INSEAD

Faculty & Research

Development and Validation of a Multi-agent Simulation of Consumers, Retailers and Manufacturers

by

David Midgley
Robert Marks
and
Dinesh Kunchamwar

2006/03/MKT
(Revised version of 2005/78/MKT)

Working Paper Series

Development and validation of a multi-agent simulation of consumers, retailers and manufacturers

Professor David Midgley
INSEAD
Boulevard de Constance
77305 Fontainebleau
FRANCE
david.midgley@insead.edu

Professor Robert Marks
AGSM
UNSW, Sydney 2052
AUSTRALIA
bobm@agsm.edu.au

&

Dinesh Kunchamwar
School of Computing
National University of Singapore
SINGAPORE
dinesh_657@yahoo.com

Abstract

We are working on an important problem in business, namely understanding the complex interactions between consumers, retailers and manufacturers. We have developed an agent-based model based on both the academic literature and our understanding of the marketing of frequently purchased products. This model has been embedded in a genetic algorithm optimiser for purposes of verification and validation. However, the results we will present from the verification stage, together with considerations of the likely data available for validation, suggest a different approach may be needed for building, verifying and validating such models. This approach will emphasis minimalism and a clear focus on validation from the start of model development.

1 Introduction

1.1 The research problem

We are working on an important problem in business, namely understanding the complex interaction between consumers, retailers and manufacturers that leads to market and economic outcomes such as consumer satisfaction, retailer and manufacturer profits. Our particular focus is on these interactions for frequently purchased items such as those typically found in supermarkets—where marketing is often achieved through vehicles such as local advertising and in-store promotions.

Individual aspects of this problem have been discussed in many literatures and from theoretical, empirical and modelling perspectives. For example, the field of marketing has a long tradition of modelling the impact of marketing actions on the sales and market share of products (e.g. Cooper and Nakanishi, 1993; Hanssens, Parsons and Schultz, 2003). Similarly, game theorists have addressed the interaction between consumers and product manufacturers (e.g. Carpenter, Cooper, Hanssens and Midgley, 1988) and between manufacturers and distributors (e.g. Iyer, 1998).

However, we do not believe that the complete system has been adequately modelled to date, especially if we require reasonable realism in the specification of each agent within the system and if our objective is to model the multiple period interactions of primary interest to managers as well as scholars. Here we define "complete system" to mean a set of consumers purchasing a category of products, the competing retailers that make these products (amongst others) available to consumers and the competing manufacturers that both supply these products to the retailers and promote them through advertising and store displays. Much of the existing literature has modelled this system through aggregate equations representing one or more of the agents and many game theoretic approaches have studied "one-period" or "two-period" games rather than interactions over multiple periods. This literature has undoubtedly added greatly to our understanding but arguably does not capture the richness and diversity of agent decision-making nor the complete dynamics of the market-place.

As a consequence, our knowledge of the interactions between consumers, retailers and manufacturers may remain incomplete and the normative prescriptions we make from this partial view of the system may be incorrect.

1.2 Our prior work

Our belief that the system has yet to be adequately modelled to date and the genesis for the approach taken in this paper stems from some of the limitations of our prior work. In Midgley, Marks and Cooper (1997) and Marks, Midgley and Cooper (forthcoming) we modelled consumers through a sophisticated empirically-based econometric model of product sales and the decision-making routines of the marketing managers of four manufacturers as co-evolving bit strings optimized by the genetic algorithm. Only one retailer was modelled—through a simple institutional policy on the number of store promotions allowed per week and a random draw if more than one manufacturer offered a promotion in any week. In subsequent work we have sophisticated the single retailer by again developing an empirically-based model, this time of the retailer's choice of promotions offered by the manufacturers. These models capture aggregate data extremely well and easily model a multiple period game—which we see as the main contribution of our bit string approach as compared to standard game theory models. However, we have a number of issues with these models.

First, econometric equations of consumer and retailer choice are only convenient representations; they do not reflect what is known about human decision-making processes. If, for example, decision-makers use lexicographic ordering or elimination by aspects in their choices, or pay selective attention to store displays and advertising, this is likely poorly represented by aggregate algebraic expressions. Similarly, our use of bit strings to represent marketing managers is also a convenience—we do not believe real managers map input information to output decisions via binary representations. These bit strings also suffer from their "black box" nature. Even though they fit the data well it is hard to understand what the optimized mapping represents.

Secondly, we believe that individual consumers, retailers and manufacturers have differing decision-making processes and behaviours. For consumers we often think of different market segments but we would also not expect Carrefour to make decisions in the same way as Tesco for reasons of history, organizational and cost structures and so forth. Incorporating heterogeneity in existing econometric approaches, while not impossible, usually results either in clumsy simplifications or equations whose solutions are intractable.

Finally, current approaches are "top down" imposing structures on markets that are useful to the researcher or analysts. In contrast, real markets are built "bottom-up" from the actions of independent agents of differing types. By imposing structures, rather than allowing interactions, we may be artificially constraining the system in a way which we may not understand and which may not reflect the real dynamics or behaviour of the system.

1.3 Objective of the current work

The objectives of our current work are thus to use a bottom-up or disaggregate approach to modelling this type of market system, in particular the ideas and techniques of agent-based modelling. In taking this approach we will build on the existing literature to specify the decision-making and interactions of the three types of agents (consumers, retailers and manufacturers) and following from our prior work we will seek to validate the model empirically.

1.4 Structure of the paper

The second section describes our first prototype model—which we call our Mark 1 Market—and its implementation in RePast. We present the philosophy behind the design of this Mark 1 Market and the detailed specifications of the three types of agent. The third section discusses issues in the verification and validation of such simulation models and the particular approach we took to these critical issues. This approach is based on the ideas of Miller (1998) and the genetic algorithm—together with the concern for empirical validation demonstrated in our prior work. We should note, however, that because of lack of adequate data we have not yet done any validation work. Instead, the fourth section discusses some results from the verification phase, results that led us to change our perspective on modelling such systems, whereas the fifth section simply discusses some of the issues we may face in validation. As a consequence, the sixth and final section of the paper sets out our reasons for considering a different and simpler approach to specifying the agents and outlines our ideas on the next step—the Mark 2 Market.

2 The Mark 1 Market

2.1 RePast and the basic modelling philosophy

We choose RePast as the development platform for our model primarily because of its basis in Java, the availability of ready-made objects for agent-based modelling and the more extensive community using it than some of the other platforms. We mention this here simply because choice of platform also implies some choice of methods and assumptions.

The basic philosophy of our RePast model is one of memory and decision rules. An agent has memory of what worked for it in the past and rules for deciding on which new opportunities to consider and how to evaluate them against already known alternatives. Essentially all agents update their beliefs as new information becomes available to them.

This basic philosophy applies to all three types of agent although the retail and manufacturer agents are concerned with profits whereas the consumer agent is concerned with consumption satisfaction. The retail and manufacturer agent are also conceptualized as having larger memory and more systematic decision-making than the consumer agent. The latter more follows the consumer behaviour literature on low involvement decision-making. The retail and manufacturer agents are fully aware of each other's proposals through their close interaction, whereas the consumer agent may only become aware of new offers through advertising or in-store promotion. Finally, following normal industry practice the retail and manufacturer agents operate around quarterly planning periods whereas the consumer agents operate in a weekly time frame. We now describe each type of agent in more detail.

2.2 The consumer agent

Each brand has three attributes, one of which is price (expressed inversely as "value-for-money"). To the uninformed consumer each attribute has a level of uncertainty. The non-price attributes may be thought of as features of the product and manufacturers are only allowed to advertise one feature during one period.

Consumer agents become aware of brands in two ways. First, they may become aware through advertising with a probability depending on the level of advertising of the brand relative to its competitors. Seeing advertising reduces the uncertainty the agent has on the advertised attribute. Second, through a simple probability of observing an in-store promotion for the brand if they visit the store on a week in which there is a promotion. Observing a promotion reduces uncertainty on the price attribute to zero.

The probability that an agent will be shopping for this product in any period is modelled as a Poisson process with an agent-specific parameter.

Agents are assumed to make screening decisions about which brand to put into their consideration set (finite memory) using a *lexicographic* rule and decisions about how to choose a brand in their consideration set using a *compensatory* rule. This follows some of the consumer behaviour literature.

To be added to the consideration set a brand the agent is "newly aware" of must be better than any brand already in the consideration set on the most important attribute to that agent (and on the second most important attribute if there is a tie between two brands on the first, etc.). "Better" implies having more of the attribute than any existing brand by an increment which simulates the cognitive cost of adding an extra brand to the agent's consideration set.

Once the agent's consideration set has been updated for any new brands to consider they apply a compensatory rule to choose which brand to buy. At the point of purchase they become certain of the real prices of the brands in their consideration set. The rule is then applied by computing an overall score for each brand (the sum of attribute levels weighted by their importance). Finally, this score is corrected for the risk of buying a brand where they have some level of uncertainty on one or more of the other attributes. The brand purchased is the one with the greatest risk-adjusted score.

Once a brand has been purchased the agent becomes aware of its true attribute levels, uncertainty drops to zero and the score is recomputed. Provided this score is above the agent's individual threshold level of satisfaction the brand is retained for consideration on the next purchase occasion, otherwise it is dropped as "unsatisfactory".

2.2 The retail agent

The focus of the retail agent is essentially on which store promotions make it the most profit. To achieve this result the retail agent retains a memory of previously successful promotions. This memory is more extensive than the consumer agent and includes two attributes of the promotion itself (discount of normal price and whether an aisle display was used or not). The agent also remembers the total category profit generated by the promotion. Note the total category focus of the retailer, which is obviously different to the focus of each manufacturer. Each quarter this memory is updated essentially to retain the M best promotions where M is the size of the retail agent's memory.

For the retail agent consideration is simple. They are aware of all the promotions being offered by all the manufacturers for the next quarter and they consider all systematically. This is more of a high involvement, systematic decision-making model.

From this set of offers the retailer chooses X weeks of promotions (where X is a state variable the agent can change from quarter to quarter) with the goal of choosing the X that maximize his profits. Implicit in this formulation is the fact that only one brand can be on promotion in any week. This is a simplification but one that is followed by some supermarket chains.

The agent's choice is operationalised in two steps. In the first step the proposed promotions are compared with those in memory brand by brand, establishing which are most similar on promotional attributes and then ascribing the category profit achieved previously to the new promotion. In the second step the agent then simply chooses the X promotions that make it the most profit without consideration of brand.

Finally, where there is no promotion for a brand in a week the normal price applies (wholesale price offered by the manufacturer plus fixed mark-up) and the retail agent charges the manufacturer a fixed fee for each week in which there is a promotion (called "slotting fees" in industry).

2.3 The manufacturer agent

The focus of the manufacturer agent is also to make the most profit but their world is more complex than the retailer.

First, the manufacturer agent can choose to change their wholesale price and their weekly advertising level from quarter to quarter. They can also choose which attribute to emphasize in their advertising each quarter and what to say about that attribute (e.g. the level they wish to communicate).

Second, they need to remember two separate classes of events corresponding to normal and promotional periods in the retail store. For normal periods, the agent's memory includes the previous MN most profitable settings of price and advertising (level and message). For promotional periods, the agent remembers the previous MP most profitable promotions (including discount, aisle display and brand profit).

Third, the manufacturer agent needs to make promotional offers to the retailer for the next quarter. They do this by first asking the retailer how many promotions will be scheduled for that quarter, X from above. They then offer their X best promotions where best is most profitable from their brand perspective—which is not necessarily the same as that of the retailer. Note also that the manufacturer agent will likely not be awarded X slots by the retailer due to competition from other manufacturers.

Finally, costs and profits are computed in a straightforward way, although each manufacturer agent can have different fixed and variable costs, with the latter also having economies of scale.

2.4 Agent heterogeneity and realism

The Mark 1 Market is currently run with one retailer, five manufacturers and 1000 consumers. These are arbitrary choices to start the project and as such were kept small for initial verification attempts. However, the model has the 37 parameters shown in Table 1 variously set to fixed values or drawn from a random distribution to generate agent heterogeneity.

Table 1 parameters of the model

	Number of parameters
Consumer agent	
Global constants	
Number of weeks for advertising awareness calculations	1
Size of consideration set	1
Mean and variance of distributions	
Attribute importance	6
Cognitive costs	2
Risk adjustment	2
Satisfaction threshold	2
Inter-purchase time Poisson lambda	2
Chance of observing a store promotion	2
Retail agent	
Global constants	
Number of best promotions remembered	1
Slotting fee	1
Quarterly increment/decrement to mark-up	1
Range of mark-ups allowed	2
Manufacturer agent	
Global constants	
Quarterly increment/decrement to wholesale price	1
Range of advertising levels allowed	2
Quarterly increment/decrement to advertising level	1
Size of normal memory	1
Size of promotional memory	1
Fixed and variable costs	2
Mean and variance of distributions	
Product attributes	6

This Mark 1 Market is built from two sources, the literature, especially that on consumer behaviour, and industry knowledge (one of the co-authors having worked on the analysis of supermarket data for many years). The resulting model is realistic at least in terms of face validity but is evidently quite complex in overall structure. Yet the model has limitations. For just a few examples of this we would note that in reality (1) manufacturers take explicit account of the actions of their competitors (here they do not, rather competition is indirectly inferred from results), (2) retailers and manufacturers negotiate over prices and promotions (here they simply accept/reject offers) and (3) consumers forget advertising (here they do not). So while the model is complex it is not fully based on either the literature or industry knowledge—it remains a considerable simplification. This point is important because as we will argue later the trade-off between realism and simplicity is a difficult one to judge correctly. Furthermore, in following normal scientific traditions and building the Mark 1 Market on the basis of existing knowledge we produced a model that is probably too complex. To see this point more clearly we need to turn to the critical topics of verification and validation.

3 Verification and validation

It is clearly important that we validate agent-based models just as we would validate other scientific models. However, with complex simulation models like the Mark 1 Market there is a step prior to validation, namely verification that the software correctly implements the model the researcher intended. To quote Boehm (1981) one has to first demonstrate that one is "solving the equations right" before moving on to demonstrate that one has "solved the right equations."

3.1 Verification

Unfortunately this is not an easy task of complex software. The researcher writes a detailed specification, the programmer turns this into code and along the way many decisions have to be made as to exactly how to implement the desired rules and interactions. For analytical models with a few equations it is usually possible to verify that the equations have been correctly solved. With hundreds of lines of Java this is not so easy. Figure 1 illustrates a part of our code that implements one important step—the addition of a newly observed brand to the consumer agent's consideration set. This is just on step in many, each of which interacts with the others.

Figure 1 an extract of the Mark 1 Market code

```
// add exactly one product to consideration set using lexico compare
model.write("scanning\n", 'c');
if(productInfoList.size() != 0)
{
    LexicoOrderedSet los = new LexicoOrderedSet(productInfoList.size(), model);
        for (int i = 0; i<productInfoList.size() ; i++)
        {
          if(!((ProductInfo)productInfoList.get(i)).awareness)
              continue;
          pi = (ProductInfo)productInfoList.get(i);

          if(considerationSet.isInsertable(pi, C, Wx, Wy, Wz))
          {

                  los.insert(pi, Wx, Wy, Wz);
```

Indeed some writers think it "almost impossible to verify totally a model for a complex system" (Kelton et al 2001) and "It is simply not possible to test the full parameter space of large, complex software process simulation models" (Shervais et al 2003).

Miller (1998) provides one solution to this problem with his ideas of automated non-linear testing systems. Miller suggests using optimizing routines to "break the model" by producing extreme or absurd outputs and demonstrates this idea by "breaking" some of the Club of Rome models. If one's model can be "broken" then the implication is that some part of the implementation was flawed.

3.2 Validation

Various approaches to validation have been suggested in the literature. Among the principle ones are "docking" (Axelrod 2003) whereby the model is implemented and compared across different software platforms and empirical validation. For the latter LeBaron (2006) suggests three steps. First, attempt to replicate difficult empirical features: does the model fit facts not otherwise explained? Second, put the parameters under evolutionary control. Here for example we might vary size of memory. Third, use the results from laboratory experiments with human subjects to validate features of the model. Moss and Edmonds (2005) suggest two levels of validation. First, micro-validate individual agent behaviour. They suggest this might be done statistically. Second, macro-validate the model's aggregate, emergent behaviour. They suggest this judgement may have qualitative aspects to it.

Figure 2 an approach to verifying & validating agent-based models

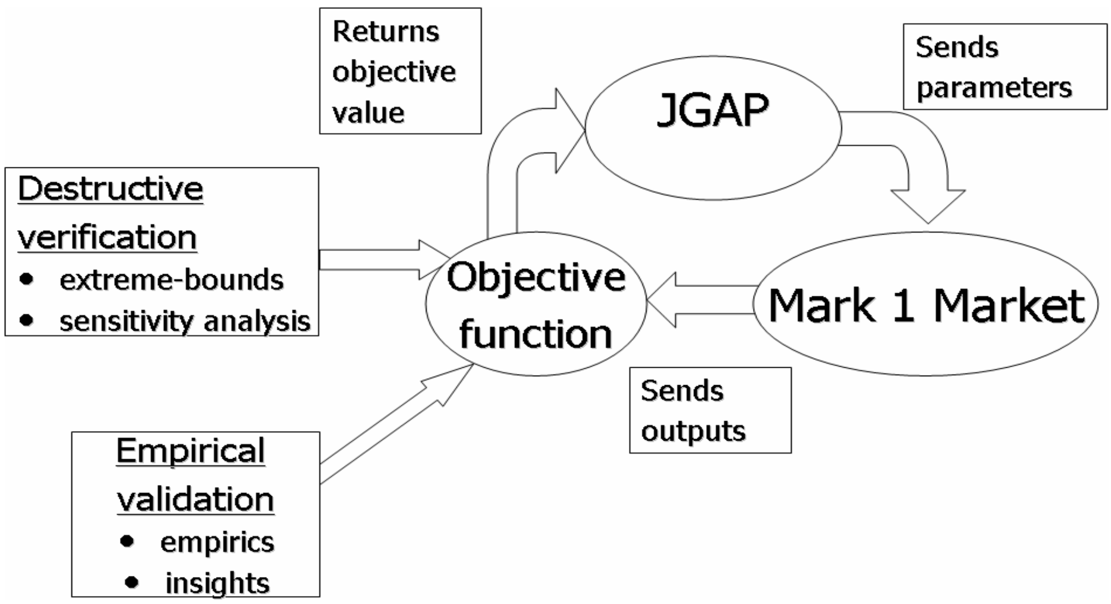

Finally paralleling Miller's ideas in verification, Schriber (2003) also suggests a role for sensitivity and extreme bounds analysis in validation.

3.3 A suggested approach to both verification and validation

Putting the ideas of Miller together with the more general discussion on validation suggests an overall approach to these problems. We suggest embedding the agent-based model in an optimisation methodology but one that allows flexibility in the specification of the objective function. For verification purposes this objective function can value extreme or absurd outputs. The goal here being to "break the model" and thus discover where implementation is flawed. For validation purposes this objective function can value the fit to real data either at the level of the individual agent or at the aggregate level. The goal here would be to examine both degree of fit and the values of the calibrated parameters for realism.

We have implemented this idea by embedding the Mark 1 Market in a genetic algorithm optimiser (JGAP). This is illustrated in Figure 2. We chose the genetic algorithm because of its robust optimization properties especially given the non-linear and complex nature of our model.

4. Verification results

We applied these ideas to verifying the Mark 1 Market. Various objective functions were tried such as:

- Maximize the market share of one manufacturer
- Equalize market shares of producers
- Maximize retailer's profit
- Maximize manufacturers' profits
- Maximize consumer satisfaction

Following Miller's advice the search space was constrained for initial analyses. First, we selected 16 parameters of key interest from the 37 available. Second, parameters were then perturbed by the optimizer within +/- 20% bounds and with discrete values. The genetic algorithm was run for 50 generations with a population of 100. Table 2 shows an extract of the output file for one such optimisation, where the bolded row indicates a set of perturbed parameters that results in a market share of 82% for one producer.

A number of issues arose during these verification exercises. First, we found code that was not invoked during the runs and needed to be examined for possible deletion. Second, we found code that was incorrect and needed to be modified. Third, we found an issue with the use of random distributions to generate the requisite degree of agent heterogeneity (particularly for consumer agents). These introduce considerable noise into the optimisation process, often such as to make it difficult to get convergence. As an expedient for these initial attempts we drastically reduced the variance of these distributions. This allows the GA to converge but obviously sacrifices some of the benefits of a disaggregate approach.

Finally, we found that the model is sensitive to some parameters but less sensitive to others. This might suggest simplifications but such a conclusion must be tempered by the relatively restricted range of verifications done to date.

5. Issues in validation

We can also use the optimiser to fit the model to real data. We have not yet done this because real data on a total system like this is not readily available. Good data on consumer purchasing patterns exists and is often integrated with data on the advertising and promotions consumers may have been exposed to. Indeed, since such data exists not only in aggregate form but also from individual household purchasing panels there is the possibility to calibrate our consumer agents at the micro-level suggested by Moss and Edmonds (2003) as well at the macro-level through aggregate sales and share data.

However, data on the retailers and manufacturers is much harder to find, particularly data on economics. This suggests that an approach here, absent access to such data, is to model observed actions and leave some factors as parameters to be estimated as a consequence of the fitting process. There is also the issue that in the consumer data we do not observe which promotional offers are made to retailers, simply the ones they accept and implement. Without additional data on promotions it may be hard to micro-validate the retail agent. Similarly we do not observe all the decision-making inputs to the manufacturer, simply the resulting actions, which again may make micro-validation difficult.

Lastly, there are interesting conceptual issues in fitting agent-based models to real data around how the differing degrees of fit to various output variables are combined and/or weighted. The existing literature in marketing has mostly either

Figure 2 extract from optimisation run

Perturbation percent																
xMean	yMean	zMean	N	WAI	WPI	SF	MUI	WxMean	WyMean	WzMean	Cmean	Rmean	Tmean	lambda	promoMean	Market Share
9	-17	15	11	-13	2	0	4	18	-18	-3	13	-20	12	6	-8	9
13	-17	15	11	-9	10	-13	4	18	-18	-17	13	-9	17	11	-8	25
-19	**6**	**0**	**11**	**-13**	**10**	**-13**	**3**	**18**	**-18**	**-16**	**13**	**-20**	**12**	**12**	**-8**	**82**
9	-17	15	13	-9	10	-13	4	18	-18	-17	13	-20	12	11	-8	2
-19	6	15	11	-10	10	-13	4	18	-11	-16	1	11	-1	0	-18	0
13	6	15	11	-3	10	-5	4	18	0	-17	13	-20	12	12	-8	53
14	20	0	1	-9	10	-13	4	18	0	-17	13	-20	12	11	-8	0
0	6	-1	11	-13	10	-13	4	18	-11	-16	7	11	-1	0	-18	27
-11	0	15	11	-13	12	-13	4	18	-18	-17	3	-20	12	12	0	1

fitted models to individual consumer data or to aggregate market data. Here we can do both simultaneously, indeed we can generate a multiplicity of outputs at different levels of analysis and observation windows. This raises question as to how we would value the degree of fit to, say, the individual behaviour of consumer agents as compared to the fit to manufacturer market share, retailer category sales, etc.

6. Conclusions and next steps

The construction and verification of the Mark 1 Market has allowed us to reach one important conclusion regarding the trade-off between realism and simplicity. Thinking about the data constraints on validation allows us to reach another shaping the formulation of these models.

6.1 Realism versus simplicity

Our experiences with verification have taught us that it is very difficult to verify even moderately complex models. And if you cannot verify your model it is not clear that you should be attempting to validate it. We came into this project with the traditional science mindset of building on the extant literature and a deep knowledge of the context. It is possible that this is the wrong approach. Any developed literature tends to emphasise nuances and sophistications leading to complexities in the model. Deep knowledge of the context tends to further add to this complexity. As a result we end up with a model with many parameters, distributional assumptions and complex interactions and housekeeping. The resulting search space for verification is large indeed and the possibility of building adequate confidence in the basic workings of the model is not that high.

In contrast, we now think the emphasis should be minimalism. For example, what are the one or two key aspects of consumer behaviour that will explain 80% of the variance in purchasing? Equally what is the simplest decision-making model for a retailer faced with competing promotional offers? And so forth. With the over-riding goal of building the simplest model that will capture a substantial part of the real phenomena.

Note this is a substantial challenge. It is easy to build simple models; it is far less easy to build simple ones that capture a substantial part of the real phenomena.

6.2 Models should be built with validation data in mind

In our particular example it has become clear that the best data will relate to individual consumer purchasing behaviour. This is simply because more commercial investment goes into collecting that data. Therefore in this area lies the best opportunity for micro-validation. In contrast, the retailer and manufacturer models will always be harder to micro-validate. This suggests to us a changed modelling approach, whereby the consumer agent is built essentially bottom-up from the data. The more assumption and parameter-based modelling might then refer to the other agents, who could be calibrated by the fitting exercise. This would also reduce the number of parameters—and thereby considerably facilitate both verification and validation.

Admittedly we have not yet fully articulated this idea. But we do think that the nature of available data should play a greater role in the formulation of agent-based models than it does in the current literature. This is not to say it should be the only determinant, theory needs also to be evident and indeed may suggest the need for new measures. However, we feel all models should be built with validation more clearly in mind.

6.3 Next steps

Our plan is to build a Mark 2 Market which is considerably simpler both in terms of the agents themselves and the interactions between them. The reduced parameter space will facilitate verification via automated testing as well as inspection. This Mark 2 Market will also be built with the available data more clearly in focus—that is, the emphasis will be on a model that can be validated more easily against real data. While we do not expect this model to be as "grounded" or "realistic" as the previous model we do hope that it can be verified and fit to data. If we know we have implemented the model we intended to, and if it exhibits a modest degree of fit, then and only then can we move on to sophisticate it.

Acknowledgements

The support of the INSEAD R&D fund is gratefully acknowledged as are comments from participants at the Marketing Science conference in Rotterdam, June 2004, the EUMAS Agent-Based Modelling workshop in Barcelona, December 2004, and the Lake Arrowhead Human Systems workshop, May 2005.

References

Axelrod, Robert. Advancing the art of simulation in the social sciences. *Japanese Journal for Management Information Systems*, 12(3), 2003.

Boehm, Barry. *Software Engineering Economics*. New York: Prentice-Hall 1981.

Carpenter, Greg, Lee Cooper, Dominique Hanssens and David Midgley. Modelling asymmetric competition. *Marketing Science*, 7(4):393-412, 1988.

Cooper, Lee and Masao Nakanishi. *Market share analysis*. Boston: Kluwer, 1993.

Hanssens, Dominique, Leonard Parsons and Randall Schultz. *Market response models: econometric and time series analysis*. Boston: Kluwer, 2003. Second Edition.

Iyer, Ganesh. Coordinating channels under price and nonprice competition. *Marketing Science*, 17(4):338-355, 1998.

JGAP. http://jgap.sourceforge.net

Kelton, David, Randall Sadowski and Deborah Sadowski, *Simulation with Arena*, New York: McGraw-Hill, 2001. Second edition.

LeBaron, Blake. Agent-based computational finance in *Handbook of Computational Economics*, Volume 2, edited by Leigh Tesfatsion and Kenneth L. Judd, Amsterdam: Elsevier Science, forthcoming 2006.

Marks, Robert, David Midgley and Lee Cooper. Co-evolving better strategies for oligopolistic price wars in *Handbook of Nature Inspired Computing for Economy and Management*, Jean-Phillippe Rennard (editor), forthcoming.

Midgley, David, Robert Marks and Lee Cooper. Breeding competitive strategies. *Management Science*, 43(3):257-275, 1997.

Miller, John. Active nonlinear tests (ANTs) of complex simulations models. *Management Science*, 44(6):820-830, 1998.

Moss, Scott and Bruce Edmonds, Sociology and simulation: statistical and qualitative cross-validation. *American Journal of Sociology*, 110(4): 1095-1131, 2005.

Schriber, Thomas and Daniel Brunner. How discrete event simulation works. *Proceedings of the simulation solutions 2003 conference*. Institute of Industrial Engineers 2003.

Shervais, Stephen, Wayne Wakeland and David Raffo. Evolutionary verification and validation of software process simulation models. Mimeo. <http://prosim.pdx.edu/prosim2004/abstract/wakeland_ext_abs.pdf>