

THE RANGE OF PREDICTIONS FOR CALIBRATED AGENT-BASED SIMULATION MODELS

DongFang Shi
Roger J. Brooks

Department of Management Science
Lancaster University Management School
Lancaster. LA1 4YX. U.K.

ABSTRACT

Agent-based simulation is increasingly used to study systems in many areas of business and science. Using agent-based simulation for prediction could be very valuable. However, these models usually have a lot of parameters which are difficult to measure directly leading to uncertainty as to the best values to use. Obtaining the values for the parameters may require calibration of the model against observed historical output data. This type of problem is an inverse problem and there may be many sets of feasible parameter values giving a wide range of predictions. The work described here investigated the extent of this problem for a word of mouth consumer model.

1 INTRODUCTION

In recent years agent-based (or individual-based) simulation has received a lot of attention. Agent-based simulation refers to simulations of systems that contain agent entities whose behaviour depends dynamically on the state of the system. This enables the agents to adapt their behaviour to changing conditions. In modeling such adaptive behaviour, agent-based simulation is a tool commonly used in complexity science (Waldrop 1992).

There is no standard definition of an agent. Some definitions list a set of properties but a better approach is perhaps simply to say that an agent is an entity for which some cognitive process is modeled (Edmonds and Mohring 2005). Usually, agents receive information from the environment (including other agents) and have internal rules that represent the cognitive decision process and determine how they respond. The rules can be a simple function of the inputs received or can be very complex incorporating various internal state parameters, and can also include a model representing the agent's worldview of some part of the environment (such as predictions of other agents' behavior). An example of a framework for complex cognitive processes is the PECS model, which has a hierarchical structure with states for physis (physical body), emotion,

cognition and social status as well as sub-components for each of these (Schmidt and Schneider 2004).

In some cases the rules governing the agents' behavior are fixed throughout the simulation and in other cases the rules can change to represent learning. The number of agents modeled can also vary from an individual agent through to a large population. Populations are usually heterogeneous with individual agents having different parameters or even quite different rules (e.g. different trading strategies in a stock market simulation). Interactions between the agents are often a key part of the behavior of the system. A very wide variety of applications have been studied using agent-based simulation including stock markets, auctions, the spread of disease, ecosystems, military battles, crowd dynamics, sports games, transport, social behaviour, social networks, the development of technology, and consumer market behaviour (such as fads).

1.1 Nature of Agent-Based Simulation

Although agent-based simulation is sometimes presented as if it were a new type of modeling, many "traditional" simulations feature some adaptive agent behavior. For example, a simple queueing simulation may include a rule that customers (the agents) will not, with some probability, join the queue if its length exceeds a certain value, or that customers leave the queue if they have to wait too long. The customers therefore have rules of behavior that depend on the prevailing conditions in the environment and so can be considered to be simple agents. Genuine novelty in agent-based simulations can lie in the complexity of the rules of behavior, in the size of the population, or in the application of simulation to new situations. One of the causes of the greater use of agent-based simulation is that increasing computing power now makes such simulations feasible. There is also an appreciation that for some systems an agent-based approach may be necessary in order to capture the dynamics of the system.

Much of the agent-based simulation work has had the aim of increasing the understanding of the type of system

rather than trying to reproduce a specific situation. Such an approach can be very valuable in producing important new insights and improving understanding. Simulation models in general are constructed by modeling local behaviour and then connecting the different parts together and allowing them to interact. Therefore any simulation model can provide useful information about the relationship between local structure and global behavior, which can increase understanding. However, relating this to a particular real system implies that the model structure is a good representation of the important parts of the real system. If this is not the case then the implications drawn may be incorrect. It is therefore important to assess the validity of the model, although in the absence of a specific real system, validation can only consist of a subjective assessment of the plausibility of the model structure and of the responses (white box validation (Pidd 1996)). For example, one of the early pieces of work was the boids simulation (Reynolds 1987), which tried to find rules for general "boid" agents to produce flocking behaviour that appeared realistic compared to the flocks, herds and schools of different animals in the real world. Some of the social science simulations are highly simplified models of virtual societies, such as the Sugarscape model of Epstein and Axtell (1996). Criticisms of these sorts of models in some quarters have been that they are too divorced from reality to provide useful information about the real world and may reflect the prejudices of the model builder (see, for example, Lansing (2002) for a discussion of this debate).

1.2 Prediction, Model Calibration and the Inverse Problem

For some applications, using agent-based simulation for prediction (rather than just better understanding) could be very powerful. For example, a company might wish to use a model of the population of their customers with word of mouth interactions to predict the sales of the product or the effect of an advertising campaign. However, the problem is that agent-based models typically have a very large number of parameters and many of these cannot be measured directly or estimated with sufficient precision. The only other information available may be historical output data from the real system. Such data can be used to calibrate the model by finding parameter values that produce a good fit with the data. This is known as an inverse problem since it consists of using the outputs to determine the inputs. The problem is that there will usually be many solutions. There are two main reasons for this. The first is that there are often many parameters and few historical data values. The second is that any model that produces a good fit should be considered acceptable. A perfect fit is not expected because any simulation is a simplification of the real system and also there may be measurement errors in the historical data.

The result is that a wide range of sets of parameter values may give an acceptable fit and are therefore feasible values. However, they may give quite different predictions. The inverse problem has been studied in other areas of science including groundwater modeling (Yeh 1986). Groundwater models are deterministic models of underground aquifers (porous and permeable rock that holds a useable amount of water, acting like an underground reservoir). These models require geological parameter values over the area modeled and yet only a few direct measurements may be available. The direct measurements are point values and only give limited information for the model since the basic model elements are cells representing quite a large area. Output values for the groundwater head levels may be available at certain points (usually wells).

The complex non-linear nature of most simulation models means that there is no simple equation for the feasible values of the parameters. Therefore methods used for tackling the inverse problem have involved running the simulation and deriving alternative predictions in some way from those runs. Brooks et al. (1994) describe an approach to find the range of predictions of a groundwater model from alternative feasible calibrations that was applied to an existing model of the Birmingham aquifer in the U.K. (Greswell et al. 1994). The aquifer was represented in the model by dividing it into rectangular cells each of which required geological parameters. The model also included input of water from rainfall and other sources (such as water mains leakage), and abstraction of water from wells. Historical data of groundwater measurements were only available at 12 sites for certain years and there were few readings before 1970. The objective was to predict groundwater levels for 2020 particularly focusing on areas where the water will be close to the surface. A prediction measure based on this objective was defined and a criterion was also established for an acceptable model based on the model output data being sufficiently close to the historical values. A search method was then used to look in the parameter space for the best and worst case predictions amongst acceptable models. The result was quite a wide range of predictions indicating that a single point prediction could be very misleading. Instead, the appropriate approach is to take account of the alternative feasible calibrations and to evaluate, in some way, the different predictions they produce. An alternative way of doing this, developed by Beven and Binley (1992), is to build up a probability distribution of predictions using the likelihood of each parameter set being the best.

2 AIM AND OVERVIEW OF APPROACH

The aim of this study was to investigate the calibration problem for an agent-based simulation, which should give an indication of the limitations of using such models for prediction. An alternative viewpoint is that the study may

indicate the amount of data required to produce a narrow range of predictions.

The approach used was to develop an agent-based model and to treat this model as the real system. Output data from this model can then be taken as measured values from the real world and, in a pseudo-modeling exercise, used to calibrate an agent-based model of the system. A method similar to that of Brooks et al. (1994) can then be used to investigate the variations in predictions. The advantage of such a pseudo-modeling exercise is that the “real system” is completely known. Consequently, the models’ predictions can be compared with the “true” future values, and the precise differences between the models and the real system are also known. In the experiments carried out here the model has exactly the same structure as the “real system” and so the uncertainty in the predictions is entirely due to the nature of the calibration process rather than any differences in the way the model and the real system work. This allows the study to isolate this effect which would not be possible if an actual real world system was used.

In most systems to which agent-based modeling is applied, the population is heterogeneous and there is a lack of detailed information about each individual in the population. There is also uncertainty as to exactly how the individuals will behave and interact in the future. A typical model therefore uses random sampling for both the characteristics of individuals and the interactions between them during the model run. Since the model is stochastic, multiple replications and statistical analysis are required to take account of this uncertainty. This creates an additional complication in the calibration process compared to the deterministic groundwater models. As a consequence, a good fit with the historical data requires comparing the measured values against values from the multiple replications, and predictions also need to be produced using multiple replications.

3 SIMULATION DESIGN

3.1 Model application

The application chosen for the research was a consumer word of mouth model. The type of product we had in mind when constructing the model was one with a short life cycle, with a high likelihood of the passing of information and opinions between consumers by word of mouth, and that is purchased as a one-off item (rather than a repeated purchase). Examples would include a computer game, a music album or a cinema ticket for a particular film. The population represented might be school or university students.

Other studies that have investigated word of mouth consumer behaviour include Baxter et al. (2003) and Ki-

jima and Hirata (2004). Baxter et al. (2003) built a generic model to allow companies to investigate the impacts of customer relationship management (CRM) strategies, with an aim of better understanding through comparison of scenarios rather than making specific predictions. Their model has 500 agents, connected in a way to mimic a social network. The agents have perception values for the price and quality of the product that change based on interactions with other agents, their experience of the product and external factors (marketing, competition, CRM). The product is a repeat purchase product (such as a subscription service) and each agent has a threshold value for the total of their price and quality perceptions above which they purchase the product. Word of mouth interactions about the product between the agents become less frequent the longer they use the product, and there is also a loss of perception at each time step to represent the effect of the competition. Customers may therefore be gained and lost as their perception values change.

The aim of Kijima and Hirata (2004) was to look at the effect of different network structures, although the precise size and structure of the networks used is unclear. Kijima and Hirata used an SIR (susceptible / infected / removed) approach, based on disease transmission, for passing information between the agents. The purchasing decision depended on the agent’s enthusiasm for the product, which is a function of the utility of the product for the agent, the reliability of the information and the agent’s attitude to risk.

3.2 Agents’ Attributes and States

Our model contains a heterogeneous population of consumers (the agents). There is no consensus in the literature as to the best way to model a consumer market with word of mouth interactions. Instead, the attributes of the agents and the interactions between agents in our model are based on our subjective views of the important factors in the real world, whilst trying to keep the model structure as simple as possible.

Each agent has two state variables that vary during the simulation:

- Knowledge (K): How much information the agent knows about the product (value between 0 and 100).
- Preference (P): How much the agent likes or dislikes the product (value between -100 and 100 with positive values indicating like).

The agents also have three fixed attributes assigned at the start of the simulation, which are selected at random for each agent from probability distributions:

- Influence (I): Represents the agent’s social standing within the population (value between 0 and 20).

- Unbiased true preference (U): Represents the preference that the agent would have about the product with complete knowledge in the absence of peer pressure (value between -100 and 100).
- Buying criteria (B): The preference value at which the agent buys the product. This represents different attitudes regarding purchasing behaviour from cautious to free spending (value between 0 and 100).

3.3 Interactions in the Model

The interactions simulated in the model, as shown in Figure 1, are conversations between agents about the product, interactions between agents and the environment regarding the product (representing seeing the product in the shops, seeing adverts or reading media articles), and agents buying the product. For simplicity, it is assumed that only pairwise interactions between two parties take place.

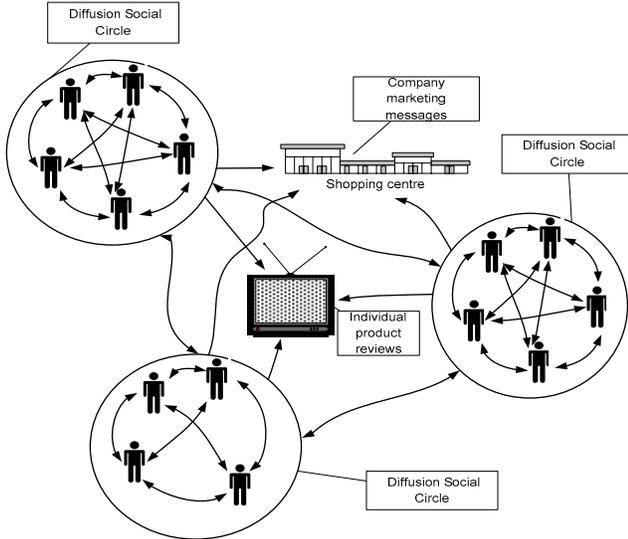


Figure 1: Interactions in the model

Fixed equations are used for each of these interactions, and the forms of the equations for agent a interacting with agent / environment b are as follows:

Change in knowledge due to the interaction:

$$K_a^{new} = K_a^{old} + \alpha(100 - K_a^{old}) \frac{K_b}{100} \quad (1)$$

Effect on preference of the change in knowledge:

$$P_a^{new} = P_a^{old} + U_a \frac{(K_a^{new} - K_a^{old})}{100} f_U \quad (2)$$

Effect on preference of peer pressure:

$$P_a^{new} = P_a^{old} + \alpha(P_b - P_a^{old}) f_K f_I \quad (3)$$

where the α values are random variables, and

$$f_U = \frac{20}{20 + |P_a^{old} - U_a K_a^{old} / 100|}$$

$$f_K = \frac{(100 + K_b - K_a)}{200}$$

$$f_I = \frac{(10 + I_b - I_a)}{20}$$

Increase in knowledge when product is bought:

$$K_a^{new} = 0.5(K_a^{old} + 100) \quad (4)$$

In each of these interactions the agents may increase their knowledge K value, representing gaining information about the product (1). The change in knowledge depends on the existing knowledge of both parties. The gain in knowledge is the proportion of knowledge not known by the agent multiplied by the knowledge of the other party multiplied by a random proportion α . For example if the agents knowledge is 70 then they will gain $0.3 \times \alpha$ of the other party's knowledge. This reflects that even if the other party knows less than the agent they will still probably have some different knowledge.

Whenever an agent's K value changes this changes its P value as a function of the U value (2). The underlying assumption is that, in the absence of peer pressure, $P = U \times K/100$ (i.e. a linear relationship between P and K). Therefore, an increase in knowledge increases the preference by a proportion of the U value (with an adjustment to take account of existing peer pressure).

The preference P value will also change due to the influence of the preference of the other agent (peer pressure) or the environment (e.g. the opinion in a magazine review) (3). The strength of both of these interactions depends on the relative knowledge and influence (the I values) of the two parties.

The agent buys the product when its preference P value reaches its buying criterion B value. This increases the knowledge of the product by half the current lack of knowledge (4). As is the case with any change in knowledge, (2) is then used to change the agent's preference.

The population is divided into groups (representing social groupings) and each agent has a much higher probability of talking about the product to other agents within the group than to other agents outside the group. Agents also have a probability of losing some knowledge each day.

Interactions with the outside environment regarding the product are split into two types (Figure 1), information from the company and information from independent sources. These interactions use Equations (1), (2) and (3) in exactly the same way as interactions with other agents.

At the start of the simulation all the agents have no knowledge and no preference about the product since it is a new product. However, the company conducts an initial marketing campaign and the agents may also see the prod-

uct in the shops or read about it in the media. These interactions enable the agents to gain knowledge and change their preference in the initial stages of the simulation. The limited time of the campaign was modeled by the probability of receiving outside information being reduced linearly down to 25% of the initial value over a period of 75 days.

3.4 Model Parameters and Modeling Scenario

Having constructed the model structure as described in Section 3.3, the next stage in the modeling was to choose the values for the parameters. The model with the default parameters is the “real system”, and so the parameters were adjusted until the model gave plausible behavior. The resulting parameters are shown in Table 1. The + and – in the table refer to the proportion of the population having negative values (i.e. dislike). The probability of an agent having a negative U value is fixed at 10%. Company and independent refer to the two types of outside information explained in the previous section. The probability of the independent source having a negative preference is fixed at 20%. The number of random conversations each day are the total number of such conversations in the population (i.e. 5 pairs of agents are picked at random from the population). The population is divided into groups of size between 2 and 8. A binomial distribution is used for the probabilities of different group sizes as it is considered to provide suitable values (probability of size x = probability of $x-1$ successes from 7 trials with probability of success 0.5, with a slight adjustment because the probability value for 0 successes is ignored).

The scenario investigated in the research was that the current time is several weeks after the product was launched (10 weeks for the first experiment but this time was varied in subsequent experiments). Data for the total sales to date have been collected and provide the best available data to use for calibration. The company wishes to predict the total sales of the product and so the model should match the sales to date and needs to forecast the final sales. The key output variable is therefore sales and all sales values are simply the number of items sold.

3.5 Behavior of “real system”

3.5.1 Sales Distribution

The model was run with 500 agents, which represents a small community such as a school. The run length was 730 days (2 years). The model with the default parameters was run 1000 times and this was assumed to represent the total population (i.e. 1000 schools). The average of these 1000 replications therefore represents the true behavior of the real system. Figure 2 shows the product life cycle of sales

per day for these replications and Figure 3 shows the same data as cumulative sales.

Table 1: Default parameters (the normal parameters respectively are the mean and standard deviation, and the uniform parameters are the minimum and maximum values)

Parameter	Value
U(+)	Normal(75,15)
U(-)	Normal(-75,15)
B	Normal(65,10)
I(company)	Uniform(0,5)
I(agent)	Normal(10,3)
I(independent)	Normal(10,3)
K(company)	Normal(60,15)
K(independent)	Normal(40,15)
P(company)	Normal(60,15)
P(+)(independent)	Normal(65,10)
P(-)(independent)	Normal(-65,10)
P(talk to same group)	10%
Random coefficient	0.01~0.15
P(receiving outside information)	Uniform (0%, 20%)
P(outside information coming from company)	80%
Proportion of K lost each day	Uniform (0%, 1%)
Random conversation per day	5
Population	500
Simulation length(days)	730

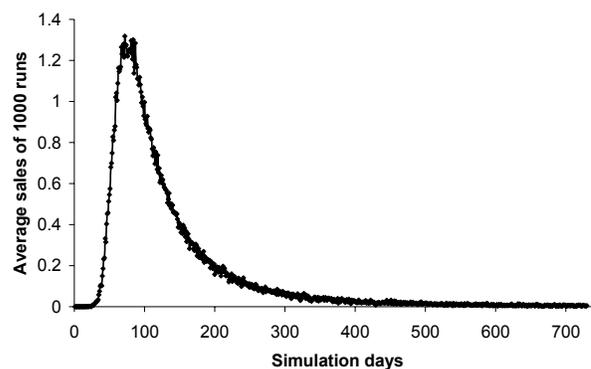


Figure 2: Product life cycle for “real system”

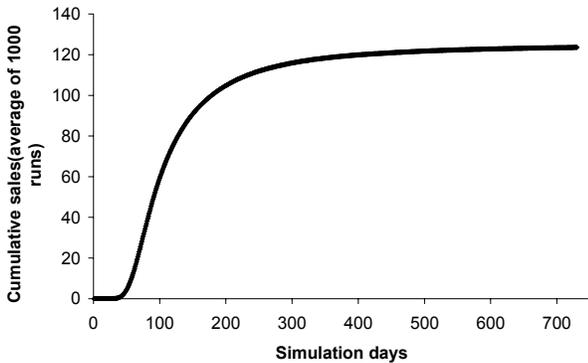


Figure 3: Cumulative sales for “real system”

The average total number of sales per replication is 124 (25% of the population). There is considerable variability across the replications and Figure 4 shows a histogram of the sales for the 1000 runs.

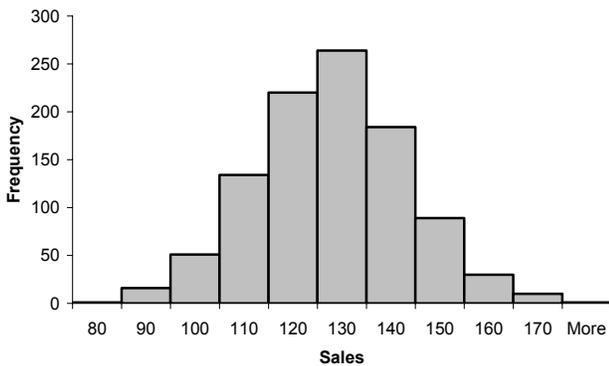


Figure 4: Histogram of sales for “real system” (the x-axis values are the upper values in each interval)

3.5.2 Further Analysis

Various additional analyses were carried out to get a better understanding of the behavior of the model. The effect of peer pressure was examined by plotting the proportion of the group that purchased the product for all the groups in one run of the model. As mentioned in Section 3.4, the group size varies between 2 and 8. The results are shown in Figure 5. The bars exclude the lower value in the range, except for the first bar which does include 0% (and, in fact all these values are 0%). All the values in the upper bar are 100%. If there was no peer pressure (i.e., no group effect) then the agents that make a purchase would be equally likely to be in any group and Figure 5 would be similar to a binomial shape with a probability of success being the proportion of sales in the population (25%). Instead there is a much higher frequency of extreme values (particularly 0% and 100%) indicating strong peer pressure resulting in similarity in the behavior within the group (e.g. there are

many groups where none purchase and many where all purchase).

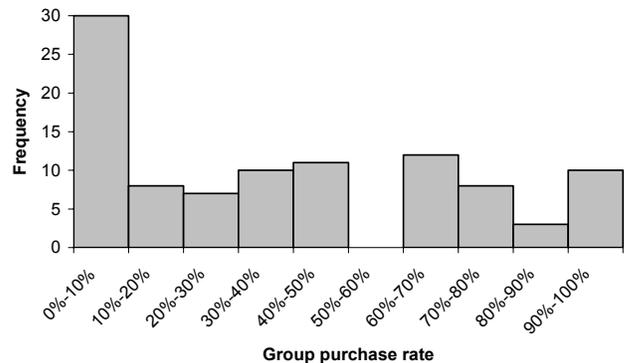


Figure 5: Purchase rate for the groups on one run

Sensitivity analysis was carried out for all the main parameters. As an example, Figure 6 shows the results for the probability of receiving outside information. This was done with a fixed probability rather than the uniform distribution of the default parameters. With a value of 0% there is no outside information and so the agents are unable to gain knowledge or preference and so there are no sales. A probability of 5% is enough to generate about 100 sales although it takes a longer time for the agents to gain knowledge and so there are very few sales in the first 10 weeks. As the probability increases above 5%, the total sales increase approximately linearly. However, even with a probability of 95% the total sales are only 200. This is because the probability reduces during the first 70 days to represent the initial advertising campaign and therefore there is a limit to the effect that increasing this parameter can have. At high values most of the sales are within the first 70 days.

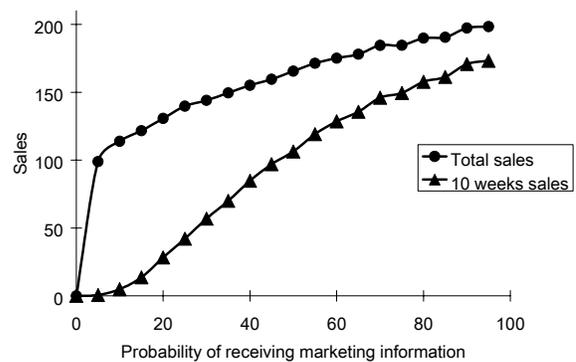


Figure 6: Sensitivity analysis on the probability of receiving outside information.

4 CALIBRATION AND PREDICTION

4.1 Process for Finding a Prediction Range

The modeling process consisted of searching for parameter values that give a good fit with the observed data (the total sales up to the assumed current time) and give the highest or lowest predictions for total sales. The same model structure as for the real system was used. In this respect the pseudo-modeling approach is removing an extra source of uncertainty compared to a real modeling situation in which the model is a simplification of the real system with many assumptions and simplifications. This has the advantage that the range of predictions must be entirely due to the calibration process.

Six parameters were chosen to be varied during the experiment as these were considered to be the most important parameters. All the other parameters were kept at the default values. The six parameters were the mean values for U for the positive and negative distributions, the mean value for the buying criterion distribution, the probability of talking to other agents in the same group, the upper value for the uniform distribution for receiving outside information and the upper value for the proportion of lost knowledge each day.

A criterion was set for the model to be give a good fit. This was that the 95% confidence interval for the estimate of the true average sales within the initial period from 100 replications must contain the true value from the real system runs (Section 3.5). A fitness function F was defined to implement this which takes the value of 0 if the true value is in the interval and $10 +$ the absolute difference between the true value and the interval if it is not in the interval.

The aim of the process is to find the highest and lowest total sales amongst the parameter values that meet the fitness criterion. There is no method that guarantees to find a global optimum for a complicated function. Instead, based on Brooks et al. (1994), the approach followed for each experiment was to run the model for a grid of three values for the six parameters (i.e., $3^6 = 729$ points in total). 100 replications were done for each point and the average sales for the initial period and the average total sales calculated. Then several local searches were carried out from different starting points on the grid and the extreme values from the local searches give the prediction range. The Nelder-Mead Downhill Simplex (Nelder and Mead 1965) was used as the local search method with an objective function to minimize $e^F + S$, or $e^F - S$ where S is the total sales. The local searches only used 10 replications so as to reduce the run time required (which was still considerable even on a high performance cluster). However, once the search appeared to have converged, the search was continued with 100 replications until a point was found that met the fitness criterion. The starting points for the local

searches were the points with the highest sales, lowest sales, three highest sales with fitness 0, three lowest sales with fitness 0 and the default parameters.

Three experiments were carried out using this method for initial periods of 70, 105 and 140 days.

4.2 Results

Table 2 shows the results for an initial period of 70 days. The total sales for the first 70 days for the real system (i.e., the average of 1000 replications) was 24.35. The results give a very wide prediction range for total number of sales (using the outer values of the 95% confidence intervals) of between 58.092 and 376.348. Examination of the parameter values may give an insight into the reason for these extreme values. The important aspect will be parameters with relatively different effects on the initial and total sales. For example, the probability of receiving outside information has more effect on sales during the initial period than on total sales (Figure 6). Part of the reason for the results could therefore be that the max S case has a high mean U value relative to B tending to produce high overall sales, whereas the low value for receiving information reduces sales in the initial period so that the model still fits. Lower U and higher probability of receiving outside information will then have the opposite effect in the min S case.

Table 2: Prediction range with 70 days as initial period

	Min S	Max S
U(+) mean of the normal distribution	55.777	91.785
U(-) mean of the normal distribution	82.407	90.898
P(talk to same group)	14.525%	5.994%
B mean of the normal distribution	60.131	50.453
P(receiving information) upper value	23.439%	15.151%
P(lose knowledge)	0.142%	0.502%
First 70 days sales (95% confidence interval)	[22.899, 24.840]	[22.727, 25.232]
Prediction (95% confidence interval)	[58.092, 60.987]	[366.091, 376.348]
Calibration benchmark	24.35	

Tables 3 and 4 show the results for initial periods of 105 and 140 days. As would be expected the prediction range for 105 days is considerably less than for 70 days. However, surprisingly, there is little difference between the ranges for 105 days and 140 days. Figure 7 shows a graph of the prediction ranges. We plan to repeat the experiments with other values for the initial period to explore this pattern further. Some of the parameter values are quite different to the 70 days results indicating a complex interaction between the parameters.

Table 3: Prediction range with 105 days as initial period

	Min <i>S</i>	Max <i>S</i>
U(+) mean of the normal distribution	88.589	71.948
U(-) mean of the normal distribution	58.455	85.100
P(talk to same group)	15.168%	6.280%
B mean of the normal distribution	77.926	48.291
P(receiving information)	24.301%	15.177%
P(lose knowledge)	0.981%	0.979%
First 105 days sales(95% confidence level)	[58.821, 63.818]	[63.451, 69.312]
Prediction(95% confidence level)	[78.764, 83.155]	[255.824, 263.835]
Calibration benchmark	63.79	

Table 4: Prediction range with 140 days as initial period

	Min <i>S</i>	Max <i>S</i>
U(+) mean of the normal distribution	56.893	87.553
U(-) mean of the normal distribution	70.038	57.662
P(talk to same group)	14.817%	15.381%
B mean of the normal distribution	60.195	77.626
P(receiving information)	33.129%	24.642%
P(lose knowledge)	0.472%	1.005%
First 140 days sales(95% confidence level)	[81.738, 86.922]	[79.364, 87.195]
Prediction(95% confidence level)	[93.357, 98.863]	[260.877, 276.682]
Calibration benchmark	86.518	

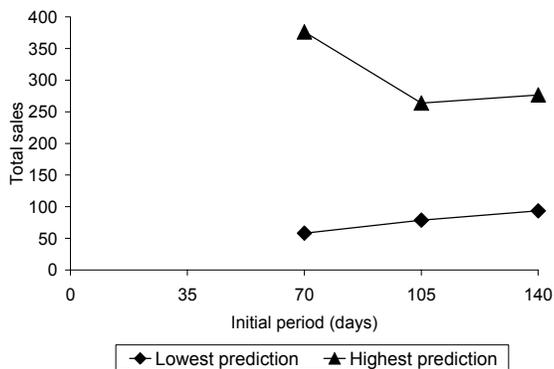


Figure 7: Prediction ranges for the experiments.

As discussed in Section 4.1, this is an idealized situation in which the model structure was identical to the sys-

tem being modeled. In a real modeling case there are additional modeling simplifications and uncertainties about the real system. Therefore a perfect fit with the observed data may not be possible, which is likely to further increase the range of predictions.

5 CONCLUSIONS

For some applications, using agent-based simulation for prediction (rather than just better understanding) could be very powerful. However, the problem is that agent-based models typically have a very large number of parameters and many of these cannot be measured directly or estimated with sufficient precision. The only other information available may be historical output data from the real system. Such data can be used to calibrate the model by finding parameter values that produce a good fit with the data. However, any model that produces a good fit should be considered to be acceptable and different acceptable models may give quite different predictions as demonstrated in this study. A method which takes account of the different feasible parameter values (such as the approach described here) needs to be used in making predictions.

The pseudo modeling approach taken here is an effective way to study and isolate specific aspects of the modeling process. In the experiments carried out here the model had exactly the same structure as the “real system” and so the prediction uncertainty is entirely due to the calibration process. Future experiments could examine many other scenarios including specific differences between the structure of the model and the “real system”. In a real life modeling project the structures will not be the same since the model will be a deliberate simplification of the real system and the complete specification of the real system cannot be known. This is likely to create additional differences when fitting the model data and the real data in the calibration process, leading to an even wider prediction range. Studies using this approach for a real system would be also be useful.

The issue examined here actually applies to any complex simulation model in which the parameter values cannot be measured directly and have to be determined by calibration. However, it is particularly relevant for agent-based simulations because in many cases they model aspects of human characteristics and behavior on which it is very difficult to obtain data. This may therefore limit the usefulness of such models for prediction.

ACKNOWLEDGMENTS

Many thanks to Richard Meek of the Marketing Department at Lancaster University for helpful discussions on consumer behavior and to the anonymous referees for their comments. This paper is partly based on Brooks, R. J. and D. Shi. 2006. The calibration of agent-based simulation

models and their use for prediction. *Proceedings of the Third Operational Research Society Simulation Workshop (SW06)*. 157-161. Birmingham, UK: The Operational Research Society.

REFERENCES

- Baxter, N., D. Collings and I. Adjali. 2003. Agent-Based Modelling – Intelligent Customer Relationship Management. *BT Technology Journal* 21(2): 126-132.
- Beven, K. J. and A. Binley A. 1992. The future of distributed models: Model calibration and uncertainty prediction. *Hydrological Processes* 6: 279-298.
- Brooks R. J, D. N. Lerner, and A. M. Tobias. 1994. Determining the Range of Predictions of a Groundwater Model which arise from Alternative Calibrations. *Water Resources Research* 30(11): 2993 – 3000.
- Edmonds, B and M. Mohring. 2005. Agent-based simulation modeling in social and organizational domains. *Simulation – Transactions of the Society For Modeling And Simulation International* 81(3): 173-174.
- Epstein, J. M. and R. L. Axtell. 1996. *Growing Artificial Societies: Social Science From the Bottom Up*. Washington D.C.: Brookings Institution Press and the MIT Press,
- Greswell, R. B., J. W. Lloyd, D. N. Lerner and C. V. Knipe. 1994. Rising Groundwater in the Birmingham area. In: *Groundwater Problems in Urban Areas* (ed: W. B. Wilkinson), 330-341. London: Thomas Telford.
- Kijima, K. and H. Hirata. 2004. Diffusion of Word-of-Mouth in Segmented Society: Agent-Based Simulation Approach. In: *Artificial Intelligence and Simulation (ed. T. G. Kim), selected papers from 13th International Conference on AI, Simulation, and Planning in High Autonomy Systems*, Lecture Notes in Computer Science (Subseries: Lecture Notes in Artificial Intelligence), 3397: 198-206.
- Lansing, J. S. 2002. ‘Artificial Societies’ and the Social Sciences. *Santa Fe Institute online working paper* No. 02-03-011, www.santafe.edu.
- Nelder, J.A. and R. Mead 1965. A simplex method for function minimisation. *Computer Journal*. 7: 308-313.
- Pidd, M. 1996. *Tools for Thinking: Modelling in Management Science*, Chichester : John Wiley & Sons.
- Reynolds, C.W. 1987. Flocks, Herds, and Schools: A Distributed Behavioral Model. *Computer Graphics* 21(4) (SIGGRAPH '87 Conference Proceedings): 25-34.
- Schmidt, B. and B. Schneider. 2004. The Reflective Control of Cognition and Emotion. In: *Proceedings of the Simulation Study Group Workshop* (ed. S. C. Brailsford, L. Oakshott, S. Robinson, S. J. E. Taylor). 145-154. Birmingham, UK: The Operational Research Society.

Waldrop, M. M. 1992. *Complexity: The Emerging Science at the Edge of Order and Chaos*, London: Viking (Penguin Group).

Yeh, W. W.-G. 1986. Review of Parameter Identification Procedures in Groundwater Hydrology: The Inverse Problem. *Water Resources Review* 22(2): 95-108.

AUTHOR BIOGRAPHIES

DONGFANG SHI Received a BSc (Hons) Economics from Beijing Normal University, China in 2000, and a MSc Management from Lancaster University in 2002. She is currently a PhD student at Lancaster University.

ROGER J. BROOKS is a lecturer in the Management Science department at Lancaster University. He received a PhD and MSc in Operational Research from Birmingham University and a B.A. (Hons) degree in mathematics from Oxford University. He is co-author of a textbook on simulation. His areas of interest include conceptual modeling, Boolean networks and agent-based simulation. His website can be found via <www.lums.lancs.ac.uk> and his e-mail is <roger.brooks@lancaster.ac.uk>