The Cost of Rational Agency

Scott Moss*
Centre for Policy Modelling
Manchester Metropolitan University
http://www.cpm.mmu.ac.uk/~scott
s.moss@mmu.ac.uk

The rational agency assumption limits systems to domains of application that have never been observed. Moreover, representing agents as being rational in the sense of maximising utility subject to some well specified constraints renders software systems virtiually unscalable. These properties of the rational agency assumption are shown to be unnecessary in representations or analogies of markets. The demonstration starts with an analysis of how the rational agency assumption limits the applicability and scalability of the IBM information filetering economy. An unrestricted specification of the information filtering economy is developed from an analysis of the properties of markets as systems and the implementation of a model based on intelligent agents. This extended information filtering economy model is used to test the analytical results on the scope for agents to act as intermediaries between human users and information sources.

* This paper arose from a series of discussions with Ed Durfee, Aaron Armstrong, Moshe Tannenholz, Michael Schroeder and Leigh Tesfatsion, none of whom necessarily agree with the argument or my conclusions. I am grateful as always to my colleague Bruce Edmonds with whom everything in the paper has been discussed. The intellectual imporatnce of SDML will be apparent from the paper. SDML has been designed and developed by Steve Wallis in the VisualWorks development environment. It also runs on top of the VisualWorks 2.5.x object engine. The development environment and permission to distribute the object engine free of charge for academic purposes are provided by our sponsors, ObjectShare (UK) Ltd.

1. Introduction

There is an understandably naïve view among computer scientists in the multiagent systems community that economic theory is about economic processes and that economists have demonstrated that and how competitive market mechanisms ensure efficient outcomes in the sense that it would not be possible to reduce the cost of one activity or increase the satisfaction of one agent without increasing the cost of some other activity or reducing the satisfaction of some other agent. This view is false and its consequence is a clear misdirection of effort in at least one area of multiagent systems research: the development of agent-based, information filtering economies to find and report to human users relevant information from databases or the Internet.

The purpose of this paper is to demonstrate that the effort is misdirected and that to hold to long-standing, demonstrably effective principles of software engineering better enables the power of the agent paradigm to be realised.

The factual position regarding the nature of economics and the reasons why economics sets a bad example for computer science are explained in the section 2. An example of the consequences of adopting economists' methods is presented in section 3. A constructive alternative approach is offered in section 4 with a demonstrator implementation described in section 5. The results from that implementation are reported in section 6. The wider lessons to be drawn for multiagent systems research more generally are argued in section 7.

2. How economics sets a bad example.

The standard computational modelling procedure in economics is to specify an arbitrary representation of one or more agents, an arbitrary specification of the relationships among agents and then to assess the goodness of the model by how well the results conform to some unobserved equilibrium.

A seminal work in this vein was Thomas Sargent's *Bounded Rationality and Macroeconomics* [10]. Sargent addressed an implausibility in a core concept in conventional economics called rational expectations. A rational expectations equilibrium is one in which every agent in the system knows the correct model of the system. The common application is to macroeconomics so that everyone in the economy knows the forecasting model that will correctly predict the future values of all macroeconomic variables. It is usually assumed in addition that every agent has the same utility function – this being the "representative agent" assumption that diversity of preferences or expectations within the economy have no significant or biasing effect on its performance in aggregate. This assumption is wholly unsupported by any evidence whatsoever. On the basis of Sargent’s, there is a growing literature in which agents are represented by some implementation of a genetic algorithm and the goodness of the models incorporating this representation is measured by the closeness of its output to a rational expectations equilibrium.

The important point here is that both the agent specifications and the measure of goodness of the models are not at any point considered in relation to any observations of the real world.

The consideration of markets in modern economics has a much longer history. In the 1880s, Leon Walras devised the general equilibrium concept in which there is a set of prices such that the demands for and supplies of every good and service produced and traded in the economy were equal. There was no excess supply of any good or service and no excess demand. The process by which this would generally be achieved has never been identified. Walras conjectured that such an equilibrium could be found by a process of *tâtonnement* in which an auctioneer would receive all bids for and offers of goods and raise the prices of those goods for which the bids exceeded the offers and reduce the prices of those goods for which the offers exceeded the bids. This process is unstable whenever there are more than two goods and two traders in the economy. A recent attempt by Binmore *et al.* [2] to devise a game theoretic model was claimed to be interesting because software agents face similar problems when they meet other software agents that are "badly programmed" meaning that they are not rational in the sense that economic agents are assumed to be rational. The agents in that game were to be considered "as a metaphor for an evolutionary process in much the same way that the auctioneer of neo-classical economics is a metaphor for some unmodeled process that eventually equates supply and demand."

This is the crux of the issue. Whether in game theory or in market models, the agents and the processes specified in economic models are metaphors for some unjustified impression invented by the modellers out of their own heads. These arbitrarily specified agents are engaged in arbitrarily specified processes in the precise sense that both agent and process specifications are entirely unvalidated. What is more, verification takes the form of nothing more than ensuring that sufficient (but not necessary) conditions exist for an equilibrium that itself has no empirical correspondence to any observations whatever. Economic theory is a collection of metaphors for unmodelled (because unobserved) processes and unobserved agents.

While these metaphors have been translated into a number of areas of multiagent systems research, the potential consequences of such translations are readily demonstrated by one example. The example developed in this paper is that of the information filtering economy in which automated information agents buy and sell information ultimately to provide human users with sets of information tailored to their individual interests and requirements. The inspiration for these information markets is the partial equilibrium theory of microeconomics.

In this theory, competitive markets are defined on a set of agents who demand a good or service and mutually exclusive set of agents who supply the good or service. All agents on both the demand side of the market and the supply side of the market are price-takers. That is, they observe the market price which is the independent variable of their demand and supply functions, respectively. If there is no production, then the demand and supply functions are derived from utility functions. The market is cleared (demands and supplies are brought into equality) by effectively instantaneous price adjustments such that, in the stable cases, a price increase reduces excess demands and a price reduces excess supplies.

Of course, we do not observe many such markets and those we do observe are the organised stock and commodity exchanges that invariably have very strict rules of behaviour and stylised mechanisms for the exchange of assets. There are in these markets both jobbers and brokers where the jobbers buy and sell the assets at prices they specify while brokers arrange transactions between sellers and jobbers and between buyers and jobbers but do not themselves buy and sell. Other markets patently do not work that way. Commodities that are specialised and produced for thin markets (such as ships) are only produced to order and are never produced to hold as inventories. The price of each unit produced is negotiated in advance. The same is true for specialised buildings such as sports stadia but less true for buildings that can be well standardised and for which markets are densely populated such as houses on housing estates which might differ only by colour. Chocolate bars, domestic soaps and other fast-moving consumer goods are sold by producers to wholesalers who sell them to retailers who sell them to the users. Large retailers can bypass the wholesalers.

Evidently, common observation indicates that the processes of exchange for different commodities are simply different.

There are three specific questions to be raised in this regard:

1. Is there any reason to suppose that information is, in some sense relevant to exchange, exactly like stocks, bonds and futures contracts?

2. What are the systems requirements to support market-style buying and selling of information?

3. Are the agent representations of economic theory appropriate models for information agents.

The first of these questions is left to information scientists to answer. A partial answer and a clear methodology are offered for the second question on the basis of an unambiguous and constructive answer in the negative to the third question.

3. Some consequences of adopting agent representations from economics

The information filtering market devised and implemented by Kephart and his colleagues at IBM [4, 5] provides a good example of the costs of assuming rational agency in the sense of conventional economics. It is a good example because of the clarity with which both the implementation of the model land its restrictions are expressed in the reports of the implementation.

The objective of the IBM model is to inform a process whereby "agents will progress naturally from being mere facilitators of electronic commerce to being financial decision-makers, at first directly controlled by humans and later with increasing autonomy and responsibility". The rationale for adopting the perspective of economics is that

> After all, economic mechanisms are arguably the best known way to adjudicate and satisfy the conflicting needs of billions of *human* agents. It is tempting to wave the invisible hand and assume that the same mechanisms will automatically carry over to *software* agents. However, automated agents are *not* people. . . . How might these differences affect the efficiency and stability of future information economies? [5, p. 1]

The truth of this statement would not conflict with the historical and factual position that the representation of economic mechanisms by economists has no empirical correspondence with observed economic mechanisms or their consequences. It follows that this statement does not justify the assumption of rational agency as specified or inspired by economics. Of course, just because economists have invented "economic" mechanisms to analyse does not mean that those mechanisms would not be useful if implemented as software systems for information filtering. That they are not useful in this regard is indicated by the IBM team's reports of their model of an information filtering economy.

A schematic description of the IBM model is given in Figure 1.

In this model,

- the information source issues one article per time step where an article is represented by an integer drawn at random from the sequence 1...J and where each integer is interpreted as an information category

- brokers ($\mathbf{B_b}$) pay a fixed price to the information source for the current time step's information and a "transportation charge" for each consumer told about the information.

- consumers ($\mathbf{C_c}$) pay brokers for the information they buy after being told about it and computation cost to decide if they want the article.

Both brokers and consumers are rational in that each has a subjective expected utility function. The expected utility of a broker is the probability of purchasing a published article *times* expected net revenue per article in the category specified by the integer representing the article. The expected utility of a consumer is the probability of being offered an article in published category *times* the utility that would be derived from such an article net of purchase and computation cost.

Figure 1: The Structure of the IBM Information Filtering Model

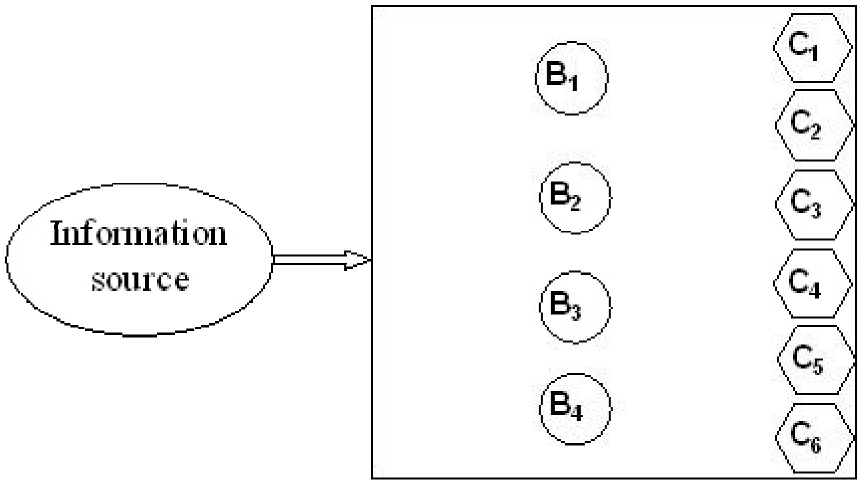

The authors point out four restrictions imposed on the model that, presumably, are to be relaxed to capture usefully scaled implementations of information filtering economies. These are

- The model disregards articles that fall into multiple categories

- The model disregards the possibility of multiple articles containing redundant information

- The model disregards differences among consumers and among articles

- The model disregards the advantages and difficulties associated with automatic evaluation of articles' semantic content

It seems reasonable to ask why, if these restrictions are undesirable, they have been built into the model at all? Though the authors do not say so, one reason will surely be tractability. The utility functions chosen for brokers and consumers, respectively, facilitate the derivation of analytical results about (a) equilibrium prices for different numbers of categories, given specified probability distributions of the occurrence of each category or (b) the optimal number of categories for a broker to offer in a given price regime. It is usual in these cases for an increase in the scale of the problem to lead to intractability or uncomputability (e.g.,[1]). To complicate the utility function by allowing articles falling into multiple categories would require a utility function the arguments of which included some representation of information contents and/or joint probability distributions of the occurrence of the various categories of articles. The value of any article when the same information could be included in articles of several categories would again involve the determination of the expected utility from any one article depending on the probability of the occurrence of other articles. The main problem of allowing consumers to have different forms of utility function is that the outcomes depend on the distribution of possible expenditure (usually due to the distribution of income and wealth). Consequently, the determination of an equilibrium requires the simultaneous determination of each consumer's expenditure ability together with, in this case, the prices and number of categories of article determined by each broker.

In all of these cases, as in economic theory, agents have to know the actual state of the world (or market) and the actions (or even more strongly, the utility functions) of every other agent. In order to relax these assumptions, the IBM team employed simulation techniques that indicated the *system* outcomes to be sensitive to the assumed initial conditions.

In summary, the IBM model of an information filtering economy entails known restrictions in order to render it tractable while adding an element of realism renders the model both analytically intractable and sensitive to initial conditions. This is a consequence not only of the rational agency assumption but the model-building procedures required by that assumption.

In building models with rational agents, there are inherently restrictions on the specification of the systems in which they are implemented. While the specification of any agent will entail some restrictions on the systems, the restrictions imposed by rational agency as evidenced in the IBM information filtering economy are undesirably (and unnecessarily) strong. An alternative is to start with the specification of the desired system and implement agents that can function usefully in that system. In the case of markets, a weak specification of agents has been known for twenty years [7, 8] to support some very general and robust analytical results on the institutional structure of markets. These results are exhibited and applied to an information filtering market in the next section.

4. Intermediation in information filtering markets.

An information filtering economy is a system in which software agents buy and sell information more or less autonomously on behalf of human users. Adopting the analogy of the market as observed in relatively uncontrolled economies, it is appropriate to allow but not to impose intermediary agents who buy and sell information but do not themselves use it (jobbers) or who neither buy nor sell on their own account but who arrange transactions between buyers and sellers (brokers). In the IBM information filtering economy the intermediaries called brokers do buy and sell information. Provided the meaning is clear, there is no harm in continuing that practice in this paper.

The first issue raised in this regard concerns the conditions in which any intermediary can function in a market. In order to get a handle on this question, the weakest possible assumptions are offered about agents. It is assumed that

A1 agents can only survive in the market if they are able to pay for the information they buy

A2 agents prefer survival in the market to exit from the market

In the real world, agents can continue to pay their bills over significant lengths of time only if their receipts are not less than their disbursements. Ignoring lending, borrowing and the associated charges and incomes, this amounts to setting a price that, on average, exceeds the unit cost of selling including both production costs (if any) and the costs of engaging in the necessary transactions. Consequently, assumption A1 is a statement of fact about all markets restated in relation to information markets in particular. Assumption A2 is the fundamental behavioural assumption which reflects observed business behaviour.

Taken together with elementary accounting identities on cash flows, assumptions A1 and A2 have the following consequents:

C1 agents will always seek to reduce the costs of their purchases if to do so entails no losses of information quality or timeliness

C2 agents will always seek to increase their sales revenue if to do so involves no additional expenditure

Lower costs and increased revenues imply larger net cash flows to provide a cushion of assets to draw down on occasions when receipts fall short of disbursements.

These two assumptions and the two consequents imply the *fundamental law of market intermediation*:

Intermediaries (brokers and jobbers) can function profitably in a market if and only if the spread between their bid and offer prices is less than the savings on transactions costs afforded to their suppliers and customers.

The demonstration of this law could hardly be simpler. The notation is:

p_o	The price at which an intermediary offers an item (the offer price)
p_b	The price for which an intermediary bids for an item (the bid price
$p^{(d)}$	The price in direct exchange (paid by a user to a producer without the intervention of an intermediary)
$t_u^{(i)}$	The transactions costs incurred by a user in purchasing an item in intermediated exchange
$t_u^{(d)}$	The transactions costs incurred by a user in purchasing an item in direct exchange
$t_s^{(i)}$	The transactions costs incurred by a source in selling an item in intermediated exchange

$t_s^{(d)}$ The transactions costs incurred by a producer in selling an item in direct exchange

t_o The transactions costs incurred by an intermediary in selling an item

t_b The transactions costs incurred by an intermediary in buying an item

It follows from consequent C1 that an agent will purchase from an intermediary only if the total cost (offer price and transactions costs) of buying from the intermediary are less than the price and transactions costs in direct exchange with a source. That is,

(1) $p_o + t_u^{(i)} < p^{(d)} + t_u^{(d)}$

From C2, an agent will sell to an intermediary only if the price net of transactions costs is greater than in direct exchange. This implies

(2) $p^{(d)} - t_s^{(d)} < p_b - t_s^{(i)}$

Combining these inequalities,

(3) $p_o + t_u^{(i)} - t_u^{(d)} < p^{(d)} < p_b - t_s^{(i)} + t_s^{(d)}$

Applying A1 and A2 to the intermediaries and solving in expression (3) for the intermediary's bid-offer price spread,

(4) $t_o + t_b < p_o - p_b << \left(t_u^{(d)} + t_s^{(d)}\right) - \left(t_u^{(i)} + t_s^{(i)}\right)$

which states in effect that any intermediary's transactions costs must be less than the bid-offer spread which must be less *a fortiori* than the *total* savings in transactions costs by sources and users as a result of the activities of the intermediary. Consequently, total transactions costs are smaller if intermediaries can survive in a market than if they are absent. Since the bid-offer spread can only appropriate a part of those savings, the total costs of trading in any market will be reduced by the profitable presence of intermediaries.

This reasoning leads to the following *hypothesis on intermediary survival*:

> *In any system involving transactions, intermediaries can survive if and only if total expenditure will be lower in their presence than in their absence.*

5. An extended information filtering economy model

A specification of the conditions in which intermediaries can survive is necessary to test the hypothesis on intermediary survival. It is evidently necessary and sufficient that intermediaries can reduce transactions costs in ways that are not available to sources and users. The natural supposition here is that intermediaries can achieve economies of scale that individual sources and users cannot. Not only is this a natural supposition based on common observation, it is one that is easily accommodated by a suitable extension of the IBM information filtering economy.

The extension required is one that does not suffer the restrictions imposed and acknowledged by the IBM team. That is, it will have the following characteristics:

- Information packets (such as newsgroup articles or composite data structures) will fall into multiple categories in the sense that different users will seek different collections of information and the same information packet will be of interest to different users because of different items of information contained by it.

- Information packets can contain redundant information in the sense that the same item of information can appear in several information packets.

- Consumers and information packets differ in that different consumers use different collections of information and each packet contains a unique set of items of information.

- The representation of information supports representations of automatic evaluation of the semantic content of the information packets and does so at a scale that would accommodate existing and likely procedures for (say) natural language processing.

5.1 Model structure

The structure of the extended information filtering model is schematised in Figure 2.
The roles of the agents in this model are:

- The information generator is a non-cognitive agent that stores a private digit string of arbitrary number base and length.

- Information sources (S_i) are non-cognitive agents that access the integer values at specified positions in the digit string at major time step or *trading cycle*. The set of digit string positions and corresponding values held by a source is an *information packet*.

- Brokers (B_b) are cognitive agents that acquire the information packets of various chosen sources and offer to customers (C_c) or to other brokers the values at the digit string positions they demand. That is, brokers can "break bulk" so that, having acquired one or more information packets from source(s), they can supply exact demands to other agents.

- Customers (C_c) are cognitive agents that either acquire packets offered by the sources or enquire of brokers for the values of digits at demanded positions.

- Cognitive agents can acquire from information sources only the whole of their data packets at the current trading cycle but they cannot access only some of the information. This is analogous to the holding of whole data records or the storing of whole newsgroup articles or web pages by information sources.

There are two sources of transactions costs: transportation and processing costs. Following Kephart *et al*. the processing costs are proportional to the number of items of information received. Extending Kephart *et al.* to take location into account, transportation costs in this paper are assumed to be proportional to the distance between the supply point (source or broker) and the acquisition point (broker or customer). This requires each source, broker and customer to be given a location on the unit square as depicted in Figure 2 with sources distributed at random on the left hand edge of the square, customers distributed at random on the right hand edge and brokers distributed at random inside the square.

Figure 2: The Structure of the Extended Information Filtering Model

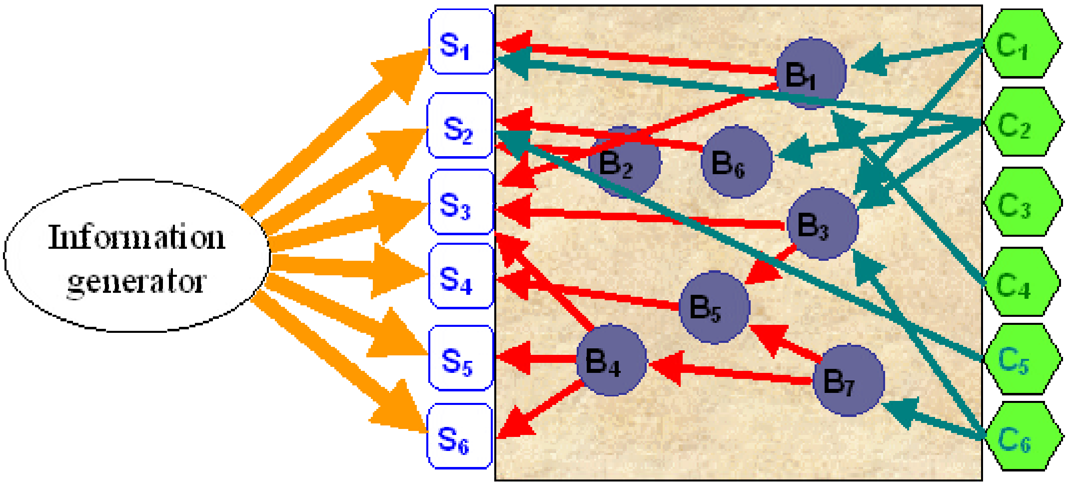

Exchange cannot take place without communication between buyer and seller. The model was implemented in SDML, the strictly declarative language that supports representations of communication by agents to assert clauses to one another's' databases. It also supports nested time levels and, within time levels, agents can act in parallel. In order not to change the state of the world for one agent by virtue of the actions of another agent, no clause asserted by one agent will take effect for any other agent until the next time step. Consequently, if *agent-i* asserts the clause (demandEnquiryFrom agent-i [3 7 2 6]) to the database of *agent-j*, then *agent-j* will not be able to access that clause until the next time step.

All cognitive agents in the model are synchronous parallel agents. To enable them to communicate in order to seek, agree and effect transactions, within each major time step, the trading cycle, is a nested time step, the *communication cycle*. A limit of six communication cycles were allowed per trading cycle though there were fewer communication cycles if all demands were filled earlier.

5.2 The representation of information

The information set in the extended information economy model is a digit string of arbitrary fixed length and arbitrary number base. This digit string is held by the information generator and is mutated at the start of each trading cycle. There is a fixed probability of mutation at each position of the string bounded from above by a user-selected maximum probability. These mutation probabilities are determined by a hyperbolic tangent transformation of the U[0,m) probability distribution where m is the maximum mutation probability. This transformation makes it possible to bunch all mutation probabilities about half the maximum or to make the realised distribution uniform or to bunch the probabilities close to zero and close to the maximum.

Each source was allocated at random a number and selection of positions on the information string, In each trading cycle it stored a pairlist of indices of the string and the current value at the indexed position. The digit string length, its number base, the number of sources and the maximum number of items of information that could be held by each source are specified by the model operator for each run.

5.3 Agent specifications

Each information user is defined by a location as indicated above, by demand for a set of information items corresponding to positions on the fixed-length digit string of the information generator, by *endorsement schema* as described below and by a set of rulebases and corresponding databases. A broker is also defined by location, by *endorsement schema*, rulebases and databases but not by demands. Brokers have rules to communicate willingness to supply items of information but information users do not. Apart from the absence of intrinsic demands and the presence of rules to support sales of information, there is little to distinguish the broker agents from the user agents. Utilising the object-oriented features of SDML, brokers are instances of type ExchangeAgent, users are instances of type InformationCustomer and both of these types are subtypes of type InformationAgent.

The endorsements schema are the key distinction between the agents implemented in the system reported here and rational agents. The endorsements schema are an alternative to utility functions. The principal conceptual difference is that an endorsements scheme is a framework for agents to reason and learn about preferences for different objects whereas utility functions represent static and unchangeable preferences by agents.

An endorsements scheme links mnemonic tokens to numerical values with a basis of comparison among different values. Clauses defining endorsements schema of information users in the extended information filtering economy model include:

> endorsementSchemeDefinition exchangeAgentEndorsementScheme [[lowCost 1] [highCost -1] [reliableSupply 2] [unreliableSupply -2] [accurateInfo 3] [inaccurateInfo -3]] 1.2\
>
> endorsementSchemeDefinition infoSourceEndorsementScheme [[lowCost 1] [highCost -1] [mostCompleteInfo 2] [accurateInfo 3] [inaccurateInfo -3]] 1.2\

The first argument in each case is an instance of type EndorsementScheme. The second argument is the pairlist of tokens and corresponding values. These tokens are associated with objects such as agents or information items or information sources by rules. The third argument is the basis of determining the endorsement value of an object – the combined value of all endorsements associated with the object. The *endorsement valuation function* is

(5) $$E = \sum_{e_i \geq 0} b^{e_i} - \sum_{e_i < 0} b^{|e_i|}$$

where the e_i are the values of the endorsements on the object and b is the third argument of the endorsement scheme definition clause.

An example of a rule endorsing an object is reproduced in Figure 3 from the main rulebase of type InformationAgent. The antecedents are satisfied if it is the first communication cycle of the trading cycle and if there was a purchase of information items from a supplier during the previous trading , items (unified with ?content) were supplied during the same trading cycle and there were no items purchased but not actually supplied. If all of that is true, then the clause endorsing that supplier as reliable is asserted to the database corresponding to the time level *tradingCycle* so that it does not have to be asserted each communication cycle.

Figure 3: Rule endorsing a supplier as reliable

```
and
    time tradingCycle ?tc\
    last tradingCycle (and
        purchaseFrom ?supplier ?indices ?price\
        infoSuppliedBy ?content ?agent)\
        not
            (and
                includes ?indices ?index\
                not
                    includes ?content [?index ?value])\
---
all tradingCycle (endorsementFor ?supplier reliableSupply [?tc])\
```

Both utility functions and endorsement mechanisms relate to agent preferences. The difference is that a utility function specifies a map of possible outcomes and preference relations among them while the endorsements mechanism is a process of developing preferred courses of action based on and changed by experience. It is not surprising that to impose fixed preference maps on agents and an algorithm to be used for making decisions whatever the prevailing circumstances, will restrict the domain of application of the system to a class of states in which the algorithm is applicable or tractable and exclude all other states. By defining a process that is itself contingent upon the state of the system, the endorsements mechanism can respond to emergent system phenomena as demonstrated in the extended information filtering economy model.

The endorsements mechanism generates emergent preferences among suppliers based on the behaviour of the suppliers and the nature, for example volatility, of the information they offer.

More direct means were used to ensure that agents did not buy from high-priced brokers. In the initial trading cycle there were no brokers and information users continued to search among information sources until all available information items in their lists of demands had been acquired. The sum of the transportation and processing costs of those acquisitions were retained in permanent memory. If, in any subsequent trading cycle, an agent received an offer of supply from a broker at a price which made the items wanted more expensive than the cost at which they were previously obtained from an information source, then the broker was endorsed as being a high-cost supplier and the current value of the information would be obtained from the source instead.

At the same time, brokers would seek to obtain high prices by building models of the effects of their mark-up strategies. The mark-up for each broker was applied to the cost of supplying the information items it expected to sell on the basis of demand enquiries in hand. When demand was strong, the mark-up would typically rise and when demand was weak it would fall. As will be seen from the simulation results, even the most successful brokers would occasionally see their prices undercut by a new, strategically placed broker so that they lost all custom and rarely recovered. In setting prices, each broker would note the maximum price the customer claimed it would pay and set an actual price per item as a random linear combination of the mark-up on anticipated costs and the maximum price.

6. Simulation results: testing the intermediary survival hypothesis

There are no production costs in the model as specified for this paper. Consequently, all costs are transactions costs resulting from transportation and processing as modified from the IBM model.

The test of the hypothesis on intermediary survival is whether intermediaries survive when their presence reduces total information costs and not otherwise. From the analysis of section 4, intermediaries can survive in the market only if there are economies that they can enjoy that cannot be enjoyed by either sources or users. The special advantage of brokers in the extended information filtering economy model is that they can break bulk (which the sources cannot) and they can sell information (which the users cannot). Consequently, brokers will be able to reduce the transactions costs of the users by selling them the information they want and not requiring them to process information to determine whether they want it. Also, brokers will be able to reduce transportation costs by locating close to the users and incurring the bulk of the transportation costs once and then lesser transportation costs on to each of the users rather than multiplying all of the transportation costs from close to the sources.

One would anticipate that location would be relatively important when transportation costs dominate transactions costs and that location would be relatively unimportant when processing costs dominate.

In the original publications of the markets analysis described in section 4, Moss [8, 9] argued that intermediation would also require the number of intermediaries to be small in relation to each side of the market. Common observation indicates that large chains of supermarkets and department stores, for example, purchase directly from producers while smaller, independent shops purchase from wholesalers. If this argument is correct, reducing the number of information users or the number of information sources in the extended information filtering economy model would reduce the scope for intermediation by reducing the scope for savings in information costs.

The parameters selected for the various simulation runs were set to test these hypotheses. Because the information sources are not cognitive agents, there was no point in varying their properties. The main aspects of comparison involved different transportation and processing costs, different rates of entry by brokers into the market and different numbers of information users.

It was common to all simulation runs that the information digit string had 40 digits to the toroidal number base 8. There were 15 information sources each of which could hold a maximum of 15 information values. Each information user could demand up to 12 items of information each trading cycle. The maximum number m of new jobbers that could enter the market was either 1 or 4 or 8 so that at each trading cycle the number of entrants was drawn at random from the U[1,m) interval. The transportation costs of one packet over the unit distance was either 1 or 1000 and the cost of processing an information item was either 100 (if the unit transportation cost was 1) or 1 (if the unit transportation cost was 1000) or, to test the effect of increasing costs 1000. The number of information users was either 100 or 25.

In order to determine the effects of intermediation in the market, no brokers were active during the initial trading cycle of each simulation and the cycle only ended when every information user had acquired the available information it demanded.

Figure 4: Information cost indices corresponding to different degrees of ease of entry by brokers – 100 users.

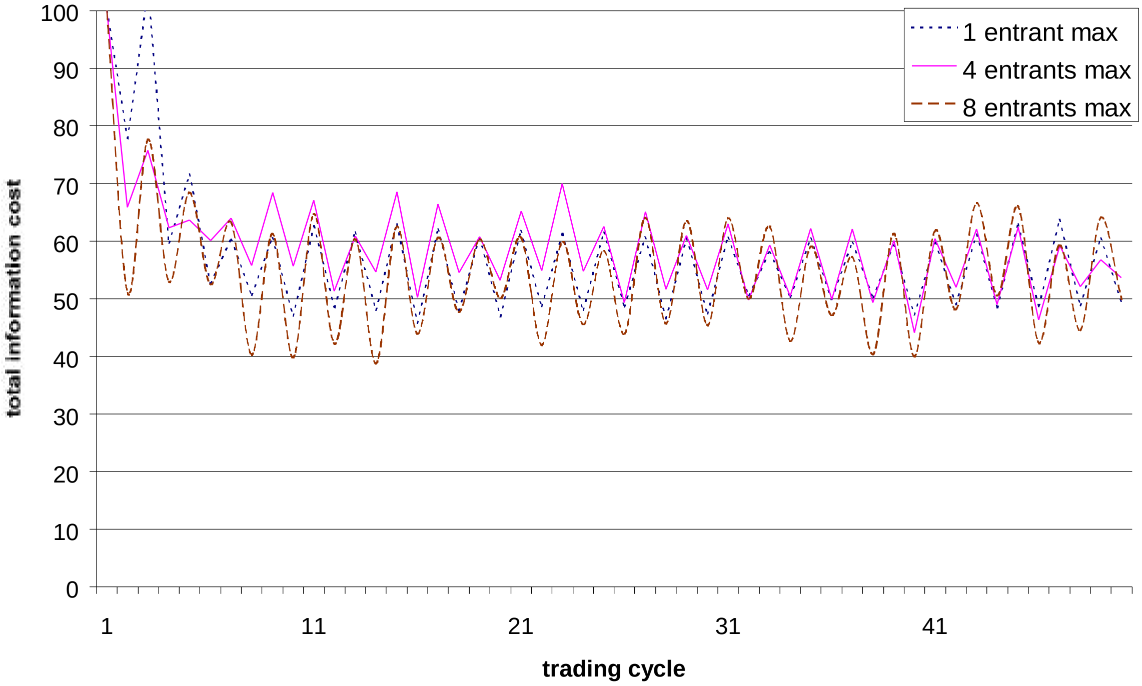

The first issue addressed was the effect of ease of entry on the efficacy and cost of information filtration. As seen from Figure 4, in the case where the unit transportation cost was 1000 and unit processing cost was 1, there were no significant differences in information costs among the three runs with maximum entry at 1, 4 or 8. The vertical axis measures the information cost index such that the index = 100 at the initial trading cycle. Evidently, the system had settled down by trading cycle 9 after which information costs oscillated mainly between 50 and 60 percent of the costs in the absence of intermediation. There is no obvious reason to conclude that increased ease of entry, as represented by higher average numbers of entrants, reduces information costs. After the first 10 cycles, the average information cost indices for maxima of 1, 4 and 8 entrants, respectively, was 54.74498, 56.93128 and 53.40589 with standard deviations of 6.537899, 6.63393 and 9.051949. The differences in the means are neither ordered according to ease of entry nor statistically significant. There is a possible relationship between variance and ease of entry since the standard deviations increase with the maximum number of entrants. There are too few observations to reach a conclusion on this conjecture.

The markets analysis suggests that a smaller number of information users will reduce the scope for economies of scale available only to intermediaries so that there will be fewer intermediaries and less cost savings. A run with 25 information users and a maximum entry of eight brokers per trading cycle generated the series for the total information cost index in Figure 5.

It is clear from Figure 5 that total information costs are systematically higher than in the initial trading cycle when there were no intermediaries. Total information costs are higher they include acquisition of information directly from sources by both brokers and users. That no intermediary survived its first trading cycle indicates that all failed to recover the costs of their own information acquisitions. The loss-making information sales to users nonetheless accounted for a substantial proportion of user purchases. Indeed, by this measure of market penetration, there were no significant differences between the run with 25 users and any of the runs with 100 users. The data series on market penetration by brokers are reported in Figure 6.

Figure 5: Total information cost index: 25 users

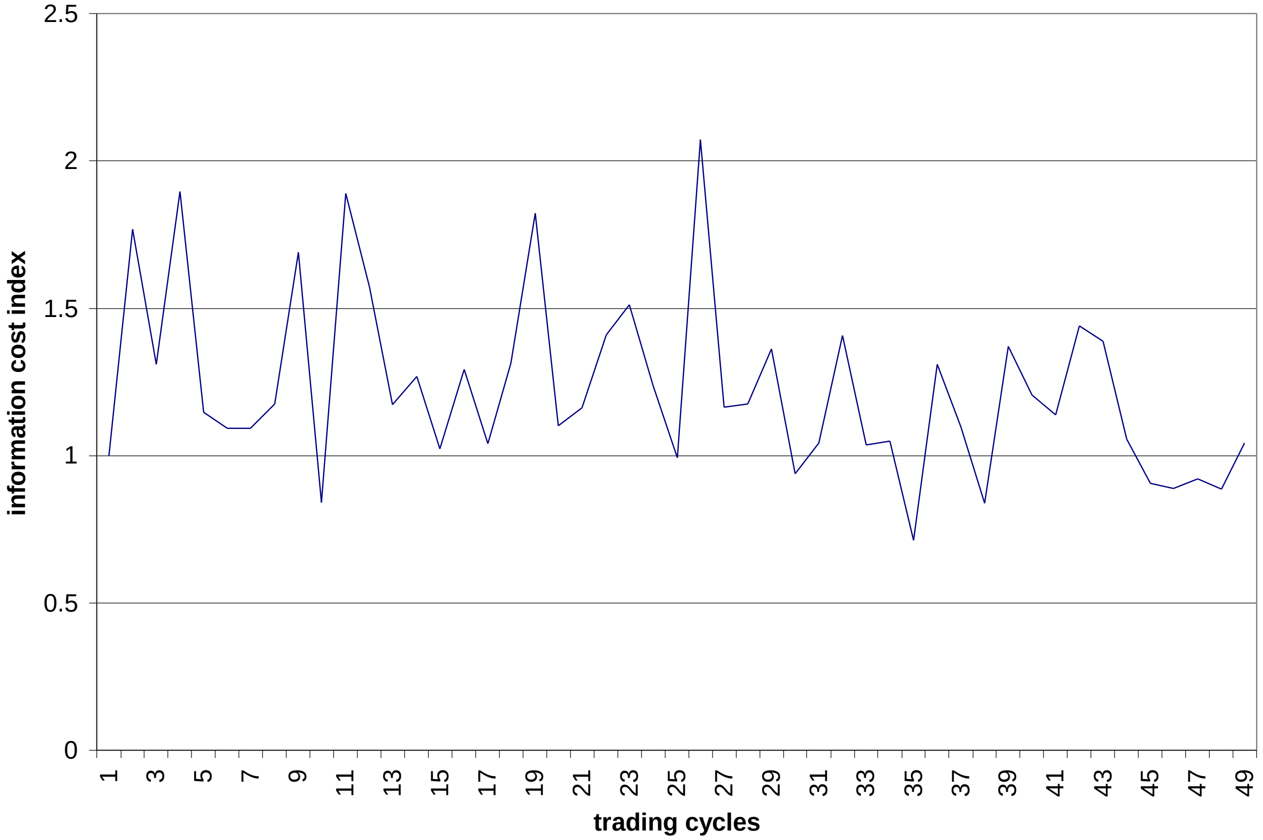

The hypotheses on intermediary survival is fully supported (though of course not proved) by the simulation results.

After trading cycle 5, no broker survived a single trading cycle. Of the six brokers that survived beyond the time of entry in the first five cycles, two survived into a second trading cycle, two into a third and one each into a fourth and fifth. None of these survived past trading cycle 8.

By contrast, intermediaries did survive in runs with larger numbers of users where the information cost index averaged little more than half the value without intermediation. These results are reported in Figure 7 and Figure 8.

With entry by one broker per trading cycle, the entrants at trading cycles 1, 3 and 4 survived profitably through the whole simulation run of 50 cycles while no other broker survived at all. The effect of increasing the maximum number of entrants per trading cycle is seen in Figure 7 and Figure 8 where the maximum entry was 8. These figures report the market shares by volume of jobbers surviving at least into a second trading cycle with at least a five percent market share in at least one trading cycle. Figure 7 is derived from a simulation run with high transportation costs and from a simulation with high processing costs. The "high" processing costs were set at 100 per item received because the average number of items in a packet was on the order of 10 (maximum 15). In consequence, the "high " processing cost per information packet was about the same magnitude as the "high" transportation cost per information packet. Evidently, there will be a small number of dominant brokers in the market specified here. Under some cost conditions there is a trend towards monopoly though the monopoly need not be permanent.

Figure 6: Market penetration by brokers

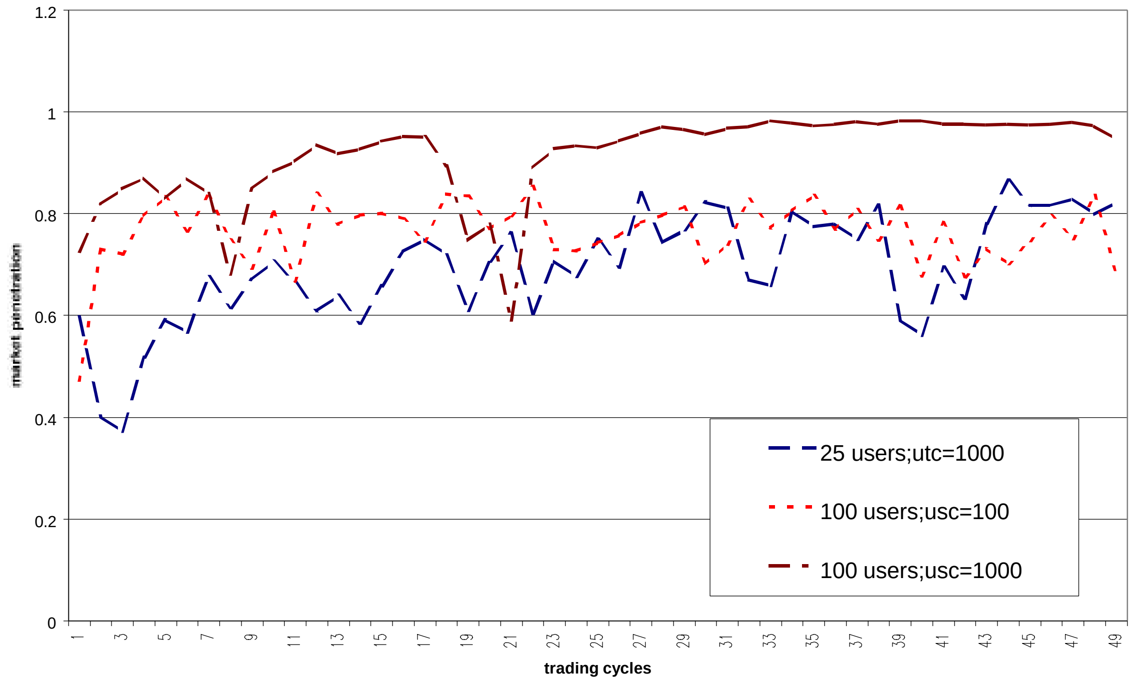

Figure 7: Market shares by volume of brokers in high transportation cost regime

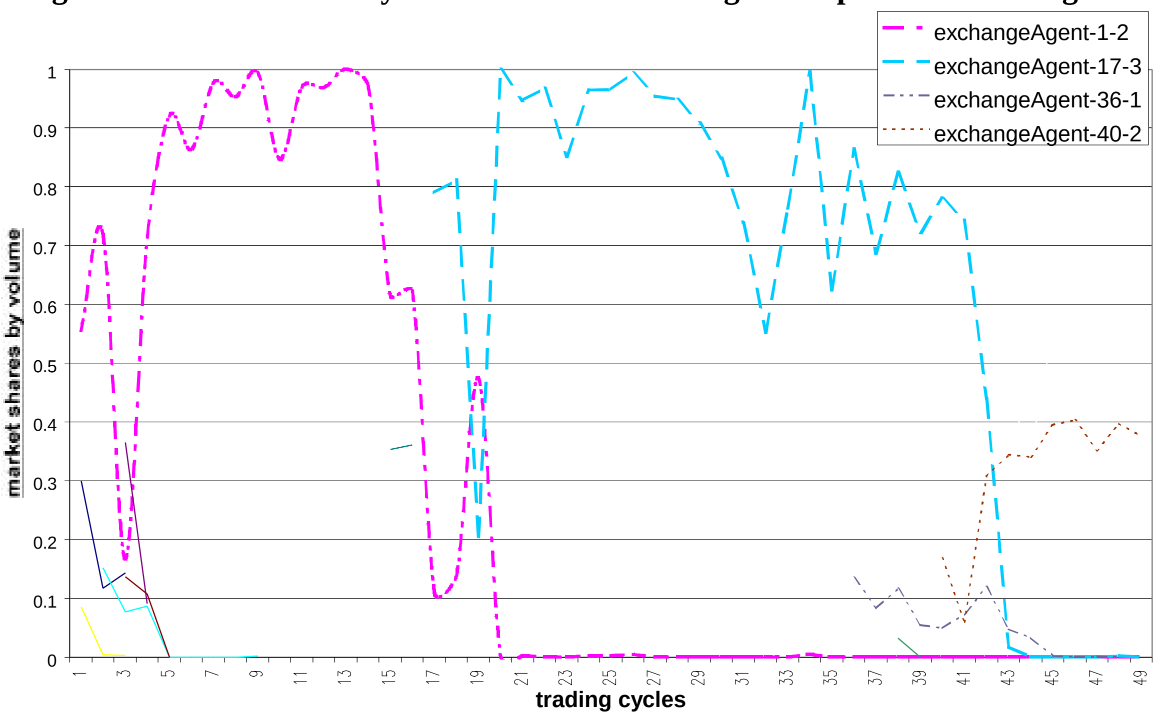

In the run generating the data for Figure 7, the first monopolist, exchangeAgent-1-2, lost its monopoly at trading cycle 17 to the new broker, exchangeAgent-17-3, and after a brief struggle was left with only sporadic sales on a small scale. Its subsequent survival was due entirely to the substantial (but declining) financial resources from its erstwhile monopoly. Towards the end of the simulation run, competition was increasing as exchangeAgent-36-1 and exchangeAgent-40-2 established footholds in the market. The randomly generated location of each of these brokers was close to the information users' edge of the market space and to each other. The mark-up by exchangeAgent-17-3 on costs to determine its minimum prices was less at entry than the mark-up of exchangeAgent-1-2. And the mark-up of exchangeAgent-40-2 at its entry was less than the prevailing mark-up of exchangeAgent-17-3. The proximity of these agents to one another meant that their costs of supplying users were much the same. The smaller mark-up of the entrant meant that its prices were lower and this was sufficient for it to establish a strong position with the customers of the incumbent.

In the run generating the data for Figure 8, with much lower transportation and higher processing costs, there were several strong brokers throughout with exchangeAgent-2-2 maintaining a strong presence and with substantial competition from exchangeAgent-1-1 early in the run and from exchangeAgent-31-4 towards the end. During the middle of the run, a number of relatively short lived brokers were successful.

Figure 8: Market shares by volume of brokers in high processing cost regime

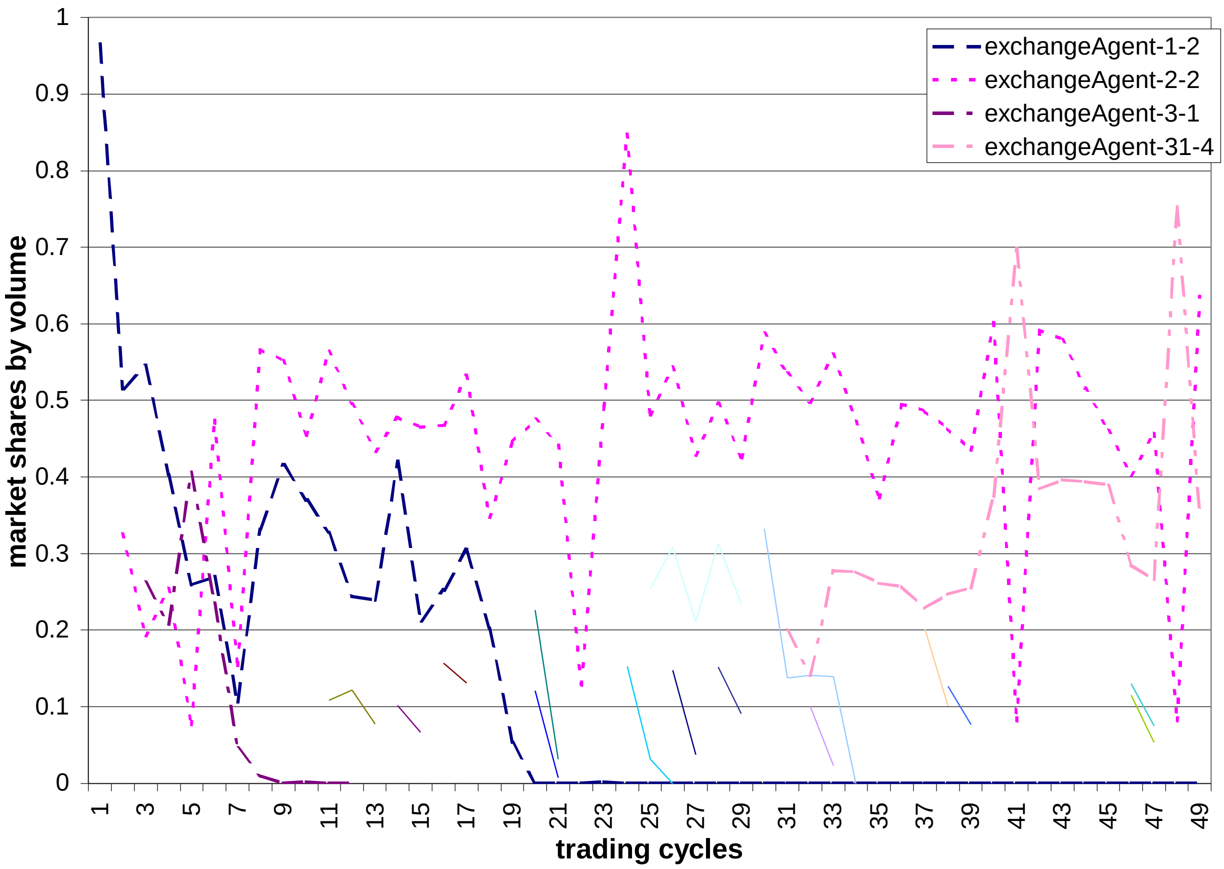

The results reported here unambiguously confirm the basic hypotheses: intermediaries can survive if and only if their activities reduce total costs in the market. In the information filtering market modelled here, all such costs are information costs. The original markets theory published some 20 years ago indicated that larger numbers of users would enhance the scope for intermediaries to reduce transactions costs and this result, too, has been confirmed. However issues concentrated on by economists, such as ease of entry into a market, appear to influence the patterns of market shares but not to influence the effects of intermediation on costs. A natural conjecture to offer in this regard is that limited cognitive capacities of information brokers will generate some relationship between numbers of surviving brokers and the scale of the information system as measured by numbers of information sources and users. This is not, however, an issue related to competitiveness. The claimed benefits of ease of entry for competitiveness and efficiency do not seem to apply to market intermediation.

7. Conclusion

The model and results reported in this paper were developed from the specification of a desired software system with a specification of information agents that was compatible with the needs of that system. Properties of the system were derived from first principles and relevant, common observation of the real-world phenomena that inspired the software system by analogy.

The endorsements mechanism of the agents supported the development of preferences by experience in a manner that was sufficiently flexible to enable them to function under a range of cost and competitive conditions. By contrast, the specification of "rational" agents in the IBM model is the likely source of the need undesirably to restrict the specification of that model.

There remains the question of how robust are the results reported here with respect to different agent specifications affect those results? While the answer is obviously beyond the purview of this paper, one approach to the investigation of that issue follows naturally from the use of the SDML modelling environment [9].

SDML is a strictly declarative language that corresponds to a fragment of Konolige's strongly grounded autoepistemic logic. This correspondence is the result of a co-evolution of the language and modelling approaches developed by the Manchester Centre for Policy Modelling. It has proved possible in simple cases to use the language to prove theorems about models' properties. More powerful theorem provers would doubtless extend the ability to prove such theorems.

Any model that runs under any programming language is sound and consistent with respect to that language. Models developed in SDML, including the model reported here, are sound and consistent with respect to SDML and therefore to strongly grounded autoepistemic logic. Consequently, the models (though not necessarily the agents) have the same formal clarity as agents represented by procedural formalisms such as expected utility maximisation or declarative formalisms such as BDI or deontic logics. Since the purpose of developing multiagent systems is to generate useful, reliable *systems* using agents as a powerful *programming* analogy, it would seem more appropriate to ensure that the properties of the systems rather than the agent representations are robust. For this reason, it would be useful to prove theorems about the effects of, for example, different degrees of ease of entry, different numbers of users *etc.* in order to determine wherever possible what properties of the system are independent of the properties of particular agent specifications.

These are, of course, issues for future research. For the present it is sufficient to have demonstrated that the cost of rational agency is a loss of scalability and applicability by virtue of the requirement to specify systems that such agents can support.

References

[1] Binmore (1987), "Modeling Rational Players – 1", *Economics and Philosophy,* Vol.3, pp.179-214

[2] Binmore, K. M. Piccione and L. Samuelson (1997), "Bargaining Between Automata" in Rosaria Conte, Rainer Hegselmann and Pietro Terna (eds.), *Simulating Social Phenomena* (Berlin: Springer-Verlag, Lecture Notes in Economics and Mathematical System), pp. 113-132

[3] Cohen, Paul R. (1985), *Heuristic Reasoning: An Artificial Intelligence Approach* (Boston: Pitman Advanced Publishing Program).

[4] Hanson, James E. and Jeffrey O. Kephart (1998), "Spontaneous Specialization in a Free-Market Economy of Agents"

[5] Kephart, Jeffrey O., James E. Hanson, David W. Levine, Benjamin N. Grosof, Jakka Sairamesh, Richard B. Segal and Steve R. White 91998), "Dynamics of an Information-Filtering Economy" *in Proc. Second International Workshop on Cooperative Information Agents (CIA-98)*.

[6] Konolige, K. (1988), "On the relation between default and autoepistemic logic", *Artificial Intelligence*, v. 35, pp. 343-382.

[7] Moss, Scott (1981a), "Technology and Market Structure in the Analysis of Exchange", *Managerial and Decision Economics*, v. 2, pp. 106-120.

[8] Moss, Scott (1981b), *An Economic Theory of Business Strategy* (Oxford: Basil Blackwell).

[9] Moss, Scott, Helen Gaylard, Steve Wallis and Bruce Edmonds (1998), "SDML: A Multi-Agent Language for Organizational Modelling", *Computational and Mathematical Organization Theory* v. 4, pp. 43-70.

[10] Sargent, Thomas (1993), *Bounded Rationality and Macroeconomics* (Oxford: Clarendon Press).