techniques with agent based models (for example, Ormerod 2007 and Ormerod and Wiltshire 2009). However, an inherent level of uncertainty remains.

An example where the connectivity of the network was crucial to the properties of the model is that of Svenska Handelsbanken (SH), Sweden's largest bank. For reasons of commercial confidentiality, the full results of this work are not in the public domain, but a description of the key features can be given.

In the early 1990s, Sweden experienced a recession which was almost as bad as that of the Great Depression period of the 1930s. Defaults on corporate loans rose sharply as a result. The default rate of SH was, however, considerably lower than that of its main competitors.

The bank wanted to know, essentially, whether this was due to pure good fortune given that the bank was much less exposed in the sectors where defaults were particularly high, or whether it was due to their own process of granting and monitoring of corporate loans.

The system at SH had evolved over many years. In contrast to the very centralised, tick-box driven approach of most banks, at SH the officer at the branch level is able to exercise considerable judgement on both the granting and the monitoring of corporate loans. A set of guidelines does exist, but they are not rigid. The officer is informed in part by his or her private information and opinion. But an additional important influence is the informal network across the whole of the bank through which knowledge, both explicit and tacit, about an industry, a firm or even the individual asking for the loan, is exchanged.

The bank not only wanted to know the answer for itself, but it was under pressure to adopt a process more in keeping with the industry standard, a much more centralised approach with little scope for individual judgement.

We developed an agent based model for the processes of granting and monitoring corporate loans. The structure of the network was approximated using the evidence from a survey designed with the assistance of the bank. The average path length of the network was low, indicating that knowledge flowed very readily across it.

We ran simulations of the model with both the approximation to the actual network and with networks with distinctly higher path lengths, even removing the network altogether in one set of solutions. The conclusion was that the network, and the efficiency with which knowledge could be transmitted across it, was an important reason for SH's good comparative performance in the early 1990s. By coincidence, we presented these findings to board members in the summer of 2008. In the recession which followed almost immediately, SH once again had a good comparative performance on corporate loan defaults.

This suggests that we obtained a reasonable approximation to the actual network and that it was important. However, an inherent amount of uncertainty remains as to the structure of this network (for example, did participants in the survey give answers which they felt would be the desired ones? The culture of the bank placed emphasis on the informal network, and they may have responded accordingly).

3.2.4 Behavioural rules

There are many examples of uncertainty relating to behavioural rules. Even the most senior executives are aware of their lack of knowledge about what are often key aspects of the business. This does not mean that they do not try to become more knowledgeable, simply that at any given point in time there are some things which they perceive 'through a glass darkly'. A standard response in economics would be that over time they would learn. But in practice, this is often not the case at all. The environment in which a firm operates changes too rapidly for systematic learning to be able to reduce uncertainty.

An ongoing example is the water industry in the UK, which we are modelling with a consortium of firms and the regulator, Ofwat. Much of the regulation is focused upon increasing the quality – measured in many dimensions – of the product.

The amount of uncertainty varies across activities. But in general, the companies do not know for certain ex ante how much increase in quality will result from a given level of capital expenditure. Further, although they believe that the overall function between

quality increases and capital spend exhibits diminishing returns, they do not know for certain either its exact shape, or precisely where they are on any given curve (the uncertainty of initial conditions once again).

Overall, therefore, we conclude that stochastic elements are an inherent feature of almost all models of corporate behaviour. It is a delusion to imagine that these can be modelled deterministically.

3.3 Dynamics

A consequence of the equilibrium paradigm is that much analysis compares a current situation (or a forecast one) as an equilibrium with another potential equilibrium after a policy change, an investment, or a 'shock'. This has been particularly prevalent in macro-economics, which finds it much easier to compare equilibrium outcomes when changes have worked through.

For example, the mainstream paradigm in academic macroeconomics is DSGE, or dynamic stochastic general equilibrium models to give the full description (for example, Smets and Wouter 2003, Tovar 2009). There are many critiques of such models, and two from an agent-based complexity perspective are Farmer et al. (2012) and Ormerod and Helbing (2012). The great difficulties such models experienced in not just forecasting but even explaining the financial crisis has not deterred the Central Bank of Iceland – and Iceland was one the countries most affected by the crisis – from deciding to develop a new DSGE model, which it describes as being 'more suitable for policy simulations' than rival models (<http://www.sedlabanki.is/?PageID=722>)

The reason that such models seem to more suitable for policy simulation is that their equilibrium properties guarantee that results can be produced. This is however putting the cart before the horse as it certainly does not guarantee that results are realistic.

Another example is the use of cost benefit analysis. In principle the idea that we make a decision based on the balance between costs and benefits is an eminently sensible one. Who could disagree with this? The questions arise when considering how to value them, especially in the case of the benefits of a proposal. A recent example of the difficulties of comparing future benefits with current costs arose in the context of the Stern report (20xx) into dealing with climate change. His calculations took a very low discount rate so future generations had a high value, while also forecasting that economic growth would continue in the world economy. The result, as several economists pointed out (for example Nordhaus 20xx), was to impose costs on poor people now to benefit rich people several generations hence, while also assuming that none of these future people would act to avoid the negative consequences, by moving away from flood plains for example.

The dynamics of feedbacks can also make a potential difference to policy outcomes. The concept of agglomeration goes back to Alfred Marshall in the 1890s and his description of how growing locations could create more effective labour markets, spillover effects and innovation. The concept is essentially dynamic as each addition to the central area creates a further benefit not just to those joining the agglomeration but to those already there. There is no sensible answer to the question of the equilibrium size of an agglomeration.

These concepts were used to support the case for a new metro line in London, Crossrail. Adding the dynamics of agglomeration made a significant difference to how it was possible to think about the benefits of the expenditure. However, the modelling had to be done in an essentially equilibrium context in order to gain acceptance, and various authors (for example, Rosental and Strange xxxx) have estimated agglomeration elasticities which essentially put this dynamic process into a marginal context.

Indeed it is possible to use a maximising framework to describe the benefits of agglomeration in which agents trade off the benefits of co-location against the disbenefits of congestion and distances to be travelled. This is the framework developed by, for example, Fujita, Krugman and Venables (1999) However, it makes more sense to think of the process of agglomeration as a dynamic and uncertain one in which the timeframe over which businesses and organisations engage with each other to create new opportunities is a more effective way to understand and to model what is happening.

Another example of such a model was constructed for the Manchester Independent Economic Review in the UK. This model looked at innovation as a dynamic and disequilibrating process which was dependent on the ability of innovators to pass on their innovation. A detailed description is in Ormerod and Rosewell (2011)

The model takes an initial innovation to be exogenous and it is taken up by one agent/organisation at the outset. The characteristics of the agents are governed by their willingness to innovate, their desire to keep innovation to themselves, and their willingness to communicate with others. We define two different methods by which innovation may be passed on via the network linkages. The first is a direct relationship between two partners, while the second is a group relationship.

First, an organisation with an innovation will provide it to another firm only if its level of secrecy, or the propensity of a firm to try to retain the benefits of its innovations, is less than the absorptive capacity, or the degree to which a firm actively engages in activities which enable it to identify and adopt new innovations, of the firms it is linked with. This method of adopting an innovation represents a mutual relationship or exchange between

firms and implies a degree of trust or collaboration. It is probable that this relationship is more likely to exist with customers, suppliers or third parties than with competitors.

The second method for spreading an innovation we describe as a copying behaviour. Here if a firm looks at the spectrum of organisations to which it is linked and finds that the proportion that have adopted an innovation is higher than their own personal threshold, they will mimic their behaviour and adopt the innovation

The network behaviour was not set *a priori*. We were able to calibrate the network across which any innovation percolates to the actual networks in a survey.

This model offers a way of thinking about innovation based on dissemination across a network where the behavioural rules of innovating organisations also have a bearing on the outcome. It enabled us to analyse the potential for cascades of innovation and features likely to limit such potential. The model shows that the distribution of parameters across groups of agents matters as well as averages. The ability to calibrate these distributions to survey results in a region of the UK enables us to investigate the importance of the distributions as well as the average values of the parameters.

The results showed that the ability to distinguish between a successful, global, cascade of innovation at the outset is very limited. Whether a cascade emerges or not is very difficult at the outset. We can see that industry surveys with similar results for innovation behaviour and attitudes can nonetheless produce very variable outcomes. Moreover, initial success cannot predict the final outcome.

Classically, policy has concentrated on identifying innovation that ought to be supported, and on creating sector support groups. In the UK, the latest policy initiative has been the establishment of NESTA (National Endowment for Science, Technology and the Arts). A recent pronouncement suggested setting up new sector groups, to look at different ways innovation takes place in 'non-traditional' sectors. How this is likely to help change the parameters of the kinds of behaviour which generate innovation take-up is not at all clear.

However, the models do suggest that there are ways to help maximise the likelihood of innovation spreading. The model shows that:

- The willingness to exchange innovation is a major driver
- Well-connected but not too dense networks help
- A willingness to copy others helps

Policy can help create networks by events and the support of the widest possible networks. Note that these need to be more broadly focused than on just one sector. Policy to support the willingness to absorb innovation is much more difficult. The policy instinct is to provide grants or tax breaks. However, there is no guarantee these would support the sort of innovation that organisations can readily use, which may be hard to identify in policy terms.

The conclusion can only be that a continual policy of supporting test and try which does not worry too much about an ex ante analysis of the potential for success is the only option.

The examples outlined above have illustrated how policy conclusions can be extended and developed using agent based simulations which have practical consequences. However, this is not enough to establish the acceptability of the conclusions reached and moreover some models describe the limits to policy rather than its effectiveness. The next section examines some of these challenges in the context of some further examples.

4 Challenges of policy using agent based models

4.1 Defining the question

In most jurisdictions there is a developed set of rules for defining and analysing policy questions. As a policy maker once said to one of the authors, we have been doing it this way for thirty years, so how can it be wrong?

Innovation in policy analysis is as likely to follow a policy choice that needs support as to follow a matter of principle. For example, the development of agglomeration as an approach to evaluation of metro railways followed the political policy decision that such a railway should be supported. Indeed it is possible that such 'gut' or instinctive responses might implicitly include more model parameters than the formal modelling is able to do.

The challenge was then to formulate the policy question in such a way that it was capable of generating a new modelling framework that could in turn be implemented.

A similar challenge was presented by a different transport investment, in this case a road bridge across the Thames in London. The policy question was to estimate the economic benefits of this bridge. A potential methodology is to consider the economic benefits of other similar bridges. Two serious problems emerged with such an approach. First, there were few bridge investments to consider and each had very specific contexts. Several were rural, one had been built for overtly political reasons, others had international contexts. No case study controlled for enough features to compare with an urban bridge in a large city and there too few cases for a statistical analysis.

Within the city, transport models have traditionally been produced on a cross sectional basis. Although models are updated, no effort is made to compare data across time and to produce it coded in a way which makes this possible. Thus the research question became one of asking what role accessibility played in generating economic and residential activity. The agents in such a model became the locations responding both to available accessibility

and to all the other multifarious factors which might affect them. A fuzzy clustering method was used to focus on the role of accessibility, while implicitly controlling for the unknown and unmeasured other effects. Thus the research method and the policy question became themselves affected by the availability of data which could be used to address them.

4.2 Parsimonious or large models?

This is a major issue on which there is already a considerable literature, and the length of the discussion here does not reflect its inherent importance. Our concern is to set the debate in the context of the need to deal with the phalanx of mainstream economic opinion in persuading decision makers that complexity-based agent based models should be used to help inform decisions.

We obviously cannot say that the absolute size of a parsimonious model has to lie below a certain number of parameters, this will be context dependent. As Dibble (2006) states: 'parsimonious modelling means selecting the simplest possible model capable of generating a phenomenon of interest'.

However, the presumption is that a parsimonious model will contain a fairly small number of parameters. An important advantage of this, given the time and resource constraints under which models have to be built in many practical decision making contexts, is that a model can be developed within the required timeframe. Further, we can explore its properties more easily than is the case with larger models.

There is, however, often an important tension in any practical context. Decision makers often believe that their particular problem can only be understood by knowing a great deal about the issue, and by taking into account a large number of factors. This is particularly the case with what we might term senior middle management, people who have spent their working lives in the industry or policy area and who will probably never make the board themselves. At the highest level, there is often much more awareness of the need to abstract key issues. Nevertheless, in most consultative processes where knowledge about a particular problem is extracted, the tendency for non-model builders is usually to want to include more detail rather than less.

4.3 Validation

A key issue relates to the validation of agent-based models⁷. A criticism frequently encountered from economists in organisations is along the lines 'this is all very well, but how do you know the model is correct? Surely there are many sets of rules which would give similar results?' Our response to the latter is usually to tell them that if it is so easy,

⁷ The current authors discuss this at greater length in Ormerod and Rosewell (2009)

they should go away and produce some other sets of rules, which buys silence, if not agreement.

But they have a point. Econometrics for example, for all its limitations, has a well developed methodological process of model validation. It has inherent limitations, in that a regression is simply plane-fitting in n dimensions and is not readily generalisable. Further, in practice most applied econometricians violate the basic precepts of statistical theory. In general, the power of statistical tests is known if a hypothesis is formulated and then confronted with the data. What happens in practice is that if a specification fails a test, the econometrician will re-specify until he or she gets a 'model' which passes the test. But by then the power of the test is theoretically unknown, although the strong presumption is that it is much weaker than that which would obtain if the precepts of classical statistical theory had been followed. However, econometricians believe that they have such a procedure, which informs them as to whether the model has been validated or not.

Agent-based modelling lacks basic tests which would indicate, for example, whether or not a model is over-fitted. Complex systems science lacks concepts such as the Akaike Information Criterion to penalise over-fitting.

A little known paper by Osgood (2006) addresses this question. This uses techniques on estimating dimensionality taking from the dynamical systems literature. We suggest that methodologies along these lines would be very useful in both helping to arbitrate on model size, and on developing a robust methodology for the validation of our models. This will obviously be different to that of, say, econometrics, but we need one in place so that we can confront the repeated objections raised on the process of validation of our models.

4.4 Prediction and Explanation

In much policy making the line is blurred between explanation and prediction. Most policy questions involve asking how a future world will be different if a particular policy is adopted. This actually involves three key propositions:

- Can I explain the current position in this policy area?
- Can I be confident how it will evolve if left alone?
- Can I know how it will change if parameters change?

These propositions are theoretically distinct. For millennia there have been models which have predicted the movements of planets and stars fairly accurately, but which explained them in terms of mystery and religious forces. Indeed a current model of long term future employment levels in London has been on average accurate for the last decade but has

essentially zero explanatory power, relying on trends in population, productivity and the economy (GLA Economics, 2011).

Macro-economic models used to explain the past evolution of the economy will link outputs such as investment and consumer spending to other elements of the model and finally to variables outside the model – such as population. So what will drive the future evolution is the behaviour of these exogenous variables, precisely what the forecaster knows least about.

Using any model to ask the ‘what if’ question requires a good understanding of what is being held constant. A model can be run with, for example, a different oil price. Is it reasonable to leave all else unchanged? What reactions might there be? Would other behaviours change? Cross section models, such as those used in transport or to study input output relationships, are very risky if used to examine large changes. And since they are expensive to build and maintain, this is what they are used for. Yet a large change in transport availability will probably change other relationships too. A large investment on a particular site will encourage other activities, and may well force or encourage innovation which will change supply chain relationships.

Economists are beloved of the assumption of ‘ceteris paribus’ – other things being equal. Too often they do not stop to consider what other things will not be equal in a particular case. The advantage of agent based modelling is that it does force the analyst to consider such questions. Which behavioural rules are relevant and which agents affect which others become highly charged questions which need considerable thought. To hold all else equal is both a lazy way out and requires making other, albeit implicit, assumptions.

5 Conclusion - Acceptance and Relevance

In making policy, there is a need for a shared set of assumptions. If every problem is built afresh from the ground up, it will be slow and hard to make any decisions. Policy makers are impatient.

The current paradigm is one where changes can be assessed in a common framework based on the principles of agents which maximise and which learn effectively. Prima facie, this appears to make sense.

Yet in practice, small changes can be so small that agents cannot react to that which is too small to see, while large changes may well change the parameters of their choices. Learning in a changing world is hard and so any form of optimisation is beyond most agents’ scope. Many policy makers would recognise such realities in a description. Nonetheless developing a new and shared set of rules which allows us to approach such realities must be recognised

as a hard problem which we are only just beginning to address. How a question is defined, what sort of model should be used to address it, the standards of validation, and what is being held constant are all hard questions to address.

Nevertheless, progress is being made. Approaches to policy problems which recognise variants on agent behaviour, or the importance of dynamics, or feedbacks are increasingly being used. A good example is the study of regulated industries. Traditional approaches have assumed that marginal incentives will be evaluated by rational, profit maximising firms. Increasingly, behavioural approaches are making sense of how different kinds of firms are motivated and constrained. Our own work in the water industry has identified different motivations depending on the attitude to risk (constrained by financial resource), and to regulatory pressure (constrained by an attitude to reputation). Behavioural economics is adding to this with increasingly definitive conclusions on the ability of agents to recognise and react to incentives.

The agent based modelling community need to recognise these steps and work to bed them in and strengthen them with an evolving set of rules of engagement.

Acknowledgements

The authors acknowledge the support of the EU Co-ordinated Action Non-Equilibrium Social Science (NESS), part of the Global Systems Science initiative, and of the FP7 ICT project Open Collaboration for Policy Modelling (OCOPOMO).

References

- R.G. Lipsey and K. Lancaster (1956-1957), 'The General Theory of Second Best', *Review of Economic Studies*, 24(1), 11–32
- D. Kahneman (2012), *Thinking, Fast and Slow*, Penguin
- G. Akerlof (1970), 'The Market for "Lemons": Quality Uncertainty and the Market Mechanism', *Quarterly Journal of Economics*, 84, 488–500
- J. Stiglitz and S. Salop (1977), 'Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion', *Review of Economic Studies*, 44, 493-510
- F.A. Hayek (1974), 'The Pretence of Knowledge', Nobel Prize lecture, (http://www.nobelprize.org/nobel_prizes/economics/laureates/1974/hayek-lecture.html)
- H.A. Simon (1955), 'A Behavioral Model of Rational Choice', *Quarterly Journal of Economics*, 69, 99-118
- J.L. Borges (1941), *El Jardín de senderos que se bifurcan*
- R. Maurice (1968), *National Accounts: Sources and Methods*, HMSO, London
- N. Kaldor (1938), 'The Cobweb Theorem', *Quarterly Journal of Economics*, 52, 255–280
- B. Kosko (1986), 'Fuzzy cognitive maps', *International Journal of Man-Machine Studies*, 24, 65-74
- J.A. Dickerson and B. Kosko (1994), 'Virtual worlds as fuzzy cognitive maps', *Presence*, 3, 173-189

- W.B. Arthur (1989), 'Competing technologies increasing returns and lock-in by historical events', *Economic Journal*, 99, 116–131
- W.A. Brock and S.N. Durlauf (2001), 'Discrete Choice with Social Interactions', *Review of Economic Studies*, 68, 235–60
- A. De Vany and W.D. Walls (1996), 'Bose-Einstein dynamics and adaptive contracting in the motion picture industry', *Economic Journal*, 106, 1493–1514
- P. Ormerod (1998), *Butterfly Economics*, Faber and Faber, London
- R.A. Bentley, P. Ormerod, and M. Batty (2011), 'Evolving social influence in large populations', *Behavioural Ecology and Sociobiology*, 65, 537–546
- M.E.J. Newman (2010), *Networks: an Introduction*, Oxford University Press
- P. Ormerod (2007), 'Extracting deep knowledge from limited information in an evolving complex system', *Physica A*, 378, 48–52
- P. Ormerod and G. Wiltshire (2009), 'Binge drinking the UK: a social network phenomenon', *Mind and Society*, 8, 135–152
- P. Ormerod and B.C. Rosewell, (2001) *The Evolution of Competition*, Volterra Consulting for British Telecom
- F. Smets and R. Wouters (2003), 'An estimated dynamic stochastic general equilibrium model of the euro area', *Journal of the European Economic Association*, 5, 1123–1175
- C.E. Tovar (2009), 'DSGE models and central banks', *Economics E-Journal*, 3, <http://dx.doi.org/10.5018/economics-ejournal.ja.2009-16>
- J.D. Farmer, M. Gallegati, C. Hommes, A.P. Kirman,, P. Ormerod, S. Cincotti, A. Sanchez and D. Helbing, (2012, forthcoming), 'A complex systems approach to constructing better models for managing financial markets and the economy', *European Physics Journal*
- P. Ormerod and D. Helbing (2012, forthcoming), 'Back to the drawing board for macroeconomics' in the *Proceedings of the Bank of England/UK Government Economic Service Workshop on: What Should Graduate Economists Know?*, February 2012
- N. Stern (2006) *The economics of climate change*, HMG
- W.D. Nordhaus, (2007) 'A review of the Stern review on the economics of climate change', *Journal of Economic Literature*, Vol XLV, Sept 2007, pp686–702
- A. Marshall (1890), *Principles of Economics*, Macmillan
- S. Rosenthal and W. Strange, (2003) '[Evidence on the Nature and Sources of Agglomeration Economies](#)', *Handbook of Urban and Regional Economics*, J.V. Henderson and J.-F. Thisse, eds,
- M. Fujita, P. Krugman and A. Venables, (1999), *The Spatial Economy*, MIT

P.Ormerod, B Rosewell and G Wiltshire (2011), Network models of innovation processes and policy implications, in C Antonelli, ed., *Handbook on the Economic Complexity of Technological Change*, Edward Elgar

C. Dibble (2006), 'Computational Laboratories for Spatial Agent-based Models, in L. Tesfatsion and

K.L. Judd, eds., *Handbook of Computational Economics: Agent-Based Computational Economics, vol II*, Elsevier

P. Ormerod and B. Rosewell (2009), 'Validation and Verification of Agent-Based Models in the Social Sciences', in F.Squazzoni, ed., *Epistemological Aspects of Computer Simulation in the Social Sciences - Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, 5466, 130-140, doi: 10.1007/978-3-642-01109-2 _10

N Osgood, (2006) 'Low-Dimensional Dynamics in Agent-Based Models', *System Dynamics*,

GLA Economics, (2011) Working Paper 51, Employment projections for London by sector and trend-based projections by borough