

Agent-based modelling — intelligent customer relationship management

N Baxter, D Collings and I Adjali

The intelligent customer relationship management (iCRM) tool — built using agent-based modelling techniques — aims to illustrate how CRM investments can influence a customer population, giving a clearer view of potential return on investment (ROI). Unlike conventional approaches, this model considers the communication of customer experiences between members of a social network, incorporating the powerful influence of word of mouth on the adoption of products and services. The tool is an advance on traditional techniques that rely on macroscopic behaviours and aggregated customer data, while neglecting important network and spatial effects.

1. Introduction

It is widely accepted that the way a company treats its customers plays a key role in future profitability. Customer relationship management (CRM) can take a variety of forms but the underlying strategy is to identify and manage the most valuable customer relationships. Unfortunately most enterprises do not fully understand how CRM interventions influence their customer base; in the past many CRM projects have failed to meet expectations.

Agent-based models (ABMs) consist of a population of discrete entities, each with individual rules of behaviour. Such models consider the relationships between each agent, and their interactions with the surrounding environment. These models can provide valuable insight into a variety of business problems.

Collings et al [1] formulate some of the basic concepts regarding agent-based modelling of product adoption. This paper presents an enhanced version of the technique and introduces the iCRM tool, a model that explores the impact of customer relationship management.

2. Agent-based modelling

Agent-based models are simulations based upon local interactions between members of a population, and between those members (agents) and the environment in which they are contained. In software these individuals are often implemented as objects and characterised by a set of parameters and behavioural rules. These rules consider changes in the environment, and define the interactions with other agents in the system. The technique has a wide

range of applications — agents can represent anything from vehicles in a traffic system to participants in a financial market.

The macroscopic behaviour of an agent-based system results from the multitude of interactions between agents and is often counterintuitive and difficult to predict. In such models an equilibrium state may never be reached, but the ability to analyse the underlying system and study the dynamics of its behaviour is of significant interest. The agent-based technique is distinct from other modelling approaches where characteristics are often aggregated and manipulated on a macroscopic level.

Agent-based models provide a natural framework to capture spatial and network effects in systems. Individuals can be associated with a geometric location and permitted to move around their environment where appropriate. The individuals can be located in a continuous space or restricted to a discrete grid-based geography. In the fields of social science and computational economics, agent-based modelling has provided an opportunity to move beyond the concept of a rational agent. The reasoning of individuals can be progressively refined, accounting for their limited knowledge and abilities. An agent-based simulation can also capture the flow of information between agents in a system, a critical factor in the construction of a social simulation. There have been many studies to investigate the links that associate one person with another within a population. This is known as a social network.

Agent-based modelling has been applied to several business areas, helping decision-makers anticipate market

dynamics, structures, and evolutionary paths. The simulations can help identify the causes of disequilibria and clarify sources of uncertainty. Recently, several agent-based models of electricity markets have been constructed [2—5] to test a variety of market structures. The electricity market complex adaptive system (EMCAS) model is a new agent-based tool developed by Argonne National Laboratory to study the complexities of restructured electricity markets [6]. The laboratory has also used the technique to simulate the health care industry, predicting future performance and evaluating alternative policies, equipment, and medications.

It is not only the research community that uses agent-based modelling; Sainsburys, the supermarket chain, has worked with SimWorld Ltd to create a model of customers in one of its stores. The results of the simulation have given the company insight into designing the optimum shop layout avoiding bottle-necks and generating more sales [7]. Macy's, the department store, has also developed a similar model in collaboration with Pricewaterhouse-Coopers to find the ultimate location for cash registers and service desks in their outlets [8].

3. Customer purchase modelling

There are numerous theories in social science and marketing concerned with product adoption, especially the diffusion of new innovations over time. These classical diffusion models take a macroscopic approach and generally make three assumptions:

- there are only two adoption states (consumers have either made their first purchase or have not yet adopted),
- the total number of potential adopters is fixed,
- there are no repeat purchases or multiple adoptions.

It has been observed that the take-up of many products and the spread of other phenomena such as disease can often be described by an S-shape curve [9—11]. In 1969 Bass [12] used diffusion theory to reproduce this S-shaped pattern for new consumer durables. He incorporated terms to account for the influence of mass media and interpersonal communications (equation (1)):

$$\frac{dN(t)}{dt} = p[\bar{N} - N(t)] + \frac{q}{N} N(t)[\bar{N} - N(t)] \quad \dots (1)$$

where $N(t)$ is the cumulative number of adopters at time t , \bar{N} is the total number of potential adopters, p is the coefficient of innovation and q is the coefficient of imitation. The first term in equation (1) represents adoptions due to innovators — buyers who are not influenced by those who have already bought. The second term describes adopters that are influenced by the number

of people who have already purchased the product; these adopters are referred to as imitators.

In order to predict or explain adoption rates for a particular product p , q and \bar{N} can be calibrated using an analogous product or a fit to existing data [13]. Bass [12] successfully applied this model to historical sales data for eleven innovative consumer durables including air conditioners, refrigerators, home freezers, black and white televisions, and power lawn-mowers. However, the macroscopic nature of diffusion models fails to capture the underlying system behaviour. The models do not identify the key levers or influencing processes within the system and the restrictive assumptions result in limited flexibility.

In an agent-based model representing a population of customers, each individual must have a cognitive process dictating their likelihood of adopting a product or service on offer. The individuals are influenced by environmental factors such as marketing and competition, and those who purchase the product will be exposed to a range of different interactions with the company. Agents can exchange their perceptions of a product or service along the links that form their social network, influencing the decisions of others within the population. This simulation approach gives control over the key drivers of customer behaviour allowing detailed exploration of the system dynamics.

3.1 The cognitive process

The simulation presented in this paper implements a multiple-stage decision process based on the adoption model outlined in Rogers [14] (see Fig 1). Each agent (consumer) within the model follows a specific instance of this process, completing a full cycle at each time step.

• Acquisition

During the first stage of the decision process an individual receives information that alters their perception of the product or service on offer. This change of perception is due to a combination of inter-agent communication and external factors such as marketing and competition.

• Decision

The individual now makes the choice of adopting or staying with the product if their perception is sufficiently high.

• Implementation

This step represents the explicit act of adoption or rejection, making the process particularly suitable for repeat purchase models.

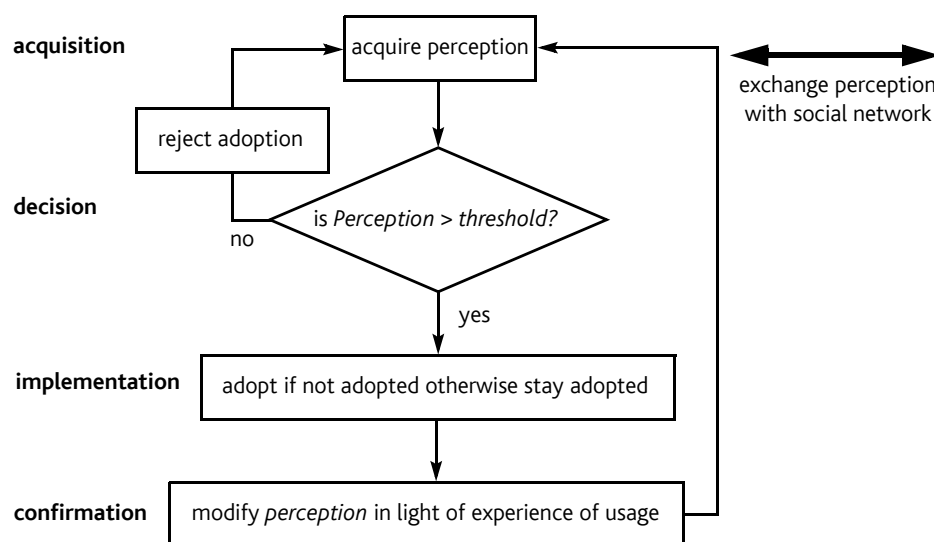


Fig 1 Flow chart outlining the decision-making process of individuals in an agent-based customer model.

- Confirmation

In this step the agents' parameters are adjusted based on experience of using the product or service.

In addition to the procedure described above it is important to consider the influence of social fads and other network externalities on product or service adoption. In these processes it is the knowledge of who within the population has adopted that generates the pressure to adopt, rather than the functionality of the product or service itself. The term network externality describes the circumstance where a product's benefit changes as the number of agents consuming it changes. A classic example of a product that exhibits network externalities is the fax machine — as more people purchase fax machines, users can communicate with a greater number of people, and the utility of the device increases.

In an agent-based customer-model, a pure fad process could replace the knowledge-driven procedure in Fig 1. However, this is unlikely to be realistic for most products and services. The concept can alternatively be implemented as an addition to the acquisition stage of the adoption process. At every time step in the simulation, individuals examine their social network and assess the proportion of their acquaintances that have adopted. For a pure fad the agent will adopt the product as soon as this proportion reaches their individual threshold value [15]; however, in the combined process, changes in the proportion of adopters will simply be an extra influence on the perception of the product.

3.2 Social network

In the early 1960s Milgram conducted one of the first empirical studies of social network structure [16]. He distributed a number of letters to a random sample of people in Nebraska. The letters were all addressed to a

single destination in Boston, Massachusetts. Milgram instructed that the letters could only travel by being passed from person to person, and could only be passed to somebody the current holder knew on a first-name basis. When the letters finally arrived in Boston, Milgram discovered that the average number of steps taken for the journey was six and labelled this phenomena six degrees of separation. Although Milgram's experiment was not scientifically rigorous, the concept of a small world is now widely accepted and is defined as a social network where chains of intermediate acquaintances are small compared to the total number of people.

The structure of networks can be modelled using graphs. Dots (vertices) are used to represent individual entities with lines (edges) connecting them. These linkages can form topologies ranging from regular lattices to completely random distributions. Early representations of social networks were often based on random graphs [17]; more recent studies have developed increasingly realistic characterisations [18].

If there are N people in a particular network — each with an average of z acquaintances — then a one-dimensional regular graph can be created by linking each node to its z nearest neighbours (Fig 2a). In contrast, a completely random graph is formed by drawing links between arbitrary pairs of individuals (Fig 2b). The variable z is known as the co-ordination number of the network.

Random graphs exhibit the small world effects observed in real social networks [19], but fail to capture clustering within a population. Realistically one would not imagine the links in a social network to be completely random; it is reasonable to expect that two people are more likely to know each other if they share a common acquaintance.




Fig 2 (a) Regular ring lattice with $N=8$, $z=4$ (each vertex is connected to its z nearest neighbours).
 (b) Random lattice with $1/2 Nz$ edges.
 (c) Watts-Strogatz [18] small-world (the dotted lines show the links that have been rewired).
 (d) Newman-Watts [20] small-world with 3 extra links (dotted lines).

The degree of grouping within a network is defined by the clustering coefficient C — the probability that two individuals with a common acquaintance will be connected. For a fully connected network C is obviously equal to 1 and for random graphs. Studies of real-world networks [18] show small-world characteristics but exhibit clustering coefficients significantly greater than for a random graph representation.

In order to model this grouped small-world structure [18], take a one-dimensional regular lattice and randomly re-wire the edges to introduce varying levels of disorder. Each link has a probability p of being re-positioned (Fig 2c); when $p = 0$ the original network is unchanged and for $p = 1$ the result is completely random. However, intermediate values of p provide a good representation of observed networks — the final topology is highly clustered with a short characteristic path length.

In a subsequent paper Newman and Watts [20] identify a number of problems with the Watts-Strogatz [18] model. Their most significant concern is that re-wiring links allows regions of the original lattice to become detached from the rest of the graph. To address this issue they suggest an alternative version of the model where links are added at random between pairs of vertices without removing any edges from the original graph (Fig 2d).

The small world model provides a reasonable network for simulating the flow of information between individuals in a social system. It can be used to construct realistic linkages when the exact structure of a social network is unknown, or create a more accurate representation when parameters such as the clustering coefficient have been determined.

4. The iCRM tool

Analysts predict that companies will spend billions of dollars in the next few years on software and services to help them manage their customer interactions more effectively. Unfortunately, most organisations do not fully understand how these investments will affect their

customer base and currently only 21% of CRM projects meet all expectations [21].

Word of mouth can be a powerful way for businesses to recruit new customers. People are often suspicious of advertisements and seek opinions from trusted friends and acquaintances before purchasing a product or service. Referrals are effective as they usually come from someone who is familiar with the product or service but has no financial motive for recommendation. Companies generally find that referred customers require less sales time to build trust and credibility, and tend to be more loyal than those whose purchases are driven by advertisement Griffin [22].

Some businesses have successfully manipulated word of mouth by running referral schemes offering benefits to existing customers who recommend friends. However, just as positive word of mouth can be a highly effective marketing tool, negative word of mouth can be destructive. Typically a dissatisfied customer will tell eight to ten people about their experience; one in five will share their dissatisfaction with twice as many [22]. This statistic is of particular interest in the CRM arena where technology-driven solutions are often sold on their ability to cut operating costs regardless of their impact on customer satisfaction.

The intelligent customer relationship management (iCRM) tool is an agent-based customer model built using the cognitive process discussed in section 3.1 and a Newman-Watts small-world social network (section 3.2). It is a decision support tool enabling organisations to visualise the impact of their CRM strategies and explore the effects of word of mouth on customer recruitment and retention.

The dynamic nature of the model makes it a unique approach to understanding customer behaviour. It can be used to recognize the future consequences of CRM implementations and estimate the nature of return on investment.

The *i*CRM tool is a generic model that captures the key drivers behind customer behaviour and choice, and can be used in a consultancy environment to facilitate dialogue. It is designed for use with a range of clients and is not specific to any company or sector. Trials show that using the tool can stimulate discussion and improve understanding of CRM issues. However, the *i*CRM model is not a precise forecasting tool. The agent-based approach offers a more realistic representation of a customer population, but it is still based on numerous assumptions and simplifications. In reality many of the parameters required to define a theoretical model are difficult to measure, and specific external events can cause unexpected behaviour. Results should be treated as illustrative but can still be used to compare the impact of different CRM strategies both in terms of market share and financial performance

4.1 Model structure

The *i*CRM model consists of a population of five hundred customers (Fig 3) and a single product. For simplicity the product is restricted to two parameters: price and quality. The agents within the model are heterogeneous and each has its own interpretation of the product's attributes, forming a distribution of perceptions within the population. When an individual has a combined perception that exceeds their internal threshold they will adopt (or re-adopt) the product on offer (section 3.1). The underlying adoption model makes this tool particularly suitable for subscription products or services where people review their decision to purchase at regular intervals.

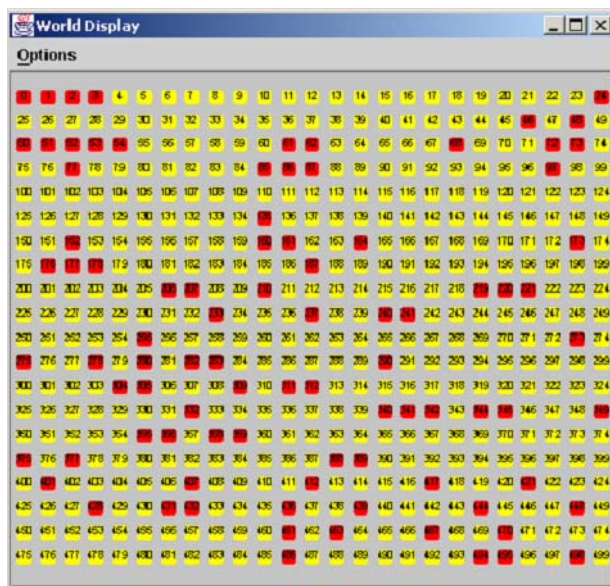


Fig 3 Screenshot from the *i*CRM tool showing a graphical representation of the customer population. The dark squares represent agents who have adopted and the pale squares are potential customers.

Each agent in the model is part of a social network through which perceptions are compared and influence

exerted on members of the population. Research [23] has shown that long-term customers are less likely to recommend a product or service than recent adopters. New customers are more conscious of their purchase — they want to talk about it. Over time, habit and familiarity take over and recommendation rates fall, an effect captured by the *i*CRM model by linking the probability of social interactions to length of adoption.

In addition to word of mouth, potential customers are influenced directly by external factors such as marketing, competition, and CRM interactions. At each arbitrary time step a proportion of agents are affected by marketing and sales material; for simplicity this is always assumed to have a positive impact on perception — the number of people targeted and the magnitude of the influence can be adjusted as appropriate. Although the model only deals with a single product, the effects of competition are an important aspect of the adoption process. This pull towards alternative products or services is captured in the model as a gradual erosion of perception at each time step and can be tuned to represent markets with differing competition levels. The most important influence in the tool is the effect of CRM experiences on consumer perception. This represents personal interactions between the company and customer such as complaints, repairs, billing, etc. These contacts could be made through a range of channels and can be defined in terms of their frequency and impact.

4.2 Using the *i*CRM tool with clients

In order to illustrate the impact of a particular CRM investment the model must first be calibrated to represent the current strategy of the company in question (referred to as the base case scenario). This is usually an iterative process that involves working with a client to populate the model's parameters with a combination of factual data and estimates based on qualitative discussions. Once the client is satisfied that the model is giving a reasonable prediction of future market share (Fig 4) they can consider CRM investments that change some or all of the parameter




Fig 4 Graph showing a fictitious base case for a business with an initial market share of approximately 25%.

values. The exact form of the base case is not critical to successful use of the model; it is improvements on this scenario due to CRM investments that are of key interest. The stochastic elements of the model are controlled using a sequence of random numbers. This list of numbers is identical for each run of the simulation so the pattern of interactions in the base case matches that in the scenario representing CRM implementation. Any changes in market share can therefore be attributed to the CRM solution rather than statistical variations. It is possible to run the model many times using a different sequence of random numbers in order to calculate the average expected impact of a particular CRM intervention.

In the model, CRM interactions are characterised by their frequency, cost, and the impact they have on customer perception. If the benefits of a CRM solution can be expressed in these terms then the appropriate parameters can be adjusted and the behaviour of the customer population examined. If the client can supply basic financial information, the tool can also compute an estimate of return on investment (ROI) at each time step.

In a simple version of the model, customer contacts can be classified as positive or negative and given fixed impacts as appropriate. Figure 5 shows the result of increasing the probability of a positive experience by 1%,




Fig 5 Graphs showing market share and ROI following CRM investment.

starting from the base case shown in Fig 4. It can be seen that there is very little impact on market share for the first year after investment — ROI increases very slowly. However, during the second year the decline in market share is greatly reduced; returns increase rapidly, reaching breakeven after approximately 3 years. This illustrates that it can often take time for improvements in customer service to have a significant effect. Customers must interact with a company over a period of time before they become aware of changes and it takes a further period for this improved perception to diffuse through the population. Although this example is based on fictional data, it illustrates the type of analysis enabled by the tool.

5. Conclusions

The way in which people adopt products and services is extremely complex. The decision to purchase is based on a multitude of influences, including complex interactions with other members of the population. Alternative techniques to the standard diffusion models are required to gain insight into the dynamics of such systems. This paper has introduced the concept of agent-based modelling and how it can improve upon conventional approaches when analysing customer populations.

The *iCRM* tool provides a generic model to explore the key drivers of customer behaviour, helping corporate decision makers understand the impact of their CRM strategies. The model allows clients to experiment with the levers that influence product adoption; the company can then identify these controls in real life and attempt to manipulate them.

The concepts applied in the *iCRM* tool could be extended to create client-based models built upon relevant qualitative and quantitative information specific to their market. These simulations could be used to analyse detailed scenarios and provide more accurate predictions.

References

- 1 Collings D, Reeder A A, Adjali I, Crocker P and Lyons M H: 'Agent based customer modelling: individuals who learn from their environment', Proceedings Congress on Evolutionary Computation 2000, San Diego, ISBN 0-7803-6375-2, pp 1492—1497 (2000).
- 2 Bower J and Bunn D W: 'A model-based comparison of pool and bilateral market mechanisms for electricity trading', Energy Journal, 21, No 3 (July 2000).
- 3 Petrov V and Sheble G B: 'Power auctions bid generation with adaptive agents using genetic programming', Proceedings of the 2000 North American Power Symposium, Institute of Electrical and Electronic Engineers, Waterloo-Ontario, Canada (October 2000).
- 4 North M J: 'SMART II: The spot market agent research tool version 2.0', Proceedings of SwarmFest 2000, Swarm Development Group, Logan, Utah (2000).
- 5 North M J: 'Agent based infrastructure modelling', Social Science Computer Review, Sage Publications, Thousand Oaks, California (Fall 2001).

- 6 North M J, Conzelmann G, Koritarov V, Macal C, Thimmapuram P and Veselka T: 'Multi-agent modeling of electricity markets', Proceedings of the Computational Analysis of Social and Organizational Science Conference 2002, Carnegie Mellon University, Pittsburgh, PA, USA — <http://www.casos.ece.cmu.edu/>
- 7 New Scientist, Firm Forecast, 24 (April 1999) — <http://www.newscientist.com/hottopics/ai/forecast.jsp>
- 8 Bonabeau E: 'Predicting the unpredictable', Harvard Business Review, 80, No 3, pp 109—116 (March 2002).
- 9 Fisher J C and Pry R H: 'A Simple Substitution Model for Technological Change, Technological Forecasting and Social Change, 2, pp 75—88 (May 1971).
- 10 Meade N and Islam T: 'Technological forecasting: model selection, model stability and combining models', Management Science, 44, pp 1115—1130 (August 1998).
- 11 Modis T: 'Predictions', New York, Simon and Schuster (1992).
- 12 Bass F M: 'A new product growth for model consumer durables', Management Science, 15, No 5 (January 1969).
- 13 Mahajan V, Muller E and Wind Y: 'New-product diffusion models', Kluwer Academic Publishers (2000).
- 14 Rogers E M: 'Diffusion of Innovations', The Free Press (1995)
- 15 Watts D J: 'A simple model of fads and cascading failures', Santa Fe Institute Working Paper (00-12-062) (2001).
- 16 Milgram S: 'The small world problem', Psychology Today, 2, pp 60—67 (1967).
- 17 Sattenspiel L and Simon C P: 'The spread and persistence of infectious diseases in structured populations', Mathematical Biosciences, 90, pp 67—383 (1988).
- 18 Watts D J and Strogatz S H: 'Collective dynamics of small-world networks', Nature, 393, pp 440—442 (4 June 1998).
- 19 Newman M E J: 'Small worlds', Santa Fe Institute Working Paper (99-12-080) (1999).
- 20 Newman M E J and Watts D J: 'Scaling and percolation in the small-world network model', Santa Fe Institute Working Paper (99-05-034) (1999).
- 21 Hewson Consultancy Group: 'Making a compelling business case for CRM (January 2000) — http://www.hewson.co.uk/sistrum/crm_management_insights/stream4/0001compelling_crm.pdf
- 22 Griffin J: 'Customer loyalty: how to earn it, how to keep it', Jossey-Bass Publishers (1995).
- 23 East R: 'Recommendation as a function of customer tenure', Conference Paper for ANZMAC 2000 (November 2000).



Nicola Baxter joined BT in 2001 after graduating with a Master of Physics degree from Manchester University.

Since joining, she has been working in the Strategic Analysis and Research unit, on the application of agent-based modelling to business and economic problems. Recent projects have been the development of a decision support tool to enhance understanding of customer relationship management and the analysis of trading arrangements in bandwidth markets.



David Collings was educated at Cambridge University and Imperial College where he was respectively awarded a PhD and first degree in Physics. In 1996 he joined BT at Adastral Park and has been working in the Strategic Analysis and Research (STAR) Business Modelling group, researching new tools and techniques for business modelling. He has published a number of papers in the fields of solid state physics and business modelling and is the author of a commercial software package.



Iqbal Adjali is a senior member of the Strategic Analysis and Research unit where he is currently leading the Computational Economics team. He has extensive experience with modelling techniques including game theory and agent-based simulations. In 2000, he led a team who provided consultancy to BT 3G Mobile and developed a simulation tool to prepare BT for the UK 3G spectrum auction. Iqbal completed an MBA at Cranfield School of Management in December 2002 and holds a PhD in Theoretical Physics from Oxford University. He worked previously in the areas

of distributed systems modelling, nonlinear dynamics and fractal modelling. He has published in fields such as evolutionary economics, game theory, fractal modelling, distributed systems modelling and particle physics.