

Social Simulation Models and Reality: Three Approaches

Scott Moss

Director, Centre for Policy Modelling, Manchester Metropolitan University, Aytoun Building
Manchester M1 3GH, United Kingdom
s.moss@mmu.ac.uk
<http://www.cpm.mmu.ac.uk/~scott>

Theories in the social sciences are informed either by sociology or by economics. That is, they either draw generalizations from verbal descriptions of social interaction or from mathematical representations of interacting agents as constrained maximizers. The three models discussed in this paper are not less rigorous than economic models and, because they are validated relative to a formal logic rather than mathematics alone, capture much more of the richness of sociological analysis than does the economics approach. Moreover, each has, in its own way, been developed to capture salient characteristics of observed regularities in specific social interactions. Together they exemplify a set of modelling techniques and an integrated methodology which capture the rigour and precision of the economist's approach to the social sciences while, at the same time, capturing the suggestiveness and richness of the sociologist's approach.

1 Introduction

It is only a little simplistic to say that theories in the social sciences derive either from sociology or from economics. That is, they either draw generalizations from verbal descriptions of social interaction or from mathematical representations of interacting agents as constrained optimizers. The three models discussed in this paper are not less rigorous than economic models and, because they are validated relative to a formal logic rather than mathematics alone, capture much more of the richness of sociological analysis. Moreover, each has, in its own way, been developed to capture salient characteristics of observed regularities in actual social interactions.

The purpose of this paper is to present these models as examples of an integrated methodology which captures the rigour and precision of the economist's approach to the social sciences and while preserving the suggestiveness and richness of the sociologist's approach to the social sciences. The construction and experimentation of these models is related explicitly to their validation and verification.

2 Model validation and verification

Validation is the process of ensuring that a computer program will not crash and also that it will perform in the manner intended by its designers and implementors. Verification is the process of assessing the goodness of fit to the characteristics of the models empirical referents. These two topics are considered in turn.

2.1 Validation

If a program runs without error in any computer programming language, then that program is consistent and sound relative to that language. That is, the program does not generate or entail mutually contradictory statements and it does not generate statements which the language does not support.

If the programming language corresponds to a logical formalism, then any program viewed as a set of statements or sentences which runs in that language will necessarily be sound and consistent relative to that logical formalism. One such language, implemented precisely to capture this feature of programs, is SDML [1] which corresponds to a fragment of strongly-grounded autoepistemic logic [2].

Programming languages generally make it easier to do some things than others. Fortran is optimized for numerical calculation; LISP for functional programming, PROLOG for backward-chaining deduction, and so on. Numerical calculations, functional programming and backward chaining can all be programmed in any of these languages though, as stated, there is a clear correspondence between the ease of programming in any of these styles and the language used.

The immediate advantage of programming in a language which corresponds to a logical formalism is that there are no hidden assumptions in the form of unstated axioms or rules of inference. In SDML, and therefore in the models reported below, each agent is defined on a rulebase and database for each period of time. Every rule in the rulebases and every clause asserted to the databases is sound and consistent relative to strongly grounded autoepistemic logic. This is not to imply that agents are individually represented as a logical formalism as in, for example, BDI agent-based models - only that the whole model is sound and consistent relative to the formalism to which SDML corresponds.

2.2 Verification

Verification by comparing model output with statistical data series is too well established to warrant detailed consideration here. Numerical forecasting models have been shown on a number of occasions to be improved by expert intervention ([3], [4], [5]). Integrating statistical forecasting models with rulebases incorporating expert judgement has been shown by Collopy and Armstrong [6] and by Moss, Artis and Ormerod [7] to improve forecasts while making the interventions explicit. Moss, Artis and Ormerod, in particular included in their system an explanations facility which provided the user with a qualitative account of the reasons for interventions by the rulebase. These qualitative reasons were couched in much the same language as that given by model operators for making similar interventions by hand.

There is therefore some precedent for including qualitatively expressed, domain expertise in models of social or economic processes and verifying the qualitative elements of such models through assessment by domain experts. A further development in this area is to integrate well verified theories from other disciplines into our computational models.

One such theory, used in two of the models reported below, is Alan Newell's unified theory of cognition [8]. The theory itself was guided by the requirement to mimic the time required by humans for cognitive acts of varying degrees of complexity. The Soar

software architecture [9] is an implementation of the Newell theory which performs well when assessed against the performance of subjects in a large number of psychological experiments.

Cooper, Fox, Farringdon and Shallice [10] showed that Soar is not the only possible implementation of the Newell theory. Moss, Gaylard, Wallis and Edmonds [1] found that reimplementing Ye's and Carley's Radar-Soar model [11] in SDML reduced the number of computations and the time required to run the model by two to three *orders of magnitude* while replicating the Radar-Soar results.

An important issue which remains is the extent to which further verification can be obtained by comparing numerical outputs from simulation models with appropriate statistical data series. The argument of this paper is that the verification of computational models with qualitative elements can and should include empirical tests of the behavioural elements of the models, assessments by domain experts and, when possible, statistical tests of the model's numerical outputs. The descriptions of the three models are used to investigate the limits to such verification for different applications of social simulation.

3 Relating Qualitative Domain Expertise to Statistical Evidence: An Intelligent Market Modelling System

A social simulation model which perfectly matched the criteria of the preceding section would represent individual agent cognition in a manner which was well validated by cognitive theory and independent experimental evidence, which was implemented in a declarative programming language corresponding to a known logical formalism and which produced numerical output series corresponding to some reliable statistical series generated by the modelled social interaction. The model reported in this section lacks the representation of individual cognition. Instead it incorporates the qualitative judgements of domain experts whose professional expertise involves some, usually highly informal and qualitative, assessment of the social conditions in which the customers for branded goods will seek one or another set of perceived brand attributes. These attributes have no objective meaning. Examples in markets for alcoholic beverages are "special", "traditional", "imported" (though usually produced domestically), "unique". The first models reported by Moss and Edmonds [12] were based on marketing practitioners' assessments of the different reasons why agents might buy such beverages. For one such market, the reasons given were that they sought the "physiological effect" — this was called "functional drinking" or that they were buying drink for social occasions such as parties or for self-reward or because they sought new and distinctive drinks. Moss and Edmonds devised a distribution function relating brand characteristics as specified by the marketing practitioners to the requirements of consumers in different social contexts. They were therefore called context-dependent attribute preference (CDAP) functions. The differences between CDAP functions and the utility functions used by economists is that utility functions relate to individual preferences and purport to mimic the decisions of individuals while CDAP functions describe the mean preferences, the importance of individual brand attributes and their tolerances to deviations from the ideal attributes for all consumers in particular social contexts. For example, the purchasers of beverages for a party will want something special but not very distinctive so that the beverages will be attractive to a wide range of tastes. On the other hand, the same individuals buying beverages to celebrate (say) a

career advancement will either want or certainly be more tolerant of distinctiveness in their purchase since it will be primarily for their own enjoyment. Specialness will still be an important consideration. Moreover, some attributes will be more important in some circumstances than others

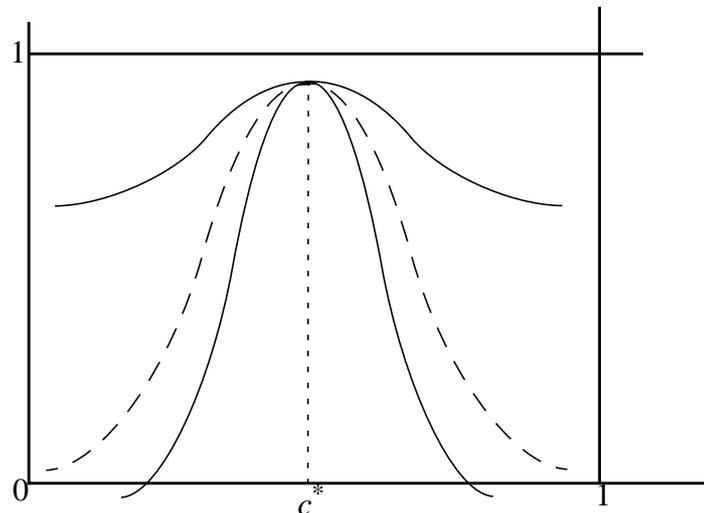


Fig. 1. The three functions (*upper, middle* and *lower*) relate the contribution to demand strength for the same ideal value of the brand attribute but are critical to different degrees. Because the values of the functions for each attribute of the brand are multiplied together in determining brand strength, the shallowest function is the least critical because deviations from the ideal reduce the contribution to strength by the smallest amount. The “variance is an index of how tolerant consumers in this social context will be to deviations of actual characteristic intensity from the ideal at c^* ”

The function which captured these *desiderata* is given in fig. 1 where the value on the vertical axis is the contribution to the strength of demand by consumers in the particular context and the horizontal axis is the intensity of the particular attribute associated with a brand. The parameters of the CDAP functions for each brand and each attribute are obtained from the marketing practitioners who enter the values on a five-point Likert scale. So the presence of an attribute could be taken to be “highly critical”, “moderately critical”, “not very critical”, “hardly matters” or “does not matter”. The ideal value is specified in similarly qualitative terms.

The model was constructed so that, given the various CDAP functions, including functions for the parameters “relative price” and “expensiveness”, the relative attractiveness of each brand together with the extent to which each brand differed in its attributes from other brands (described as inter-brand distances) was used to determine the market share of each brand given the prices of all brands considered.

In the first generation of these models, the marketing practitioners specified the contexts (social, self-rewarding, functional, *etc.*) and the characteristic ideals, tolerances and criticality of each brand attribute in each context. They also specified their estimates of the proportion of demand accounted for by consumers in each context. The model incorporating these judgements and estimates was then run over a subset of the

data available on prices and sales of each brand. Usually, the data was obtained from supermarket scanners for a geographical region such as the UK or metropolitan areas of the United States. A binary search algorithm was used to obtain the best fit in terms of smallest root mean squared errors (RMSEs) or minimum absolute percentage errors (MAPEs). The search algorithm changed the specified sensitivities of consumers to prices, to relative distances among brands and to the relative market strengths implied by the brand attributes relative to the CDAP functions. It then changed the proportions of demands accounted for by consumption in each context and, finally, it changed the ideal values of the various attributes in each context. The result was a set of modifications to the qualitative judgements of the domain experts together with statistical measures of the consistency of those judgements with the numerical data. The RMSEs and MAPEs obtained with these models was in every case far better at tracking market shares over the holdout set (the data not used in parameterizing the models) than were the best ordinary least-squares models.

A second generation of these market models reported by Edmonds and Moss [13] incorporated only the practitioner judgements of the important brand attributes in different markets and used a genetic programming algorithm to identify contexts, the relative importance of each context and corresponding CDAP function parameters to minimize RMSEs. The contexts were not given mnemonic names by the model but they provided an important input to the development by marketing practitioners of their own understanding of the markets that had been modelled.

Clearly, the CDAP functions are entirely procedural and there is no explicit representation of agent cognition in the model. In a model which aggregates context-dependent demands rather than agents, there is no immediate scope for representing individual agents, much less their cognitive processes. However, to extend these models to the development of new markets arising from new product introductions such as the so-called “alcopops” or the development of new classes of consumers in the developing or emerging-market economies, some representation of agent cognition would be essential to model emerging tastes and, so, demands for brand attributes in various social contexts. In such extensions to the models described here, the behaviour emerging from populations of software agents will also yield numerical simulation outputs amenable to statistical comparisons with the empirical record. In such cases, not only will we have well validated representations of cognition but also the models themselves, including their representations of qualitative phenomena, will be as well verified as statistical technique allows.

4 Relating Domain Expertise to “Statistical Signatures”: Modelling the Management of Critical Incidents

The model described in this section, and reported in detail by Moss *et. al.* [1], was devised to investigate the extent to which improved communication within an organization can prevent critical incidents from becoming full-scale crises. For these purposes, a critical incident is one which threatens to disrupt or actually disrupts normal operations but which is contained and resolved using the existing assets and procedures of the organization. A crisis is an interruption of the activities of the organization sufficiently extensive as to threaten its survival and which cannot be resolved with the existing assets and procedures of the organization.

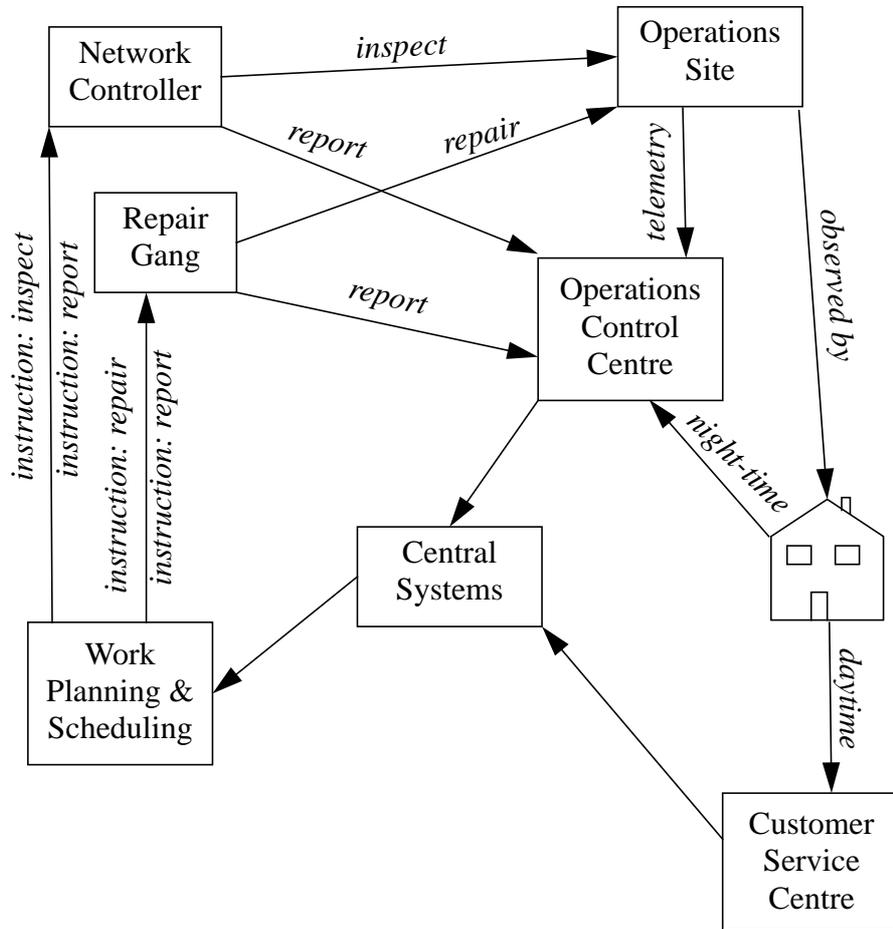


Fig. 2. The labels on the arrows describe the flows of information or actions among the various sources and repositories of information. The operations sites are not staffed and communicate by telemetry with the operations control centre. *Central Systems* and *Work Planning & Scheduling* are computer systems. The rest are of the boxes represent people or the persons grouped as departments or sections

This notion of the “statistical signature” as defined by Arthur *et. al.* [14] seems to be a statement about the visual appearance of a line chart. Though a useful notion, it is not sufficiently well defined to provide as clear a verification of a model as was obtained in the models of the previous section. In this section, a model is reported in which the output is characterised by a plausible series of outcomes conforming to the statistical signature associated with observations of the empirical referent of the model. The model also contains an explicit and well validated representation of agent cognition together with an accurate description of the relevant information systems and organizational structure of an actual company. The model is concerned with the systems and procedures for responding to critical incidents in the water and sewage services industry. Critical incidents include those which are likely to interrupt the provision

of these services to the public or which will cause environmental damage or pollution but are containable with the existing assets, systems and procedures normally available to managers of the company.

The systems and organizational structure of the company as they relate to critical incidents is depicted in fig. 2. The cognitive agents in the model are the network controllers and the operations control centre. The model was implemented in SDML with the container structure depicted in fig. 3. The model cycles over days and, within each day, 18 task cycles. At the end of every six task cycles, the network controller changes although the same three controllers are active in the same rotation each day.

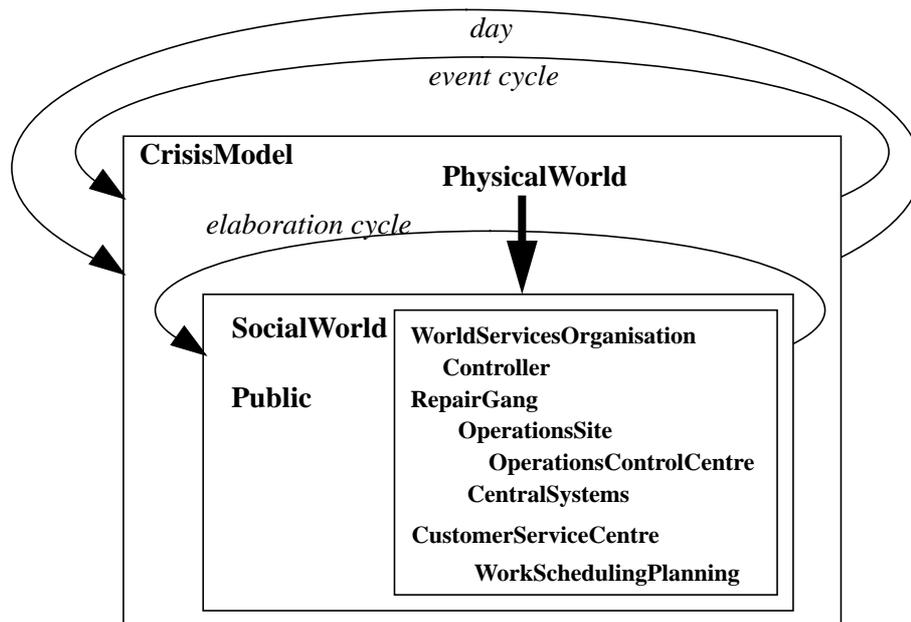


Fig. 3. The agent of type PhysicalWorld decides which, if any, primary events will occur spontaneously according to the specified probabilities and, if any such event does occur, assigns it to an operating site at random. If there are already events occurring at an operating site, the PhysicalWorld propagates consequential events at random from the specified probabilities and assigns them to the appropriate operations site. After the PhysicalWorld determines the state of the world, all of the agents in the social world fire their rules in parallel. They cycle over the time period *elaborationCycle* which is a step in elaborating mental models. After each cognitive of artificially intelligent agent determines any mental models it requires and takes such actions as are implied by those mental models, if is finished for that task cycle

The causes of specific events are not in practice known to the individuals involved in critical incident management until the manifestations of the incident have been observed and the situation analysed. Even then, they might make the wrong diagnosis. For these reasons, it would not be appropriate to post the causes of a particular incident to a database accessible by all agents. Consequently, the relevant clauses are asserted privately by the PhysicalWorld to its own databases and the fact of the occurrence of the event is asserted to the database of the operation site at which the event is to occur.

This assertion is achieved by the explicit addressing of the clause *eventOccuring event* where *event* is actually *fire*, *pumpFailure*, *contaminationIncident*, or the like. Once one such event has been allocated to the operating site, then with the appropriate probability all of the consequential events and their consequential events, *etc.* are also allocated at random to that site. In subsequent time frames, secondary consequences are allocated to that site with the specified probabilities in every event cycle while events with those secondary consequences continue. Events, once allocated, remain a feature of the site until they are remedied, if there are remedies, and the events which gave rise to them have been eliminated.

The operating sites (in practice mystified) recognize two kinds of event: telemetric and publicly observable events. When a site has had a telemetered event asserted to its databases, it sends a message stating that it has that event to the OperationsControlCentre. When a site has a publicly observable event asserted to its databases, it selects at random a percentage of households and asserts the occurrence of that event to their databases. In the simulations reported here, each household of 100 had a 10 per cent probability of being selected to receive such a message. Because the information contained in those assertions refers to actual information which is available selectively, once again explicit addressing of the assertions is appropriate.

The OperationsControlCentre agent forwards the telemetry and public reports to the CentralSystems agent who decides on the actions to be taken. The instructions to take these actions are addressed explicitly to the WorkPlanningAndScheduling agent who allocates the work by addressing the instructions of an agent of type RepairGang or Controller as appropriate. The reports by the repair gangs or controllers are addressed to CentralSystems agent. Repairs take the form of actions asserted by the repair gang to the operating site and then read from the operating site's databases by the Physical-World instance.

The cognitive behaviour in this model is by the instances of type Controller, and the workPlanningAndScheduling, OperationControl and CentralSystems instances. These agents learn by specifying and testing models which are held on their respective databases as private clauses — *i.e.* clauses which can be asserted by an agent only to its own database and read only by that agent.

4.1 Agent cognition: learning as modelling

Agent cognition is represented as a process of model generation and testing within a means-ends framework [15]. This approach has its origins in the cognitive science literature, classically Soar and, in a slightly different vein, Anderson's ACT-R [16]. Both rely on problem-space architectures which are in effect relationships between goals and the sub-goals needed to achieve those goals.

In the critical-incident model, the controllers build models relating remediable causes to consequences. They are assumed to know which individual events are causes of which other individual events but not the associated probabilities. Because of the differing probabilities of some events occurring as direct and immediate results of other events and some occurring as an indirect consequence and at subsequent times, it is not always clear which causal events are the most important to remedy first. The procedure they follow to formulate models is, in the absence of any applicable model, to generate a new model which postulates as a single cause the kind of event which is a cause of the largest number of the other observed events at the same site. For example, early in

one simulation run, two events occurred spontaneously at one of the operating sites: an intruder forced entry and the water pressure dropped. As a result, virtually everything else that could happen did happen. These events included a fire, a chlorine leak, a power supply failure, discoloured water, contamination and pollution, low water levels, no water to customers and a water taste or odour. The controller sent to inspect the site concluded that the key event to resolve was the presence of the intruder because among the observed events, more of them had intrusion as a possible cause than they had for any other causal event. The runner-up as the key event was fire which came second to intrusion only because one possible cause of a fire is an intrusion but fires do not cause intrusion.

The models were used by the controllers to identify and report to central systems the primary cause or causes of an incident. If, as a result of that report, the remedy applied on the instruction of central systems eliminated at least one of the events identified by the model, then the model was endorsed as having reduced the severity of the incident. If the result of applying the model was to eliminate all of the events covered by the model (*i.e.* all of the causes and all of the effects), then there was a further endorsement to that effect.

4.2 Results

The results of the simulations indicated that communication reduces the time required to resolve critical incidents but that events which are too complex to be resolved by the usual procedures are not affected by improving normal communications. In terms of outputs which correspond to observable statistics, the number of event cycles elapsing from the onset of an incident at an operating site until the absence of any events at the same site was recorded for every incident over the various simulation runs. When network controllers shared their successful mental models, the percentage of incidents resolved within two event cycles increased from some 50 percent to about 60 percent. The proportion that took more than two but not more than four event cycles increased from 5.9 to 7.65 per cent on average (with a confidence interval of 0.99). However, model sharing made no significant difference to the percentage of events that took more than eight event cycles to resolve.

Now the model was set up to yield very large numbers of incidents of greater and lesser complexity. The relative probabilities of the occurrence of each type of event were obtained from the company concerned. The absolute probabilities were much higher so that statistically meaningful results could be obtained relatively quickly with the available computational resources. By “much higher” here is meant a level which, if realized by the company concerned, would involve loss of licenses, jail terms for the directors and public enquiries. However, we did obtain a “statistical signature” derived from a well validated representation of agent cognition and an empirically accurate specification of company structure, systems and procedures.

Is this statistical signature in some sense accurate? We have no criteria by which to give an answer. The statistical signature, like the stylized fact, is clearly a useful notion but, for the sake of clarity and to give us confidence in the implications of our models, we should pin down what we mean by it in some formal or consistent way. One possibility is to show that models representing the time scales involved more realistically

can be parameterized to yield low RMSEs and MAPEs or similar measures of simulation accuracy and that condensing the time scales and either aggregating or reducing the numbers of agents have predictable effects.

5 Relating Emergent Behaviour to Non-Statistical Data: Modelling the Transition Economy of the Russian Federation

Using the same representation of cognition as in the critical-incident model, another statistical signature was sought in a project to capture qualitatively described events in the Russian Federation. In general, it is well known that enterprises are not paying their workers wages and inter-enterprise debt is growing very fast and inflation is substantial. There is an interesting question here in how agents learn to cope in such an unstable social environment. However, one problem with modelling such environments is that we do not find reliable data series emanating from them due, in large measure, precisely to their instability.

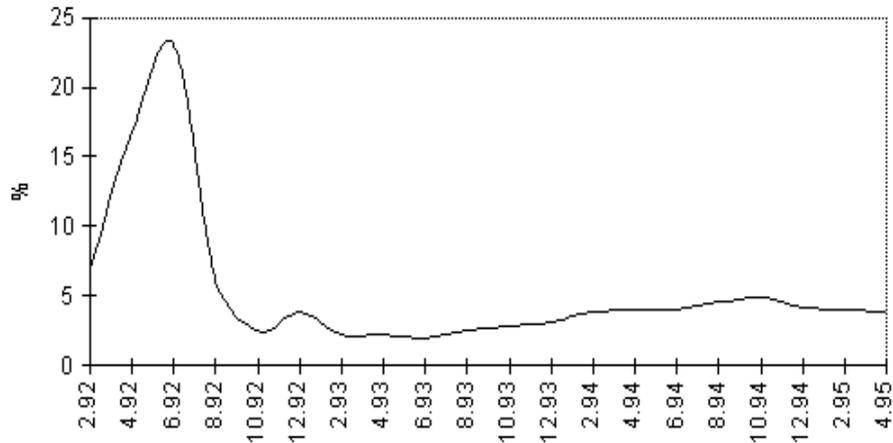


Fig. 4. The percentage monthly growth rate of inter-enterprise debt arrears in the Russian Federation from February, 1992 to April, 1995. Source: Finansy promyshlennosti Rossii, The Monthly Bulletin of the Working Centre on Economic Reform of the Government of the Russian Federation, no.1, June 1995, p.2

The particular period considered was that of the “arrears crisis” of 1992. The time pattern of the growth of inter-enterprise arrears is shown in fig. 4.

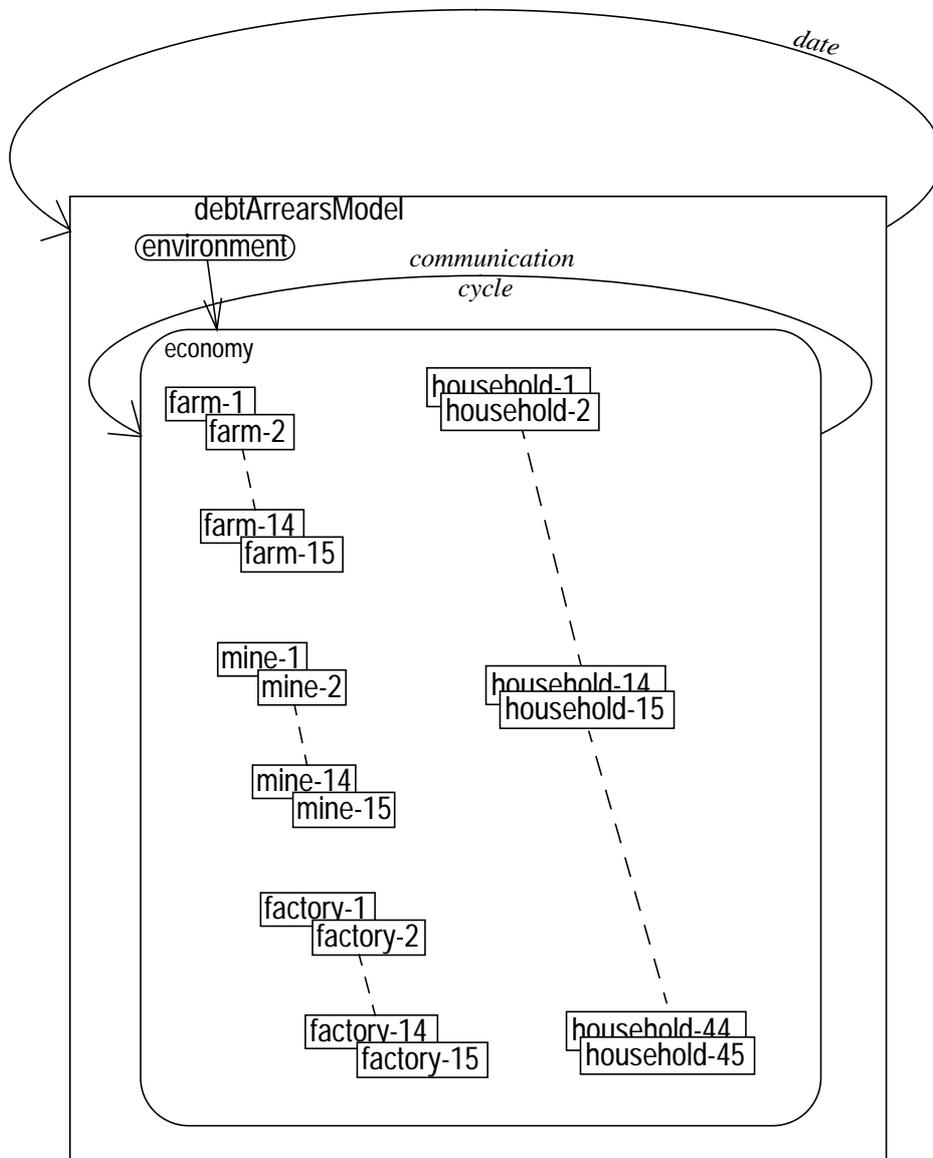


Fig. 5. The structure of the model of inflation and inter-enterprise debt arrears for the economy of the Russian Federation.

The evolution of the arrears crisis provides a clear example of how enterprises cope with situations characterised by significant uncertainty. In the Russian case, enterprises were forced to adopt a survival strategy giving priority to existing, recognized constraints. There was no possibility to maximize anything in the framework of those constraints [17]. Indeed, the scale of the accumulated debt and its persistence suggest that enterprises that debt reduction was not a high priority. For one thing, both the liquid assets of enterprises and their debts have been growing simultaneously. The debt

became an element of their survival strategy being instrumental in prolonging the existence of a business environment they were accustomed to, i.e., the one governed by soft budget constraints. Because the accumulated bad debt grew out of proportion on the national scale and became commonplace in all industrial sectors, this important business indicator ceased to be seen as a symptom of poor management efficiency.

There were three production sectors: factories producing outputs for both the mining and agricultural sectors, mines producing outputs for the manufacturing sector and farms producing outputs for the agricultural sector itself and for consumption by households. The model cycled over dates and within each date the agents in the economy cycled over communication cycles during which they made offers to buy and agreements to sell products

With this model, it was by no means difficult to generate rapidly growing inter-enterprise debt arrears. In fact, all models had these arrears but the early models had flat price, employment and output series. It was clear that some important aspect of the cognitive processes of enterprise managers was being missed. The mental models of the agents were generated initially by random combinations of actions and implied consequences for the goals of the enterprises which included sales volumes and cash holdings. Because of the extent to which the modelled economy captured the instability of the actual Russian Federation economy, there were no signals from this generate-and-test approach to developing mental models to identify a small set of models which would inform goal-enhancing behaviour. Noise but no signals were being provided by the economy itself.

We then looked for other sources of information which we found in the transactions process itself. The information available to enterprises included the prices it was being charged by its suppliers as well as the supplier's record of filling orders for their outputs. They also knew which of their own customers were paying for the goods they acquired and the prices they charged. A natural hypothesis was that enterprises would imitate the behaviour of the most successful enterprises known to them. The instability and limited information sources and general lack of reliability of information in the economy indicated that the best source of information was observation and the observations were of the suppliers and customers of the enterprise. Paying bills without holdings of cash was not a possibility even if this were a mark of success in customers. But raising prices was always possible and this was the behaviour that emerged. A result of assuming that mental models were informed by direct observation, the simulation model generated a volatile price inflation series with a strong upward trend. It turned out, though we did not have the published inflation figures at the time, that the official price series is indeed marked by the sort of volatility we found. Once again, and acknowledging its defects, the notion of the statistical signature provided a useful target for social simulation modelling.

6 Conclusion

All of the models described here represented agents or their behaviour in ways which were validated independently of the models. The behaviour of consumers in the intelligent market modelling system conforms to the independent views of marketing practitioners. The representations of cognition in the critical-incident and transition-economy models conformed to important, common aspects of the main software architectures corresponding to experimentally verified theories of cognition. In all cases, the

models captured empirically observed characteristics of the agents and their environments. The differences in the verification of each model was that we had exceptionally good statistical data with which to verify the market models by means of standard statistical measures whereas the other two models relied on stylized facts and statistical signatures which entail no clear formal criteria. The reasons were different in each case. In the critical-incidents model, the condensation of time gave us more useful results even though, with some expense, it would have been possible to obtain data about the actual duration of critical incidents. This would be a good case in which to develop techniques for mapping numerical social simulation output into actual statistical data. It has not, as far as I know, been attempted in the social simulation domain. In the case of the arrears model, there is no reliable statistical evidence so that all we have to rely on are stylized facts or statistical signatures. Validation of the model by appeal to domain experts and perhaps survey evidence concerning sources of information and the extent to which enterprise managers rely on different sources of information will doubtless help to validate our representations and verify qualitative model outputs.

Clearly the models discussed in this paper are distinguished by their role as representations of actual systems: markets for given sets of competing brands, a set of functions of a particular company, specified empirical problems facing a much-studied if little-understood economy. There are many models in this vein, though not all are concerned with management and economic issues. Anthropological simulations by Doran and his colleagues or models of emergent behaviour undertaken by Cress at Surrey, or the simulations of the emergence of altruistic behaviour by Pedone and Parisi are examples of relatively abstract social simulation models concerned with empirical phenomena not directly related to either economics or management issues.

Social simulation frequently entails emulation of real phenomena. The point of this paper has been partly to suggest some means of making the links between the models and their empirical referents clear and, where possible, measurable; but more importantly to put a case for the importance of relating social simulation models explicitly to the empirical phenomena. If this case were accepted, then the development of procedures for making those links in ways that avoid mere handwaving and unsupported assertions of relevance is by implication an important thread in the continuing development of the social simulation literature.

References

- [1] Moss, Gaylard, Wallis and Edmonds (1998), "SDML: A multi-agent language for organisational modelling", *Computational and Mathematical Organization Theory*, in press.
- [2] Konolige, K. (1988), "On the relation between default and autoepistemic logic", *Artificial Intelligence*, v. 35, pp. 343-382.
- [3] Armstrong, J.S. (1978) Forecasting with econometric methods: Folklore *versus* fact", *Journal of Business*, v. 51, pp.
- [4] Bunn, D. and Wright, G. 1991 "Interaction of judgmental and statistical forecasting methods: issues and analysis", *Management Science*, v.37, pp.501-518
- [5] Turner, D.S. (1990), "The role of judgment in macroeconomic forecasting", *Journal of Forecasting*, v. 9, pp.315-345.
- [6] Collopy, F. and Armstrong, J.S. (1992), "Rule-based forecasting: Development and validation of an expert-systems approach to combining time-series extrapolations", *Management Science*, v.38, pp.1394-1414.

- [7] Moss, S., M. Artis and P. Ormerod (1994), A Smart Macroeconomic Forecasting System. *Journal of Forecasting* 13 (3), 299-312.
- [8] Newell, A.(1990), *Unified Theories of Cognition*, (Cambridge MA: Harvard University Press).
- [9] Laird, J.E., A. Newell and P.S. Rosenbloom (1987), "Soar: An architecture for general intelligence", *Artificial Intelligence*, v. 33, pp. 1-64.
- [10] Cooper, R., J. Fox, J. Farringdon and T. Shallice (1996), "A systematic methodology for cognitive modelling", *Artificial Intelligence*, v. 85, pp. 3-44.
- [11] Ye, M. and K.E. Carley (1995), "Radar Soar: Towards an Artificial Organization Composed of Intelligent Agents", *Journal of Mathematical Sociology*, v. 20, pp. 219-246.
- [12] Moss, S. and B. Edmonds (1997), A Knowledge-based Model of Context-Dependent Attribute Preferences for Fast Moving Consumer Goods, *Omega*, 25, 155-169.
- [13] Edmonds and Moss (1997), "Artificially Intelligent Specification and Analysis of Context-Dependent Attribute Preferences", Centre for Policy Modelling Technical Report 97-28
- [14] Arthur, W.B., J. Holland, B. LeBaron, R. Palmer and P. Tayler (1996), "Asset Pricing under Endogenous Expectations in an Artificial Stock Market", University of Wisconsin-Madison, Social Systems Research Institute 9625.
- [15] Moss and Edmonds (1997), "Modelling economic learning as modelling", *Cybernetics and Systems*.
- [16] Anderson, J.R. (1993), *Rules of the Mind* (Hillsdale NJ: Lawrence Erlbaum Associates).
- [17] Aukutsionek, S. (1995) "The Market Reform and Transitional Crisis", *The Russian Economic Barometer*, vol.4, no.1, pp.22-41.