Verification and Validation of Agent Based Models in the Social Sciences

Paul Ormerod and Bridget Rosewell, Volterra Consulting, London

pormerod@volterra.co.uk

Abstract

This paper considers some of the difficulties in establishing verification and validation of agent based models. Simulation based results limit the ability to verify and blur the boundary between verification and validation. We suggest that a clear description of the phenomena to be explained by the model and testing for the simplest possible agent rules of behaviour are key to the successful validation of ABMs and will provide the strongest base for enabling model comparison and acceptance. In particular, the empirical evidence that in general agents act intuitively rather than rationally is now strong. This implies that models which assign high levels of cognition to their agents require particularly strong justification if they are to be considered valid.

"Where do correct ideas come from? Do they drop from the skies? No. Are they innate in the mind? No. They come from social practice, and from it alone; they come from three kinds of social practice, the struggle for production, the class struggle and scientific experiment" Mao Tse Tung, 1963

Introduction

This last kind of 'social practice', namely scientific experiment, is one of the few areas where the tradition of English empiricism blends with the more dogmatic nature of Marxist thought. In this paper, we argue that the best way to make progress in the validation of ABMs involves two clearly defined stages. The first is to construct correct criteria by which the output of a model is to be judged. The second is by exploration, by trying various approaches and seeing what works in practice to best meet these criteria. In this context, theory is led by practice. Abstract musing is of little or no value.

Agent based models (ABMs) in the social sciences are relatively new. In general they rely on numerical simulation to discover their properties instead of the more traditional route, within economics at least if not the other social sciences, of analytical solution. And numerical simulation in turn relies on the massively enhanced power of the personal computer. In short, ABMs are essentially a recent innovation which has only become feasible to implement over the past ten years or so.

For this reason, no firm conclusions have been reached on the appropriate way to verify or validate such models. We are in a period of rapid evolutionary change following an innovation. The innovation is ABMs, and the evolutionary process which is taking place involves a wide range of methodological issues around ABMs, one of which is the method of validation.

From the outset, it is important to emphasise that verification and validation are two distinct processes. Much of the literature focuses on validation, but verification is equally important. Pilch et al.(2000)[1] and McNamara et al.(2007)[2] characterise the distinction as follows:

- **Verification**
    - The process of determining that a computational software implementation correctly represents a model of a process
    - The process of determining that the equations are solved correctly
- **Validation**
    - The process of assessing the degree to which a computer model is an accurate representation of the real world from the perspective of the models intended applications
    - The process of determining that we are using the correct equations

[1] Pilch, M., Trucano, T., Moya, J., Groehlich, G., Hodges, A., Peercy, D. . 'Guidelines for Sandia ASCI Verification and Validation Plans – Content and Format: Version 2.0.' Albuquerque, NM: Sandia National Laboratories, SAND 2000-3101, 2000

[2] McNamara L., Trucano T and Backus G, 'Verification and Validation as Applied Epistemology', *Proceedings of the UCLA Conference on Human Complex Systems,* April 2007, http://hcs.ucla.edu/arrowhead.htm

Verification

Simulation and Solution

We start with the issue of verification. This is essentially the question: does the model do what we think it is supposed to do? This is conceptually quite distinct from the issue of validation. Whenever a model has an analytical solution, a condition which embraces almost all conventional economic theory, verification is a matter of checking the mathematics.

But there is no inherent superiority in this approach to that of numerical simulation, although not all mathematicians are prepared to accept that simulation based proofs have the same status as analytical ones. A notorious example of this is the proof of the Four Colour Problem, which asked whether a map could be coloured with only four colours so that no two adjacent areas had the same colour. No one has yet found an analytical solution but the advent of computers meant that a brute force simulation based solution could be produced by looking at over 2000 possible cases. Such a proof, written out, would require so long to check that it could not effectively be verified and not all mathematicians accept that this has the status of a proof at all[3].

For most of human history, the inability to conduct such simulation tests has meant that analytical solution has been the only possible approach to model verification. Suppose for a moment, however, that the Greeks had invented the computer and statistical theory rather than algebra. We might now have not an analytical but a statistical concept of proof. Imagine that since the time of Antiquity there had been a stream of developments in computing power. A 'proof' of Fermat's Last Theorem might well be given by calculating that the relevant equality was not satisfied by any power (beyond 2) up to and including ten to the million raised to the million, a number unimaginably larger than any number which might ever be used in practice.

---

[3] Described in Ian Stewart, *From Here to Infinity*, Oxford University Press, 1996,

Returning to the issue of numerical solution and simulation, in the physical sciences this approach is now used widely in the context of partial (non-linear) differential equations, where analytical solutions simply cannot be found. This enables us to extend our understanding of the world. It is not inferior to analytical solution. It complements analysis. Mainstream economics has limited itself by its inability to go beyond analytical results. This also constrains the modelling strategy itself as we show when considering the validation issues.

The downside of ABMs in the verification process is that it is not as straightforward as with analytical solutions. We do not even have widely recognised routines such as the Runge-Kutta approach to numerical solutions of differential equations.

Replication

A key aspect of verification is replicability. Rand and Wilensky (2006)[4] discuss this in detail, in particular the three criteria developed by Axelrod (1997)[5]. These are, first, numerical identity, showing that the original and replicated model produce exactly the same results. This seems less relevant now than perhaps it was ten years ago, because the behavioural rules in many ABMs typically contain stochastic elements.

The second point, distributional equivalence, covers this point. The properties of the original and replicated model should be statistically indistinguishable from each other. Note that the replication needs to be carried out over the same number of solutions as the original, especially when distributions of outcomes are right-skewed so that solving the replicated model for, say, two orders of magnitude more times than the original might generate a really extreme outcome not encompassed in the original results.

This in turn leads to the separate question of the number of times a stochastic ABM needs to be solved in order to establish its properties. There is little discussion of this in the

[4] W Rand and U Wilensky, 'Verification and Validation through Replication: A Case Study Using Axelrod and Hammond's Ethnocentrism Model', *Proc. NAACSOS*, 2006

[5] R Axelrod, ' Advancing the Art of Simulation in the Social Sciences' in *Simulating Social Phenomena'*, eds. R.Conte, R. Hegelsmann and P. Terna. Berlin, Springer-Verlag**:** 21-40, 1997

literature, and in practice, certainly with small ABMs, 500 or 1,000 solutions are often quoted. This is an area which deserves more attention.

Axelrod's third point is 'relational alignment' by which is meant if input variable *x* is increased in both models then if output variable *y* increases in the original model it should also increase in the replicated model. Stated like this, the criterion seems rather weak, and should be extended to be 'if input variable *x* is increased in both models by a given amount, the distribution observed in the changes in output variable *y* should be statistically indistinguishable.

There is a parallel in replication with the world of econometrics. The ability to perform regressions, except of the simplest kind, was severely limited until the 1970s by access to computer power. There was then a proliferation of articles on applied econometrics, which turned into a torrent during the 1980s with the development of the personal computer. Initially, published papers carried the results but no data. Even the descriptions of data sources were often cursory if not cryptic. Replicating someone else's results was a time consuming task, which often failed (the career of one of the present authors started as an applied econometrician). Gradually, a better code of practice in describing data sources evolved, followed by an increasing insistence by the leading journals that the actual data used in the regressions (including all the transformations used) be made available to other researchers, initially in the paper itself and now via the web.

So econometrics has seen a gradual evolution of better replication practice, enabled by successive waves of technology. Even so, we conjecture that only a tiny fraction of published results are ever replicated by third parties, and these will be ones that in general emerge as being considered important by the applied econometric community[6]. A similar sort of process will presumably take place with ABMs. The *Journal of Artificial*

[6] Occasionally, of course, a very bad but influential articles will attract a lot of replication. A notorious example is a paper by Benjamin and Kochin in the *Journal of Political Economy* in the early 1980s which claimed that unemployment in inter-war Britain was caused by high benefits paid to the unemployed. It attracted many replications and rebuttals.

Societies and Social Simulation, for example, requires that the code for published models is made accessible to other researchers. This is an excellent move.

But for the model builders themselves, apart from checks within the group that the code appears to be written as intended, verification is by no means straightforward. A variant of Axelrod's 'relational alignment' is of some use. Namely, if input variable x is increased, does the direction (and possibly size) of the change in output variable y accord with prior beliefs. This is only a partial check, because one of the properties of an ABM may be the emergence at the macro-level of a phenomenon which could not be deduced from the micro-behavioural rules.

This also suggests that testing ABMs may blur the boundary with which we started between verification and validation. If we test the ABM in question by changing the various input variables to test the range of possible output values that the model produces then we have to have some means of judging whether that range is consistent with the model operating correctly or not. In most cases, this can only be done by considering the plausibility of the outcome with reference to the input ranges that have been chosen. This in turn will generally be based on the range of reality that the model is attempting to explain.

This is one reason why a clear description of the problem to be modelled and how the output should be judged is essential to both verification and validation.

Validation

Describing the Problem

It is easy to overlook this stage of validation as modellers plunge into the difficulty of setting up a set of rules and building a model. Yet the process of validation requires a clear view of what the model is attempting to explain and for what purpose. What are the key facts that the model needs to explain and how well must it do it?

Without some answer to this question, it becomes impossible to judge the validation process. Since no model outputs will ever completely describe reality, then we need to know both something about what parts of reality we are trying to get most grip on.

There are likely also to be trade-offs in this description. A model which is trying to capture the range of consumer preferences will need to explain the variety of products but may be unable to precise about product innovation and growth while a model focused on how cities concentrate will focus on a measure of clustering that may have a very simplified geography.

An aspect of a problem description which is often overlooked and becomes even more central for ABMs than other modelling strategies is that of time. ABMs are typically solved in steps. But what is the equivalent in real time with the step in any given ABM?

Conventional economic theory rarely if ever faces this crucial question. An analytical solution is obtained, and an equilibrium solved for a certain set of parameters and exogenous variables values. A change is posited, and a new equilibrium solution calculated. The time taken between these equilibria is almost never considered as an issue. A brilliant exception to this was published as early as 1969 by Atkinson[7]. He showed, inter alia, that the typical time scale of transition from one equilibrium growth path to another in the Solow model was over 100 years. But this article appears to have been exorcised from student reading lists; its implication that economies, even in a strictly neo-classical world, spend a long time out of equilibrium presumably being too disturbing.

Economists often regard the difficulty of translating steps in a model into real time as a weakness of ABMs. But in fact this difficulty is a great strength, in two ways. First, mapping a step into a real time equivalent can be a useful part of the model calibration process. For example, Ormerod (2002, 2004)[8] has a simple ABM from which several

---

[7] AB Atkinson, 'The Timescale of Economic Models: How Long is the Long Run?', *Review of Economic Studies,* April 1969

[8] P Ormerod, 'The US business cycle: power law scaling for interacting units with complex internal structure', *Physica A*, 314, 774-785, 2002, and 'Information cascades and the distribution of economic recessions in capitalist economies', *Physica A*, 2004

key empirical features of the US business cycle emerge. In calculating the size of the growth rate of aggregate output in the model, which emerges from the decisions of firms, the actual range experienced by the US economy during the 20th century is a key way of fixing a step in the model to be equivalent to a year in real time. Second, and more important, it is a valuable aspect of model validation. A sensible rationale has to be provided, in many applications, for the real time equivalent of each step in the model and this is an important of establishing what phenomena are under investigation. In the example quoted above, the model aimed to explain the characteristics of the business cycle over years, so the ability to replicate annual time steps was a useful test.

Another aspect of problem description is to establish what kind of agents are under consideration. As Windrum et al. (2007)[9] note, economists have reacted to the success of ABMs by extending their own framework to incorporate certain aspects of, for example, agent heterogeneity, bounded rationality and increasing returns. All theories, even those like quantum physics which have been subjected to extremely rigorous empirical testing, are approximations to reality. The question is always, how reasonable are the approximations to the features of any particular problem being considered. By restricting themselves in general to models with analytical solutions, economists restrict the set of areas where their assumptions might be reasonable approximations to reality.

A way of thinking about this question is the table below.

Type of theory	**Ability of agents to gather information**	**Ability of agents to process information**
Rational	full	maximise
Bounded rational	partial	maximise
Behavioural ABM	partial	rule of thumb

[9] P Windrum, G Fagiolo and A Moneta,' Empirical Validation of Agent-Based Models: Alternatives and Prospects', *Journal of Artificial Societies and Social Simulation* vol. 10, no. 2, 8, 2007

In standard theory, agents are assumed to gather all relevant information, This immediately restricts the range of problems which can be usefully examined to ones where the dimension is relatively low. Bounded rationality potentially extends this dramatically, by allowing only partial (and asymmetric) information gathering. But with these types of models, the assumption that the decision processing rule is one of maximisation further restricts the closeness of the approximation of the model to reality. Almost the entire discipline of psychology, for example, suggests that agents in general do not behave in this way.

Any individual modelling strategy has to decide what kind of agents are to modelled. Are the agents going to be required to process large amounts of information or can the problem be addressed by simpler ones.

In general, we consider that simpler agents with simpler rules are to be preferred. The simpler the rule, the easier it becomes to test the model and discover its implications. Occam's razor should apply. The key aspect to validation is that the outcomes of the model explain the phenomenon. If the model explains the phenomenon under consideration better than previous models do, it becomes the current best explanation. This is the best we can expect to do.

Testing the outcomes

The key distinguishing feature of ABMs is that the macro properties of the system under consideration emerge from the behavioural rules which are assigned to agents at the micro level. Conventional economists are often resistant to ABMs because of the proliferation of decision making rules which appears to be possible. As Vernon Smith pointed out in his Nobel lecture[10] 'Within economics there is essentially only one model to be adapted to every application: optimization subject to constraints due to resource limitations, institutional rules and /or the behavior of others, as in Cournot-Nash equilibria.'. But Smith followed this immediately with the sentence 'The economic

[10] VL Smith, ''Constructivist and Ecological Rationality in Economics', *American Economic Review*, 93, pp 465-508, 2003

literature is not the best place to find new inspiration beyond these traditional technical methods of modeling'. So modelers developing the ABM tradition can hardly expect to find understanding within mainstream economics.

Both the present authors, and the experience is widely shared amongst ABM modelers, have encountered from economists a view which can be summarized as follows: you have presented one set of behavioural rules to explain your chosen phenomenon, but there must be many such sets which do this, so how do you know yours are correct? Some economists even go on to imply that it is easy to construct successful ABMs, an opinion which merely reveals their ignorance of the difficulties involved. The fact that they do not appear to appreciate that the rules of behaviour incorporated into much standard economic analysis are of a very special kind which does not stand up well to experiment should not blind us to the need to provide a sound basis for the choice of the rules of behaviour for any individual ABM application.

One key test is that the behavioural rules should be capable of justification using evidence from outside the model.[11] The better this evidence, the more credible the rules. If evidence is absent, then the appropriate response is to create simple agents whose rules of behaviour are easy to understand. For example, if the agents are firms then there is a large body of evidence that profit seeking behaviour will be observed, even if it is modified and overlaid by other motivations. On the other hand, if the problem is to consider how cities grow, where there will be multifarious agents whose rules of behaviour would be hard to describe, it may be better to test growth patterns by choosing simple agents, such as a unit of geography and allow it just to choose its mix of activities[12].

Another test is to compare the success in meeting a description of the phenomenon under consideration of different models. At present, there is little competition between ABMs

[11] P Ormerod and B Rosewell,' On the Methodology of Assessing Agent Based Models in the Social Sciences', in J Stanley Metcalfe and J Foster, *Evolution and Economi Complexity, Edward Elgar, 2004*
[12] B Rosewell and A Horton, 'Neighbours are not enough', paper presented at the European Conference on Complex Systems, Oxford, September 2006.

attempting to explain the same set of macro phenomena. We are still at the stage where only one ABM exists which accounts for any given set of macro features.

This does not, however, mean that such an ABM is not scientific. On the contrary, the inherent methodology of ABMs is far more scientific than that of conventional economics. We identify a set of empirical macro features, and plausible behavioural rules are designed from which the macro properties emerge. This is a much more scientific methodology than econometrics, for example, much of which is mere curve fitting and is not modeling in any real sense of the word. And outcomes of ABMs can be compared with the explanations given by more traditional models.

In an important sense, the current process of building ABMs is a discovery process, of discovering the types of behavioural rules for agents which appear to be consistent with phenomena we observe. Once we leave behind the comfort blanket of maximization, we need to find out what works. We return to this point below.

Comparing model outcomes

One area where competing agent based models exist is the business cycle. Three are described here. All are successful in that from simple behavioural rules, a range of complex macro features of reality emerge. But they differ in scope and scale.

The simplest is the model of Ormerod[13]. In this model, the only agents are firms. We know that ex post most of the fluctuations in output over the course of the cycle arise from the corporate sector, so this is a reasonable approximation to reality. The agents are heterogeneous, both in size and in how they take into account uncertainty. They are myopic, and look only one period ahead with limited information. Firms use very simple rules to decide their rate of growth of output and their sentiment about the future, their 'animal spirits' in Keynes' phrase. From these simple behavioural rules emerge several key features of the business cycle:

- Positive correlations in output growth between individual firms over the cycle
- The autocorrelation properties of real US output growth in the 20th century

[13] Op.cit.

- The power spectrum of real US output growth in the 20th century
- The exponential distribution of the cumulative size of recessions

The model of Dosi et.al. (2006)[14] contains considerably more features. In addition to firms, workers/consumers also appear in the model. There are two types of firm, one of which carries out research and development and produces machine tools, the other buys machine tools and produces the consumption good. Profits and wages are determined within the model. There is technical change and firm entry and exit. This model embraces more features of capitalist economies, but the behavioural rules within the model remain simple. Rather, the agents operate under imperfect information, are myopic and use rules of thumb behaviour rather than maximization.

From this model emerge considerably more empirical phenomena observed in capitalism, but the set of phenomena reported does not overlap with the previous, much smaller model. An important extension of the Dosi model from that of Ormerod is that the latter assumes (with strong empirical justification) that most of the fluctuations in output arise from firms rather than from the consumption of the personal sector. This stylized fact emerges from the rules of the Dosi model rather than being assumed. In addition, the Dosi model is consistent with the following stylized facts:

- Investment and consumption tend to be pro-cyclical and coincident variables
- Employment and unemployment are lagging variables, the former being pro-cyclical and the latter anti-cyclical
- Productivity dispersion across firms is large at any point in time, and is persistent over time
- Firm size distributions are considerably right skewed
- Firm growth rate distributions are not Gaussian and can be proxied by fat-tailed tent-shaped densities

These two models, then, appear at first sight to be diverse. But closer inspection shows that they have a lot in common. In each model, a key feature of its ability to replicate macro phenomena is the heterogeneity amongst individual agents. The ‘representative agent’ of neo-classical theory has to be discarded completely. Further, in each model decision makers are myopic, and operate under both uncertainty and imperfect

[14] G Dosi, G Fagiolo and A Roventini, ‘An Evolutionary Model of Endogenous Business Cycles’, *Computational Economics*, 2006

information. It is as if decision makers in each of the models operates with low cognition. They are short-sighted, gather uncertain and limited information, and use simple rules with which to make decisions. The results of both models also contain a more complete description of the business cycle than non agent based models.

The model of Wright (2005)[15] reduces the level of cognition assigned to agents even further. The model has nine rules which attempt to capture the social relations of production under capitalism. Almost all of these are purely stochastic, so it is as if agents are operating with zero cognitive ability. The deterministic aspect of the model is simply that firms must have enough money to pay the labour force, and if not they fire workers (at random) until this constraint is satisfied. In the model, a small class of capitalists employ a large class of workers organized within firms of various sizes that produce goods and services for sale in the marketplace. Capitalist owners of firms receive revenue and workers receive a proportion of the revenue as wages.

Again, the stylized facts which emerge from the model contain some overlap with the Dosi et.al. model, but additional ones as well:

- The power law distribution of firm size
- The Laplace distribution of firm growth
- The lognormal distribution (when aggregated over time) of firm extinctions
- The lognormal-Pareto distribution of income
- The gamma-like distribution of the rate of profit of firms

Deciding which of these models is more valid (we assume that they are equally verified) is not straightforward. They each produce results which are consistent with what is observed and do so with agent decision rules which are easy to describe and explain. The elements of reality which are explained are in each case more comprehensive than the standard model but add different levels of explanation.

Simple agents, simple rules

[15] I Wright, 'The social architecture of capitalism', *Physica A*, 346, 589-620, 2005

One criterion that might be used to determine performance is simplicity of behaviour, on the principle that if simple agent rules can produce a good description, this is better than having complicated ones.

Another way of expressing this is to ensure that agents are only required to have the minimum necessary ability to process information or to learn. The issue of low, or even zero, cognition of agents is, we believe, an important aspect of the validation of ABMs. The idea that agents act, in general, in a way far removed from the precepts of rational choice economics is now supported by an impressive body of evidence.[16] It is models which assign, implicitly or explicitly, a high level of cognition to agents which need special justification rather than those which do not.

We suggest that a model that can only work if agents have high levels of cognition (information processing and learning) will only be able to capture a limited set of outcomes. Indeed one way of testing an ABM in the social sciences is to assign increasing levels of cognition to agents to see at what point the model ceases to provide a description of reality.

An example here is Ormerod and Roswell (2003)[17]. Two key stylised facts have been established about the extinction patterns of firms. First, the probability of extinction is highest at the start of the firm's existence, but soon becomes more or less invariant to the age of the firm. Second, a recent finding, that the relationship between the size and frequency of firm extinctions is closely approximated by a power law. The model is in the spirit of the Wright model, in that the focus is on the structural relationships between firms rather than on specific rules of agent behaviour.

In the basic model, firms by definition have zero cognition. Firms are assigned at random an initial level of fitness. A matrix, whose elements are initially chosen at random, describes both the sign and the size of the impact of all other firms on the fitness of any given firm. The actions of a firm can increase the fitness of another (e.g. if it is an efficient supplier to that firm) or decrease it (e.g. if they are direct competitors in the

[16] For example, Kahneman's Nobel lecture, D Kahenman, 'Maps of Bounded Rationality: A Perspective on Intuitive Judgement and Choice', *American Economic Review*, 2003

[17] P Ormerod and B Rosewell, 'What can firms know?' *Proc NAACSOS conference*, 2003

labour market). Despite the stark simplicity of the model, it contains a key feature of reality which is absent from most of conventional economics. The focus of the latter is upon competition between firms, whether in the output or the labour markets. But in reality, a great deal of economic activity is business-to-business in which the success of any given firm is beneficial to those which it supplies.

In each step of the model, one of the elements of the matrix of connections is updated at random for each firm. A firm is deemed extinct if its fitness falls below zero. So the firms are unable by definition to acquire knowledge about either the true impact of other firms' strategies on its own fitness, or the true impact of changes to its own strategies on its fitness.

From this model with simple rules of behaviour, the observed stylised facts emerge. We then go on to examine the effects of allowing firms different amounts of knowledge about the effects of strategy in the context of the agent-based evolutionary model. There are very considerable returns in the model to acquiring knowledge. There is a sharp increase in the mean agent age at extinction for agents with even a small amount of knowledge compared to those without. Indeed, we find that as both the amount of knowledge available to firms increases and as the number of firms capable of acquiring such knowledge rises, the lifespan of agents begins to approach the limiting, full information paradigm of neo-classical theory in which agents live for ever. However, even with relatively low levels of knowledge and numbers of agents capable of acquiring it, the model ceases to have properties which are compatible with the two key stylised facts on firm extinctions. The clear implication is that firms have very limited capacities to acquire knowledge about the true impact of their strategies.

By starting with simple rules, we establish that key facts emerge from the interactions of the model. Giving the agents more complicated rules in fact turns out to destroy this feature. Only by starting with the simplest possible configuration could this be discovered.

Conclusion

Agent Based Models face a variety of issues in verification and validation which are new precisely because ABMs offer the opportunity to model a wider class of phenomena than has been possible before.

Because such models are based on simulation, the lack of an analytical solution (in general) means that verification is harder, since there is no one result the model must match. Moreover, testing the range of model outcomes provides a test only in respect to a prior judgment on the plausibility of the potential range of outcomes.

In this sense, verification blends into validation. We take the Popperian position that validation can never be proof and that we are therefore seeking for models which explain more than their predecessors and are not falsified. Thus we stress that an important part of validation is a clear description of what is being explained. This should also ideally include a description of what is not explained by the current best practice.

Our second stress is on simplicity. The validation of models with complicated agents and complicated rules is impossible in our current state of knowledge. It may be that over time the validation of simpler models will then lead to the ability to make them more complicated as levels of validation are built up. Proofs in mathematics now often require this nested approach.

However, agent based models have not yet reached this position and in our view will not do so unless there is rigorous testing of the simplest possible models at the outset. Thus we do not believe that a model with complicated agents should be accepted unless and until it has been shown that simpler (lower information, lower cognition, less processing, less learning) ones will not explain the phenomenon just as well. In particular, we believe that the growing empirical body of evidence which shows that in general agents act in an intuitive rather than rational way is an essential part of model validation. Models which assign high levels of cognition to their agents need particularly strong justification if they are to be considered valid.