From Complexity to Agent Modelling and Back Again - some implications for economics

Bruce Edmonds,
Centre for Policy Modelling,
Manchester Metropolitan University.
http://www.cpm.mmu.ac.uk/~bruce

1 What is Complexity?

1.1 "Complexity" in economics

Since Herbert Simon's distinction between substantive and procedurally rationality [18], complexity in economics has (until recently) simply meant *not* assuming that an economic agent acted as if it had the computational resources to completely cope with the demand placed on it by its environment (for a survey of this area see [11]). In a sense this is intended as a negative definition to contrast with the usual simplifying assumptions of mainstream economics. The use of the word "complexity" here reflects the greater analytic difficulty which is caused by relaxing these assumptions. However, as generally used, this usage is vague - it does not necessarily distinguish between different kinds of limitations on the agent (memory, inference, ability to generalise, access to information, etc.) or different kinds of overloading caused by the environment. It also encourages the conflation of the complexity from the agent's point of view and from the modeller's view. Thus although it has performed a useful role in critique it is not a very helpful usage for more focused analytic discourse.

1.2 The "Sciences of Complexity"

Recently there has been an explosion of fields and techniques loosely grouped around the "Sciences of Complexity" banner. Even though these have been considerably over-hyped they include many new useful techniques and metaphors. They are not, however, based on any coherent "complexity theory". Key practitioners in these fields recognise this and have frequently exhorted, anticipated and worked towards such a theory (e.g. [2, 12]), but at the moment there is no such body of general theory which has a useful analytical meaning outside individual component fields.

What these widely differing areas *do* share is a tendency to:

- use newer formal techniques (logics, automata, topological models, etc.);
- deal with systems where the behaviour of interest arises, at least partly, out of some contingent occurrences (e.g. evolutionary computation);
- tend to use modelling techniques which tend to predict and capture mainly second-order properties (e.g. it may not be able to predict a value at a particular future time but may be able to tell you some statistical properties or reveal some qualitative aspects of the process).

Thus in a broad sense these techniques try to capture some more abstract properties of systems in areas where the application of more traditional techniques is not feasible. Thus this is also a negative definition, indicating that newer approaches had had to be applied. Again, if one needs to attempt a more exact and meaningful analysis this is an inadequate characterisation of the term.

1.3 Complexity *per se*

So what is complexity *per se*? Let us approach this via a series of considerations.

Firstly one has to distinguish what complexity *is* and what may *cause* it. Without an idea of the former it will be very difficult to get a clear idea of the later. In general there will be *many* possible cause of complexity and there well may be no *overall* characterisation of such causes.

Secondly, I argue that complexity is not a property *usefully* attributed to natural systems but only to our models of such systems. The reasons for this include:

- If natural systems have a complexity it would be beyond us; it is always possible to arbitrarily increase the complexity of a system by considering more aspects or more detail;
- The complexity of things varies *critically* with the model chosen;

In this way it is similar to the property of “primality” - primality is a property of (some) numbers and not of things enumerated by numbers, even though the selection of groups to be represented by numbers can effect whether the property holds (merely because it can change the number being considered).

Thirdly, the complexity is a comparative thing, frequently we want to be able to say “A is more complex than B which is, in turn, more complex than C”.

Fourthly, complexity is relative to the framework you are modelling in. This includes the language of representation of your model, your general framework (what is given what you are trying to formulate) and your goals in modelling.

Finally complexity is usefully distinguished from ignorance. Ignorance can cause complexity in many ways, including the misframing of a problem, and complexity can certainly cause ignorance, for example where there are only limited problem solving resources available. However, it is easy to give examples where the two concepts diverge.

Packaging these all up into a definition we get:

> *“Complexity is that property of models which make it difficult to formulate its overall behaviour in a given language of representation, even when given almost complete information about its components and their inter-relations” [6]*

Thus I have relativised complexity to the model, the type of difficulty, the language of modelling and how one characterises “overall behaviour” and “atomic components”. Many of the confusions that have abounded with the use of this word have occurred because of unshared assumptions about these relativisations. Thus you will get different kinds of complexity for different kinds of “difficulty”, different modelling languages etc.

Frequently the complexity of a pattern, system or data model is taken to be the complexity of the most “suitable” model given a certain framework. The definition above does not mean that you can make complexity to mean whatever you want. When a framework has been agreed (either explicitly or, as frequently occurs in the hard sciences, implicitly), complexity can be objectively measured and attributed, just as whether the number of cows in a particular field is prime is an objective question despite the fact that “primality” refers to our (numerical) model of the cows and not the cows themselves.

Such a definition imposes some obligation to explicitly make clear the chosen relativisations but this is no bad thing. The justification of such an approach to defining complexity ultimately comes down to the extent it makes the analysis of the effects and causes of complexity clear (for more on this see [9]). We will now apply it to modelling agents which economically interact.

2 The effects of complexity on modelling by agents

We will define “environmental complexity” as a special case of the above definition as:

> *The difficulty of making correct predictions about its environment (measured by its error rate) for an agent using the best model it can infer from the information available to it given its computational resources.*

This is a more precise version of the “economic complexity” described in section 1.1. We will define the “model complexity”, pertaining to a particular model of an agent as:

> *The computational difficulty of reaching and testing a model given the constraints of the language and the known data it has to fit.*

I will consider four situations representing cases of increasing environmental complexity. The environmental complexity will affect how the agent needs to evaluate its models, in particular the kind of trade-offs between the model complexity, specificity and error.

2.1 Ideal rationality and perfect information

If an agent is in a situation where it effectively has all the time it needs for relevant computation and learning with respect to its environment (e.g. if the environment is relatively static and the agent has a long time between decisions or if the whole population acts as a sort of massively parallel search and inference mechanism and the agent has access to the “results”) then it can be treated as if it had ideal rationality and perfect information. If you have an agent in such a position then the only relevant criteria there is for judging alternative models about its environment is that of the *accuracy* (or conversely error) of their predictions.

2.2 Ideal rationality and noisy information

A slight increase in the environmental complexity for the agent it when is has noisy data but enough of it to determine the extent of this noise. This requires only a slight increase in sophistication by the agent - it has enough data and computational resources to determine the form of the correct model but might need to estimate its parameters. This situation has been studied in economics as representing a simple learning process. Here the criterion of accuracy is not needed only for the determination of its best model but also as a characterisation of its resulting behaviour.

2.3 Ideal rationality and inadequate information

Let us increase the prediction complexity further by now only giving it inadequate (i.e. noisy and insufficient) information about its environment. Let us assume that the agent can posit deliberately imprecise models, i.e. it can include room for inexact predictions using some mechanism like error terms in its language of representation[1]. Now the most appropriate model an agent can infer will typically be, at least somewhat, imprecise, so that as well as *accuracy* the agent also has to take into account the *specificity* of its models. Since the agent has insufficient information it has no way of certainly distinguishing noisy data from very complex behaviour so there will be an inevitable trade-off between the accuracy and specificity of candidate models. Although some [19] have argued that particular trade-offs can be principled, in general the nature of this trade-off will depend on the goals of the agent (e.g. its tolerance to risk).

2.4 Bounded rationality and inadequate information

Finally we come to the situation which is most environmentally complex. Here, not only does the agent have inaccurate and insufficient information but also that it does not have the computational resources to search the space of possible models for the optimal one. In other words the *model complexity* has to be taken into account along with *accuracy* and *specificity* in the evaluation of candidate models, for there will be a limit to the complexity of models it can consider. Now we have a three-way trade off to consider. - the nature of this will depend upon the agent’s goals (some examples of this are given in section 4.2 below).

3 The effects of modelling by agents on complexity

For us modellers of these agents, there is the difficulty of capturing an agent’s (or population of agent’s) resulting behaviour. I will call this the “behavioural complexity”, which I define as:

> *The difficulty of finding a model of the overall population’s agents behaviour given knowledge of the setup, algorithms and structure of each agent.*

1. If it only had a completely precise modelling language (i.e. its models could only make exact predictions) it would tend to overfit noisy data, unless the language was deliberately restricted to avoid this.

The need to include more about the process of modelling by agents in our models of them increases the behavioural complexity of such agents.

3.1 Ignoring the process of modelling by economic agents

Since *we* are agents with distinctly bounded rationality with access only to imperfection information (as described in section 2.4 above), it makes sense for us to start with simple models and only progress to more complex ones if wee need to, i.e. if our simpler models have inadequate predictive and explanatory power. If the agents are in a situation where the environment is changing sufficiently slowly and the agent has effective access to good information one can conflate the agent's best model of its environment and the true model of the environment, as the agent has plenty of time and resource to effectively estimate the true model.

In such circumstances it is sensible for us to choose the simpler model and ignore the agent's modelling process - it is notable that if *we* did not have distinctly bounded rationality and inadequate information there would be no reason for us *not* to search through models of all complexities for the most accurate one.

3.2 Including the process of modelling by economic agents

The other case is where the conditions described above do not hold, i.e. where the process whereby the agent models its environment itself has significant effects. That such cases may exist in real life is surely not in doubt (e.g. in emerging markets, during periods of great volatility in the stock market or the behaviour of economists themselves). What *is* debateable is:

1. the extent of these situations;
2. the tractability of them;
3. whether such an area of study can be said to be part of the tradition of economics as it has developed.

Although point (1) is the most important, to be decided by empirical evidence, I do not have either the knowledge or the time to consider this. I will only comment that since the assumption that one does not loose anything important by conflating the agents model with the true model is a very strong one, the burden of evidence should weigh more heavily on it as opposed to its converse.

Point (2) is a valid point, and one which I am highlighting here in this paper. The introduction of including processes of modelling by agents in our models, *greatly* increases the behavioural complexity of the situation. In fact such an introduction seems to introduce a qualitative jump in the whole enterprise of such modelling[2]. On the other hand it does seem, that at least some aspects of such situations are meaningfully formalisable and analysable (for example both [1] and [5] exhibit credible models of economic systems which exhibit both rational expectations and more complex dynamics in clearly defined circumstances).

Point (3) is in many ways trivial. The qualitative leap highlighted above may justify categorising such modelling under a separate heading, but unless such models turn out to be significantly less successful than more traditional ones, this would not justify this not being a substantial and legitimate line of research.

4 Towards dealing with the complexity of modelling agents - modelling modelling

Thus I have built up a picture of the connections in this context between types of complexity:.

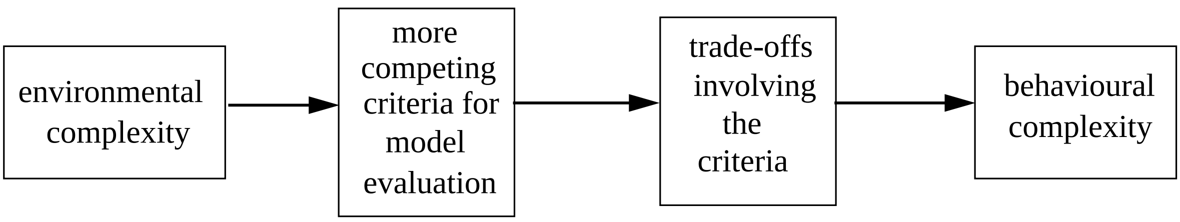

Figure 1: A diagram of the relationship between relevant types of complexity

2. Comparable, perhaps to the qualitative difference induced by the presence of some self-modelling ability.

If what I have said is correct, we are faced with a task of great behavioural complexity. So how can we proceed? In particular how can we succeed in a way that is not entirely specific to each situation? Such an enterprise comes down to nothing less than establishing a framework for modelling the process of modelling itself. Below I briefly discuss some of the key features that such a framework might have (for more details on this see [15]).

4.1 The Form - meaning distinction

The most fundamental distinction, as discussed above, is between the form of an agent's model and its meaning. Where the *meaning* of a model is determined by the mapping of the input and resulting predictions to the model. This is very similar to the syntax-semantics distinction in formal logic[3]. This distinction is critical where we are dealing with agents of bounded rationality. For example, although two models may be equivalent in terms of its predictions, one could be so costly to use that it may be all but useless in deciding on an appropriate action.

In a very general way this can be formalised by reference to the space of all relevant possibilities (similar to phase space diagrams used in physics). Models, goals, a priori knowledge, observations, goals and actions can all be associated with subspaces of this (see figure 2).

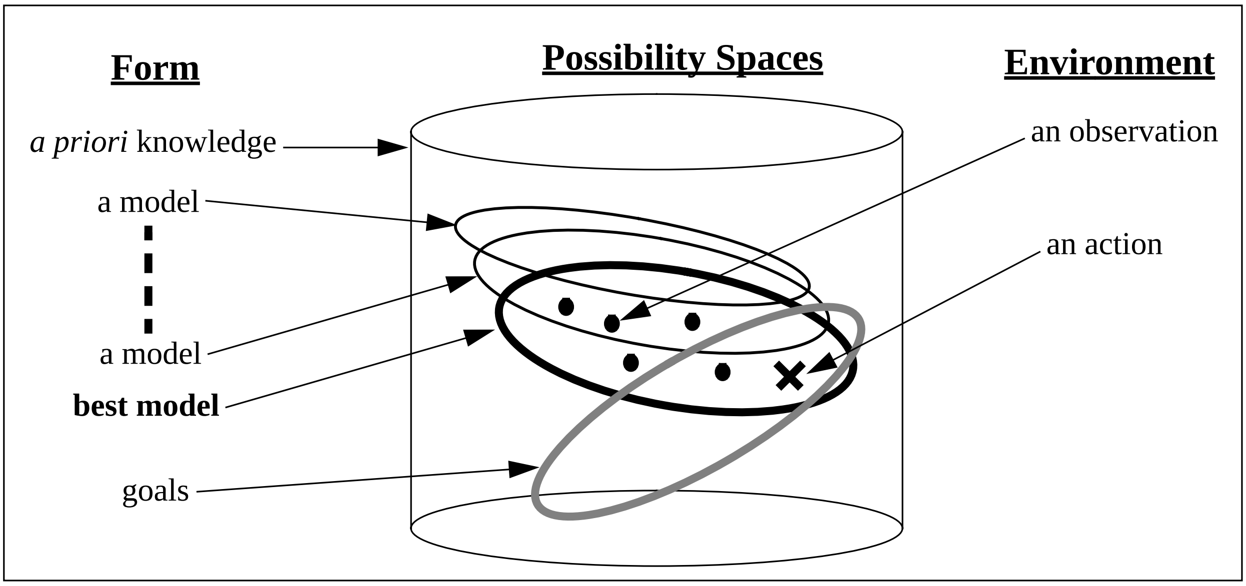

Figure 2: Form and meaning distinction

4.2 The complexity, specificity, error trade-off

As mentioned above, in situations of greater environmental complexity it is sensible for the agent to accept some trade-offs between model complexity, specificity and error.

There are several desirable formal restrictions on these measures (which can be found in [15]) and there are philosophical justifications and consequences for this three-dimensional analysis of models [9]. Here I will just look at a three examples to illustrate the different possibilities that can result.

1. A *risk-averse* agent might take the error rate as the primary criterion whilst accepting the model complexity as representing only a limitation on its resources. It would be very tolerant to vagueness (i.e. low specificity) only accepting a more specific model where one could be found without losing any significant accuracy. This would be particularly appropriate in safety-critical situations.
2. On the other hand if an agent is merely trying to predict some out-of-sample data as well as possible *on average*, then there is no particular reason not to choose quite a specific model. In addition (and depending somewhat on the problem domain) one might not want to choose an overly complex or simple model for fear of overfitting or overpredicting (e.g. as illustrated by [17]).

3. In fact this distinction can be used to induce a logical structure between the agent's models [15].

3. Finally, if an agent had an external source of candidate models and it *knew* that there would have been a tendency to elaborate (make more complex) these when they were relatively unsuccessful on in-sample data, applying a heuristic of preferring simplicity when the evidence is equal would be appropriate.

4.3 The modelling language

The collection of all possible model forms can be considered as a *language* in its broadest sense. This might correspond to a natural or formal language, but might also be something like a range of possible real-valued vectors (as in some connectionist models). What is clear is that the modelling language can critically effect the nature and "success" of the modelling process. For example if the agent has a relatively inexpressive language that can only capture some of the real possibilities then it can be genuinely surprised by observations from the environment (as distinct from merely experiencing what it considered to be a low probability event). Here it literally can not model certain combinations of possibilities, even if it *can* model the whole possibility space.

For example, in chaotic processes one can be forced to models of ever greater size as one approaches the critical point at onset of chaos if one restricts oneself to using finite automata but this is simply handlable when using [4]. Despite its importance in learning processes and calls for it to be further studied [20], very little is known about the effect of different modelling languages.

4.4 Processes of model development

So far I have talked about the universe of such models (the modelling language) and the way candidate models are evaluated (the complexity-specificity-accuracy trade-off). The last main aspect of capturing such modelling is the mechanism by which new models are developed.

Despite the fact that there are a multitude of candidate mechanisms to choose from in AI and cognitive science two mechanisms have dominated economic models of learning: those of optimization and evolution (GA/GP/EP etc.). Of these approaches those based on the GP paradigm [13] appear the most appropriate, due to the potentially expressive format of the genome. Using such an approach one can represent an agent as a population of competing mental models [8]; a possible architecture for such an agent is shown in figure 3.

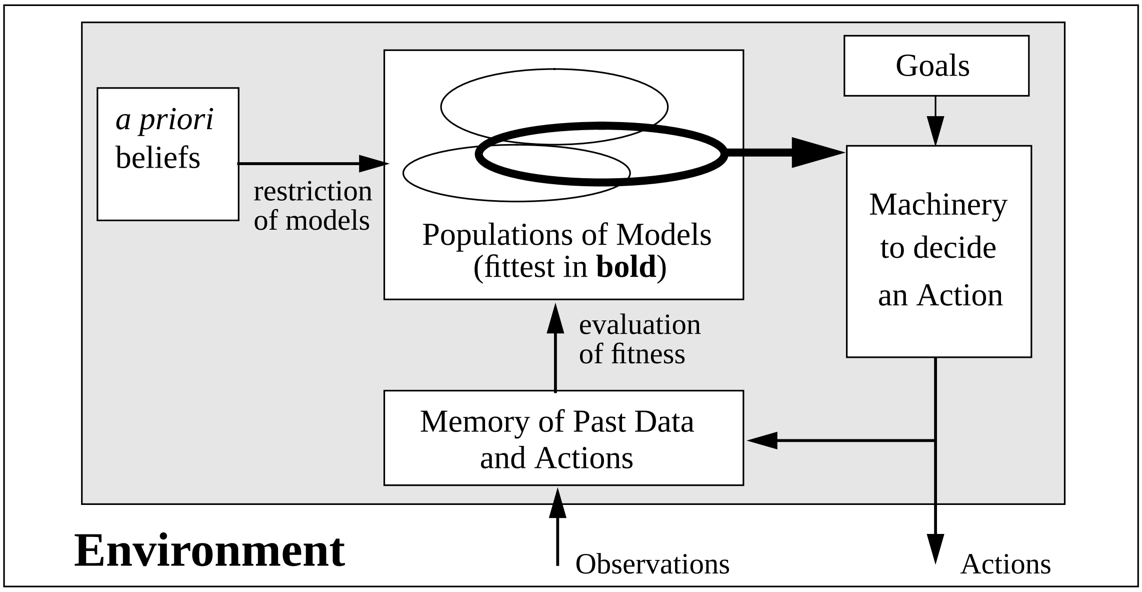

Figure 3: Using a genetic population to model an economic agent

The evolutionary approach has the weakness that the development of variation is *blind* to the process of selection. In other words it avoids the intentionality inherent in human learning. There are, however many other credible models for learning [3].

5 Some future directions for economic modelling

I will end with some speculation as to possible future avenues that the modelling of behaviourally complex systems could go.

5.1 Applying our model of modelling to ourselves

Finally let us apply the above model of modelling to ourselves as modellers with bounded rationality and inadequate information. Further let us suppose that our current economic models gave us an unacceptable level of error when applied to situations of greater environmental and hence behavioural complexity. What avenues might this suggest?

1. Accept the level of error as inevitable.
2. Trade in specificity for accuracy, i.e. accept vaguer models of the same type to decrease error.
3. Trade simplicity for accuracy, i.e. expend more effort in search for more complex models of the same type.
4. Change the modelling language, i.e. consider new types of models.

These are roughly ordered in terms of increasing optimism. (1) represents a position of extreme pessimism, and is probably only held by a few holists. (2) is a classical defensive posture for the less specific your models are the less *refutable* they are. (3) represents an optimistic approach - it assumes that the type of model is fine but just needs refining (e.g. by adding a “missing” variable). (4) represents the most radical response, appropriate when the existing modelling language can not even describe the behaviour to be modelled. It is perhaps this option that is usually what is indicated by authors when the “complexity” banner is unfurled.

5.2 Relatively new (non-numerical) techniques

The history of economics has been dominated by the use of numerical, statistical and game-theoretic modelling languages. Here are three other possibilities that have been taken up by the new sciences:

1. Formal logics, which allow for a much greater expressivity, especially when one is dealing with qualitative as well as quantitative properties. For example in [14] the assumptions behind some organisational theories are examined. Although these are difficult to use, computer based tools to aid in their use are becoming increasingly accessible for the normal user (e.g. SDML [7]).
2. Network based models can be used to capture topological properties in a flexible way, for example in [16] we produced a network model of R&D development which captured some of the context-sensitive nature of the dependencies between technologies.
3. Abstract formal languages, can used to capture generative and modelling behaviours using artificial grammars. I am not aware of an application in economics, but this has been applied to chaotic process [4] and in biology [10].

6 Conclusion - complexity again

I have analysed complexity and distinguished three different kinds of complexity relevant to the modelling of economic agents: environmental complexity, the complexity of an agent’s models and the behavioural complexity. The move to considering situations of greater environmental complexity leads to the consideration of agents that evaluate their models using more than just the accuracy of their models and this leads to a greater behavioural complexity. In response to such complexity we may need to use new styles of models that can capture the new behaviour and processes, i.e. the modelling language we use might need to become more sophisticated and hence also the models we build in it.

We have gone from a more closely defined version of economic complexity - environmental complexity all the way to the “sciences of complexity”. All these kinds of complexity are *linked* but are far from identical. Hopefully this paper has elucidated the nature of these links.

References

[1] Arthur, B. et al. (1996). Asset pricing under endogenous expectations in an artificial stock market. University of Wisconsin-Madison *SSRI report 9625*, 1180 Observatory Drive, Madison, WI 53706, USA.

[2] Casti, J.L. (1992). The simply complex: trendy buzzword or emerging new science, Bulletin of the Santa Fe Institute, 7, 10-13.

[3] Cohen, P. R. (1985). Heuristic Reasoning about Uncertainty: An Artificial Intelligence Approach. Pitman, Boston, MA.

[4] Crutchfield, J.P. (1994). The calculi of emergence: computation, dynamics and induction. *Physica D*, 75, 11-54.

[5] Darley, V.M. and Kauffman, S.A. (forthcoming). Natural Rationality. In Arthur, WB et al. (eds) *The Economy as a Complex Evolving System*, Volume II. Addison-Wesley, Reading MA, forthcoming. Also available as a Santa Fe Institute report SFI 96-08-071 (http://www.santafe.edu/sfi/publications/Working-Papers/96-08-071.ps)

[6] Edmonds, B. (forthcoming). What is Complexity? The philosophy of Complexity per se with application to some examples in evolution. In F. Heylighen & D. Aerts (eds.), *The Evolution of Complexity*, Kluwer, Dordrecht. (http://www.cpm.mmu.ac.uk/~bruce/evolcomp)

[7] Edmonds, B., Moss, S.J. and Wallis, S. (1996). Logic, Reasoning and a Programming Language for Simulating Economic and Business Processes with Artificial Intelligent Agents. In Ein-Dor, Phillip (ed.): *Artificial Intelligence in Economics and Management*. Boston: Kluwer Academic Publishers, 221-230. (http://www.cpm.mmu.ac.uk/cpm/cpmrep09.html)

[8] Edmonds, B. (1997). The Introduction of Learning into the Modelling of Boundedly Rational Economic Agents using the Genetic Programming Paradigm. *AISB'97 workshop on Evolutionary Computation*, Manchester, April 1997. (http://www.cpm.mmu.ac.uk/cpm/cpmrep10.html)

[9] Edmonds, B. (1997).Complexity and Scientific Modelling. *CPM Report 97-23*, MMU, 1997. (http://www.cpm.mmu.ac.uk/cpm/cpmrep23.html)

[10] Green, D.G. (1990). Syntactic modelling and simulation. *Simulation* 54(6), 281-286. (http://life.csu.edu.au/~dgreen/papers/simul.html)

[11] Holm, H.J. (1993). *Complexity in Economic Theory*, University of Lund: Lund Economic Studies, Lund, Sweden.

[12] Kauffman, S.A. (1993). *The Origins of Order.* Oxford University Press, New York.

[13] Koza, J.R. (1992). *Genetic Programming*. MIT Press, Cambridge, MA.

[14] Masuch, M. and Huang, Z. (1994). A logical deconstruction of organizational action: formalising J.D. Thompson's 'Organisations in Action' in a multi-agent action logic. *CCSOM working paper 94-120*, Department of Statistics and Methodology, University of Amsterdam.

[15] Moss, S. and Edmonds, B. (1994). Modelling Learning as Modelling, *CPM Report 94-03*, MMU, 1994. (http://www.cpm.mmu.ac.uk/cpm/cpmrep03.html)

[16] Moss, S., Edmonds, B and Gaylard, H. (1996). Modelling R&D Strategy as a Network Search Problem. *Workshop on the Multiple Linkages between Technological Change, Human Capital and the Economy*, University 'Tor Vergata' of Rome, March 1996. (http://www.cpm.mmu.ac.uk/cpm/cpmrep11.html)

[17] Murphy, P.M. and Pazzani, M.J. (1994). Exploring the Decision Forest: an empirical investigation of Occam's razor in decision tree induction, *Journal of Artificial Intelligence Research*, 1, 257-275. (http://www.cs.washington.edu/research/jair/abstracts/murphy94a.html)

[18] Simon, H.A. (1992). *Economics, Bounded Rationality and the Cognitive Revolution*, Edward Elgar, Brookfield, Vermont

[19] Rissanen, J. (1990). Complexity of Models. In *Complexity, Entropy and the Physics of Information*, Zurek, W.H. (ed), Addison-Wesley, Redwood City, CA, 117-125.

[20] Russell, S. and Subramanian, D. (1990). Mutual constraints on Representation and Inference. In *Machine Learning, Meta-Reasoning and Logics*, Brazdil, P.B. and Konolige, K. (eds.), Kluwer Academic, Boston, 85-106.