

---

**Peter Furness**

*runs his own business specialising in the fields of decision analytics, modelling and data mining. A mathematician by background, Peter has been closely involved with the development and application of new analytical techniques in marketing and customer value management.*

# Applications of Monte Carlo Simulation in marketing analytics

**Peter Furness**

*Received (in revised form): 14th July 2011*

## Abstract

Monte Carlo Simulation (MCS), originally developed in the 1940s for use in nuclear weapons design, is playing an increasing role in commercial applications, including marketing and Customer Relationship Management (CRM). It provides an efficient way to simulate processes involving chance and uncertainty and can be applied in areas as diverse as market sizing, customer lifetime value measurement and customer service management. This paper examines the history of MCS and presents an illustrative example to explain the basic principles of the technique. Three case studies from marketing and CRM, which underline the importance of MCS to marketing analytics, are described. Some key issues, including the drawbacks and pitfalls of MCS, are covered. The paper also considers the future, with MCS applied in the digital world, and concludes with a review of relevant software tools.

*Journal of Direct, Data and Digital Marketing Practice* (2011) **13**, 132–147.  
doi:10.1057/dddmp.2011.25

**Keywords:** Monte Carlo Simulation, customer service management, predictive analytics, software testing, CRM

## Introduction

The technique of Monte Carlo Simulation (MCS) was originally developed for use in nuclear weapons design. It provides an efficient way to simulate processes involving chance and uncertainty and can be applied in areas as diverse as market sizing, customer lifetime value measurement and customer service management. MCS is playing an increasing role in commercial applications, including marketing and Customer Relationship Management (CRM).

The importance of MCS to marketing analytics can be summed up in one word ‘integration’. It enables you to integrate together a number of components, where each of the components may have been developed using conventional methods (such as predictive or descriptive analytics), to provide a complete model. Such integration may be very difficult, inefficient or even intractable, using more traditional methods. The classic example is customer value modelling, where a number of models of customer behaviour need to be integrated together to provide a complete customer view over a period of time.

## The importance of Monte Carlo Simulation (MCS) lies in integration

Peter Furness  
Peter Furness Limited  
E-mail: [furnesspm@gmail.com](mailto:furnesspm@gmail.com);  
Skype: [furnesspm](#);  
Website: <http://www.peterfurness.co.uk>

In this paper, we shall look at the history of MCS; present an illustrative example; examine three case studies from marketing and CRM, which underline the importance of MCS; look at some key issues, including the drawbacks and pitfalls of MCS; consider the future with MCS applied in the digital world; and, finally, review relevant software tools.

**MCS can help solve problems that are intractable using traditional mathematical analysis**

### Monte Carlo Simulation

MCS was invented around 1946 in the USA by two émigré mathematicians — Stanislaw Ulam and John von Neumann — who, at that time, were working on the early stages of the hydrogen bomb development.<sup>1</sup> Ulam realized that one could model the diffusion processes involved in nuclear explosions by tracking the paths of a sample of individual particles, rather than by using the traditional mathematical analysis (involving the numerical solution of systems of partial differential equations), which had proved intractable. Each particle would be subject to random collisions and other interactions, but could be followed through time and the results subject to statistical analysis. Together with von Neumann, he worked out the theoretical basis for MCS and coined the name ‘Monte Carlo’ because of the way chance is handled by the technique.

MCS is a computer-intensive technique and, for this reason, its adoption has gone hand-in-hand with the increasing cost effectiveness of computers. It is very widely used in areas as diverse as traffic flow simulation, deep coal mining logistics, post office queuing protocols, hazard analysis of chemical plants, marketing and CRM, and risk analysis in global finance. See<sup>2,3</sup> for more information.

**The discrete event-based simulation has a number of key ingredients**

MCS comes in many flavours, but perhaps the simplest — and the one we shall be studying — is the so-called ‘discrete event-based simulation’, which has the following key ingredients:

*A clock:* which measures time and is used to determine when things happen in the simulation.

*Entities:* are things about which we need to record information (attributes) during the course of the simulation. For example, customers in a post office would be modelled as ‘entities’, as would the counter clerks. Entities may interact — for example, when a customer is being served.

*Events:* are measured by the clock and coincide with changes in the attributes of entities, or the birth and death of entities. For example, a customer’s arrival in the post office would be an event.

*Processes:* a process is a set of rules, which is triggered by an event, which may generate subsequent events and which is associated with one or more entities. For example, the serving of a customer by a counter clerk would be a process.

*Random Number Generator:* is used in conjunction with various statistical distributions to generate events and execute processes. The Generator underpins the probabilistic (hence ‘Monte Carlo’) nature of the simulation.

*Event Scheduler*: is used in conjunction with the other elements of the model to manage the generation of events and the execution of processes. The event scheduler is usually rule based.

**The evolution of entity attributes is tracked over time and stored for analysis**

MCS is run (usually on a computer) with a set of starting entities and attribute values and its evolution tracked over time by storing the attribute values of selected entities in a data set for subsequent analysis. Some summary statistics are also tracked over time (eg queue size in the case of our post office example). Depending on the assumptions being evaluated, multiple runs may be needed, each with a different random number stream so that the ‘volatility’ and ‘average’ behaviour of the underlying process can be understood. The following example will illustrate the principles in more detail.

**Snakes & Ladders is ideal for illustrating the basic concepts of MCS**

### **An illustrative example — Snakes & Ladders!**

Anyone with young (or not so young) children, or who has to entertain children, will be familiar with the board game of Snakes & Ladders. It is ideal for illustrating the basic concepts of MCS and you may learn some interesting facts about the game — which can stand you in good stead the next time you play it!

Snakes & Ladders is a game of pure chance for one or more players. Each player has a counter, which moves around a board consisting of a 10x10 grid of squares. A dice is thrown to determine how many squares each player moves in a turn. If a player lands at the foot of a ladder, they jump to the top of the ladder; if they land on the tail of a snake they go down to the head of the snake. If a six is thrown, the player has another move. The winner is the first player to land on square 100 in an exact throw. If players overshoot, then they bounce back by the overshoot amount. Players can share the same position on the board. A schematic of the playing board for a popular version of the game is shown in Figure 1.

**MCS can provide answers to important questions about the game**

Some interesting questions arise: How long is the average game? How long can a game get? How much longer (or shorter) are games with 3, 4, 5 and 6 players? Is there an advantage in starting first? These are all important questions for parents who have to plan mealtimes around the game, or who have to arbitrate in disputes about potential advantages in the order of play!

We can use MCS to model the game as a simulation. In this case, the entities are the players’ counters and we can represent their position on the board (the only attribute) as an array. The clock is represented by fixed units of 20 seconds, which is the assumed time it takes for each player to make their move. Each dice throw is an event, and the moves around the board are the processes. Random numbers in the range 1–6 are used to represent the dice throws. A diagram illustrating the simulation process is given in Figure 2.

The simplest way to simulate the game is to run a large number of single-player games — say 100,000 — and then group the results in pairs to mimic a two-player game, in threes to mimic a three-player game and so on. This is valid because the players do not interact

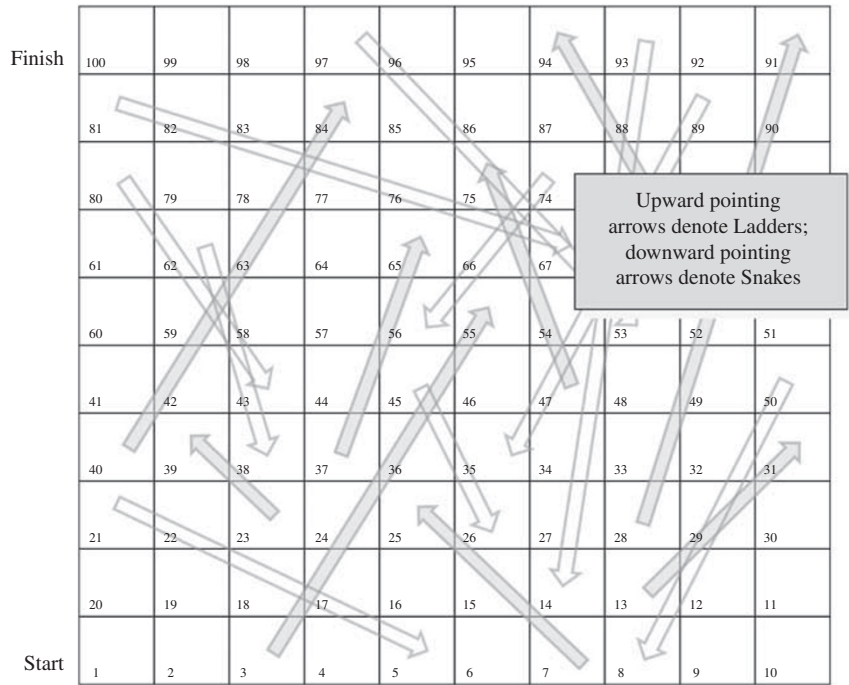


Figure 1: Schematic for Snakes & Ladders game

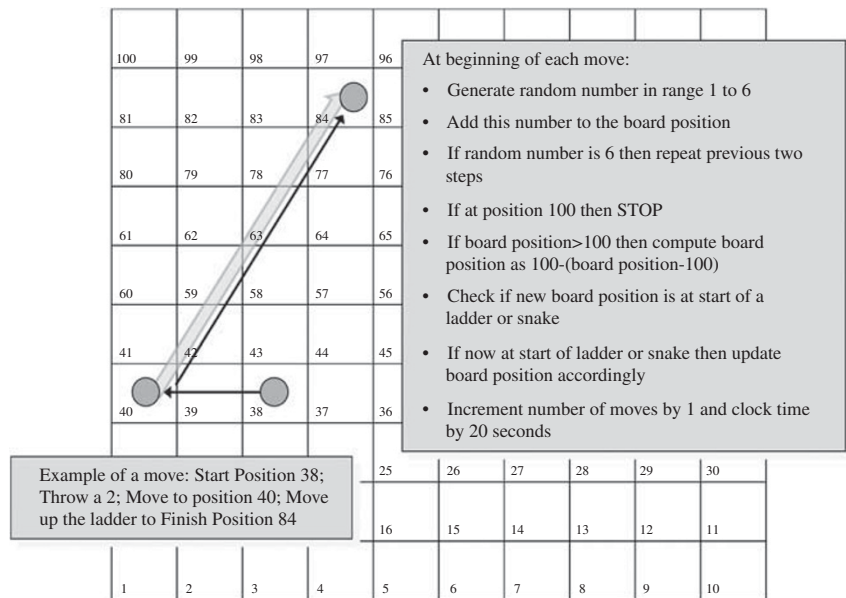


Figure 2: Process flow for Snakes & Ladders simulation

during the play (at least not in terms of counters on the board!). In the case of a two-player game, the game is deemed to end at the earliest ending of one of the two games. The results are stored in a data set and used to generate the summary statistics we need.

Figure 3 shows the distribution of game lengths in terms of number of moves for a two-player game. It also has some summary statistics calculated on the basis that each player's move takes 20 seconds. This

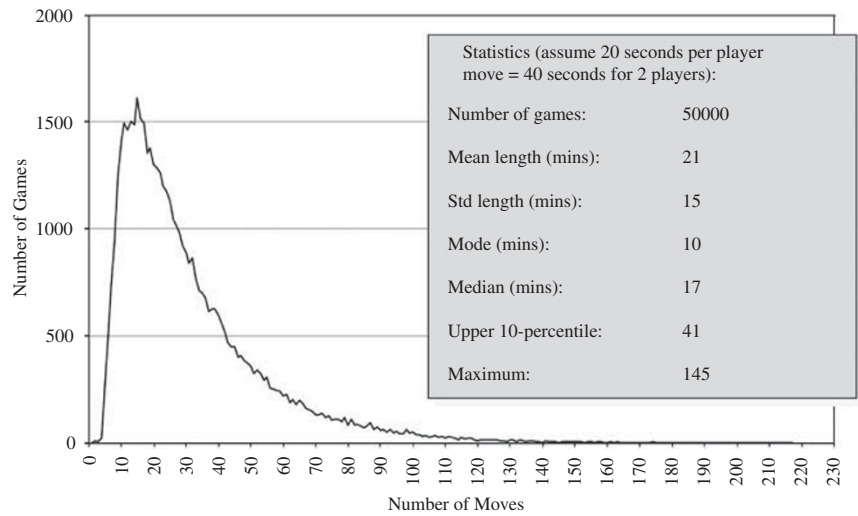


Figure 3: Distribution of game lengths — two-player game

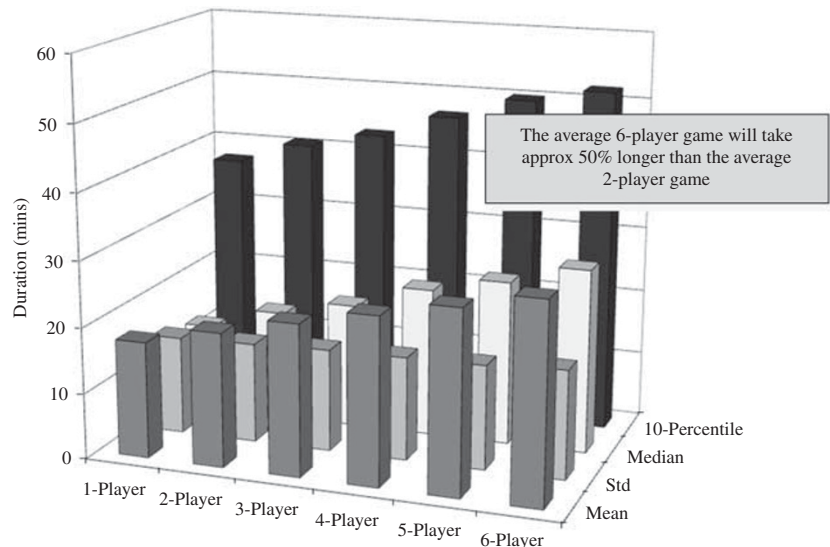


Figure 4: Game durations — multi-player games

shows that the average two-player game will take around 21 minutes. However, 10 per cent of games can take more than 40 minutes and at least one game out of the 50,000 took a staggering 145 minutes (though games of any long period of time are theoretically possible!).

Figure 4 shows the results for multi-player games of increasing number of players. The more players, the longer the games tend to be — a six-player game, for example, takes about 30 minutes, which is 50 per cent longer than the average two-player game. It turns out that the number of moves (all players taking a throw in each move) actually reduces with more players because the chances of at least one player hitting the jackpot in any given move increase. However, this is more than balanced out by the extra time needed for all the players to take their throws in each move.

**MCS can answer questions about the game which are intractable by any other method**

Our final question relates to the advantages of going first. We have answered it in the case of the two-player game. Common sense tells you that there must be some advantage, because the first player still wins even if both players finish in the same move. Calculating the probability of both players finishing in the same move is impractical using any method other than an MCS. We performed a simulation of one million games and found that the first player to move wins in a proportion of 0.5046 games, which is approximately 1 per cent more than the player who goes second — and a smaller percentage than many people would expect!

**The simulation can easily be adapted to cover variations on the basic game**

It would be very easy to extend the simulation to cover issues such as a non-constant time taken to make throws, or different board structure, or different rules about double sixes and so forth. One could even define measures of ‘customer satisfaction’ — such as the proportion of total player moves involving snake/ladder events. MCS is a very cost-effective technique for exploring these issues in the context of the Snakes & Ladders game.

In summary, we can see that MCS is a very straightforward way to analyse quite a complicated system in which chance plays a key role — in this case the Snakes & Ladders game. In addition, variations on the basic system can be evaluated quickly and easily. The author is not aware of any practical alternative to MCS in this context.

## Case studies

**Customer behaviour evolves as the products owned and the marketing received by the customer change through time**

### Case study 1 — Integrating models of customer behaviour

One of the challenges in developing estimates of customer future value is the complexity of the future purchasing behaviour of customers. As we move forward in time, a customer’s behaviour will be strongly associated with the products they have acquired up to that point and the contact strategies used by the company to promote product purchase. For example, if a bank customer acquires a mortgage, then their propensity to take up a personal loan, or other form of credit, is likely to be different from that of a customer who has not acquired a mortgage. The propensity will clearly also be affected by direct marketing and other forms of advertising of personal loans. Often companies will have propensity models conditioned upon existing product holdings (as well as other factors such as demographics, contact history, lifestyle, etc, if appropriate data are available); the problem is then how to integrate these models together to obtain estimates of future value.

**MCS can be used to model customer value over a specified time period at the individual level**

One way to do it is to use MCS. For each customer, we can simulate their product-purchasing trajectory over time, say, over 5 years. This is shown schematically in Figure 5.

By assigning values to initial product holdings and by running the simulation many times with different random number streams, we can obtain a distribution of future value — from which we can derive metrics such as average future value, standard deviation, median and percentiles at a customer level. Such customer-level estimates are then available for targeting purposes. The simulation proceeds in, say,



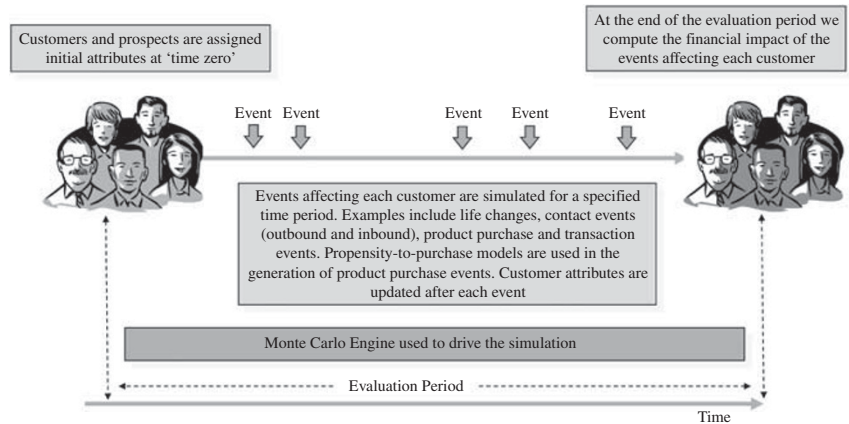


Figure 5: Using MCS to integrate predictive models

quarterly time steps and the available propensity models are used to determine product take-up in each time step. Each propensity model provides a probability, say  $P$ , of take-up in the current time step; we decide on whether or not the product is actually taken up by sampling from a binomial distribution with parameter  $P$  (this can be implemented by choosing a uniformly distributed random number, say  $x$ , in the range  $[0, 1]$ ; if  $x$  is less than  $P$ , then the product is taken up, otherwise it is not). The new product-holding characteristics are carried forward to the next time step. Rule sets are used to arbitrate between products when multiple product take-up in any time step is allowed.

Producing value metrics at a customer level using MCS is valuable for targeting applications but is also computing intensive, as many separate runs — each with a different random number stream — are required. However, for many applications, where aggregate behaviour at a group or segment level is required, it is not necessary to do repeated simulation runs for each customer. Some examples would be:

- Calculating customer value by demographic or product segment. If the segments are reasonably large — say, several hundred or a few thousand customers — then aggregating single simulations across each segment will yield reliable metrics at the segment level. Confidence intervals for each metric can also be computed in this case.
- Evaluating the financial impact of two or more alternative contact strategies. In this case, a separate simulation is run for each contact strategy. A representative but modestly sized sample of customers should be sufficient to provide reliable levels of financial discrimination for deciding between strategies.
- Producing synthetic customer data for a variety of applications, including:
  - Training marketing teams in the financial evaluation of campaigns. The simulation can be used to generate synthetic campaign outcome

**Modelling using MCS at the segment level is less computing intensive and offers many valuable applications**

data under a wide variety of assumptions about the effectiveness of particular contact strategies. This is an excellent way to create detailed numerical examples, which the staff under training can use in ROI and other financial calculations.

- Training in the use of marketing analytics and other CRM software tools — where a carefully controlled data environment — is essential — or where the company cannot mobilize the relevant data from its existing systems.
- Business modelling. A synthetic customer base, which mimics the company's own base — but which also embodies behavioural assumptions and models — can be used to explore 'what-if' scenarios. For example, exploring the impact of changes in business practices or developments in the external marketplace. Such an artificial base will typically incorporate actual customer data in conjunction with synthetic data from an MCS.

### Case study 2 — Customer service

Our second case study concerns call centre resource scheduling for a mobile telecommunications operator. The company needed to examine resourcing (ie call agents and their shift patterns) in relation to a number of key marketing issues including:

- The impact of new customer prioritization scenarios, for example, based on customer value segmentation.
- The trade-off between resource cost and quality of customer service.
- The effects of additional call volume stimulated by outbound direct marketing.

This analysis had to be done within the context of an existing volatile inbound call volume, which varied by time of day and day of week (see Figure 6).

We created an MCS for the call centre operation in which parameters were set using available operational data (eg call arrival patterns and agent service time distributions). The overall structure of the simulation is shown in Figure 7. Note that it would have been risky and expensive to have experimented with the actual call centre, hence the need for a simulation model to explore new ways of working in a cost-effective manner.

In this simulation, the principal entities included:

- *Caller*: with attributes such as caller ID and caller segment code (eg customer segmentation characteristics linkable using the caller's mobile number).
- *Call*: with attributes such as initial incidence time, time in queue (or time to abandon), time in service, etc.
- *Agent*: with attributes such as agent ID, agent segment code and call centre identifier.
- *Service*: with attributes such as start time, finish time, etc.
- *Queue Status*: with attributes such as number of callers in queue.

**MCS can help answer many important business questions about call centre resourcing and the impact this has on customer service**

**Simulation offers a cost-effective alternative to risky experimentation in the 'real world'**






Figure 6: Call centre — Inbound call profile

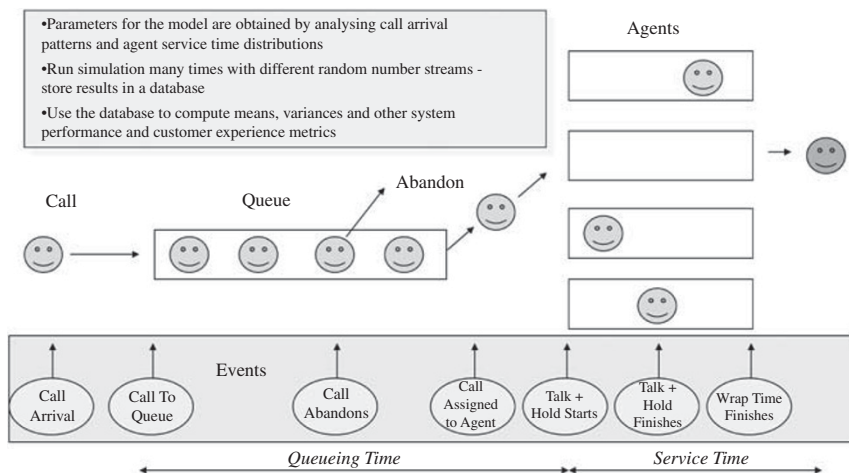


Figure 7: Call centre simulation — Overview

The events included:

- New call arrival
- Arriving call to queue
- Call abandons in queue
- Call assigned to agent
- Call talk and hold time commences
- Call talk and hold time finishes
- Call wrap time finishes

The processes included:

- Incoming call process: determines when the next call event will occur.
- Call to queue process: takes a call waiting to be queued, adds a start time for placement in the queue and time of abandon once in the queue. The current queue status is written out.

- Call abandon process: causes calls to be abandoned if the abandon time has been reached. Outputs a call record to the call data set. The current queue status is written out.
- Queue process: assigns a call to a server. The current queue status is written out.
- Service process: sets up the service time (talk + hold + wrap) for a call with an agent.
- Call finish process: outputs a call record to the call data set and updates the status of the call to completed and the service status to in wrap.
- Service finish process: outputs a service record to the service data set after updating the service status to completed and frees up the agent slot.

The simulation was run many times with different random number streams and the results stored in a database. Selected key metrics such as call waiting time distributions, call abandon volumes, agent utilization and service times could then be computed. By then running the simulation with different parameter settings designed to explore the above marketing issues, it was possible to develop a good future plan for agent resourcing, as well as to examine the operational impact of additional call centre capacity.

It is also worth noting that this simulation could be adapted relatively easily to explore the impact of a number of other factors that influence call centre operation, sales and customer service levels, for example:

- Varying offers by customer segment or by customer service experience in order to boost conversion rates.
- Targeting outbound marketing in order to stimulate a particular pattern of inbound traffic at the call centre.
- Changing the procedures for call handling at the agent level, including the use of differentiated messaging and associated scripting, in order to improve workflows.

### Case study 3 — CRM software testing

This case study looks at how one can apply actual customer data in conjunction with synthetic customer data, generated using MCS, to the evaluation of customer analysis software.

The client was implementing a CRM programme, part of which was the acquisition of CRM software to manage sales and marketing, customer service and product installation. The marketing requirements included the need for analytics tools for building descriptive and predictive customer models. We assisted with the evaluation by designing and implementing a series of tests of the analytics capabilities of the CRM tools.

One of the tests involved the CRM vendors using their toolsets to build a predictive customer model using a combination of actual and synthetic customer data, which we supplied. A sample of 25,000

**The basic simulation can be adapted relatively easily to answer many more marketing and operational questions**

**Synthetic customer data generated using MCS can be used for testing the performance of analysis software**

**Synthetic data provides a 'best possible' benchmark against which to compare the performance of a range of tools**

**In common with other modelling techniques, MCS faces the important challenges of 'verification' and 'validation'**

**Although MCS requires a lot of computing power, modern desktop systems are adequate for many marketing applications**

customers was created representing customers contacted as part of an upsell campaign, together with a flag to indicate their take-up/buy. A hold-out sample without the flags was also supplied for vendors to score. The customer data included real product usage history, as well as demographics and geodemographics. The take-up flag was synthetic and created using a very simple MCS, which took a propensity-to-buy function (essentially a probability), which we had defined and then, for each customer, sampled from a binomial distribution with this probability to obtain the flag, using essentially the same method as in Case study 1.

The propensity-to-buy variable was removed from the data set before it was supplied to the vendors. Vendors then had to build a propensity-to-buy model and return a scored hold-out sample. We then merged the synthetic buy/not buy flags back into the vendor-supplied hold-out sample and created gains charts and other performance metrics such as the Gini coefficient and maximum lift. Crucially, our propensity-to-buy function provided an absolute 'best possible' benchmark against which to compare each of the vendors' offerings. This proved to be a very severe test for all the vendors.

Table 1 shows the anonymized results for the evaluation of toolsets from four vendors. As can be seen, the performance on the hold-out sample was just one of the six criteria used in the evaluation. Each criterion carries a weight and the weighted scores across all criteria are used to compute a total score for each vendor. In this case, the performance on the hold-out sample carried 50 per cent of the total weight.

### Key issues

As with any modelling approach, MCS faces the important challenges of 'verification' and 'validation'. Verification concerns ensuring that the implemented simulation follows the specification designed for it. As such, this is no different from verifying that any computer program follows its specification. Housekeeping the progress of entities in the model can be especially problematic, and specialized simulation modelling software can help enormously with this task.

Validation involves making sure that the simulation satisfactorily represents the real world that it is intended to model. The representation will always be an approximation; the important thing is to ensure that the representation is sufficiently close, so that the relevant business questions can still be answered. There are two elements to validation — validating individual components of the model (eg checking that the average move duration is 20 seconds in our Snakes & Ladders simulation) — and validating the overall system. Both elements of the validation require empirical data about the real-world system we are trying to validate — we then compare this data with the results from the simulation (of the current system) to check whether the model is sufficiently accurate.

A third key issue is that of computational burden. Running a simulation multiple times, perhaps with many thousands of different random number streams, makes significant demands of computer power

**Table 1:** Marketing analytics tool testing results

Criterion	Vendor 1	Vendor 2	Vendor 3	Vendor 4	Criterion weight	Comment
1. Marketing scenario fit	63	81	81	73	0.125	Percentage of the total analytics requirements across the six marketing scenarios, which are supported by the tool.
2. Ease of use of tool	20	85	70	65	0.125	Ease of use of the tool in carrying out each of the analysis steps. This will be a subjective assessment with a percentage score for ease of use based on what we see at the workshops.
3. Usefulness of segments	30	30	30	30	0.125	Usefulness of the segments produced in Step 1 of the analysis including the quality of the tables/graphics used to describe the segments. Again this is partly subjective but we will use the results of our own analysis as a benchmark. A percentage score for usefulness will be given.
4. Model building methodology	10	70	70	70	0.125	The model-building methodology of Step 2 and completeness of the model performance metrics will be assessed as a percentage score. Providing the basic principles of training set/test set partitioning and evaluation of metrics on the test set are adopted then a high score should be obtained.
5. Performance on hold-out sample	51	62	75	83	0.500	The actual performance of the propensity to respond model on Dataset 3 will be judged against the 'best possible' model. We will do this by computing the Gini Coefficient (an overall measure of model performance) for the supplied scores and comparing this with the 'best possible' Gini obtained from the simulation used to generate Dataset 3. A percentage will be obtained from the ratio.
Weighted total (out of 100)	41	64	69	71	1.000	

if the results are required within a reasonable time. This is one of the major drawbacks of MCS, and computational requirements need to be carefully planned from the start. However, with the increasing power of computers — especially desktop computers — the scope of simulations, which can realistically be implemented by marketers, continues to grow.

**Looking to the future**

Another area where MCS has an important role is in the so-called Bayesian Network Analysis. Bayesian Networks (named in honour of the 18th Century mathematician, the Reverend Thomas Bayes, whose famous 'inverse probability theorem' is a cornerstone of probability theory) provide a framework for modelling the causal links between variables in a problem (see Jensen<sup>4</sup> for more background). In the context of customer modelling, Bayesian Networks provide a mechanism for updating belief (expressed, eg as a probability) in the way a customer will behave as new evidence arises. This is particularly

**Bayesian methods are playing an increasingly important role in the digital world**

relevant in the digital world where customers leave a trail of new information in real time.

In a Bayesian Network, variables are represented as the nodes in a graph or network and the causality is defined by the links between nodes. You have the notion of parent and child nodes of each node in the network. Cycles in which you can follow one node back onto itself through the network are forbidden (this is equivalent to respecting common-sense notions of causality).

Each variable can be in one of a number of states. Most early work in Bayesian Networks assumed finite state variables, although, more recently, continuous and infinite-discrete state spaces are being modelled.<sup>3</sup> Conditional probability distributions are used to define the relationships between a variable and those that link to it.

In a Bayesian Network, it is assumed that the nodes that are not directly descended from each other are independent in a statistical sense. The crucial theorem states that from the conditional probability distribution of each variable given its parents, one can compute the joint distribution of all the variables in the network.

Once all these conditional probability distributions have been defined, then the network can be used for ‘probabilistic inference’. If, for example, you know the values of one or more variables (this is known as ‘evidence’), you can use the network to compute the probability distributions of other variables conditional upon this new evidence.

**Bayesian networks can be used to encode prior knowledge about customer attributes and their interrelationships**

As already noted, Bayesian Networks can be applied to customer modelling. Figure 8 shows a hypothetical example in which our prior knowledge about customer demographics and product take-up is encoded in a Bayesian Network.

The network captures the common-sense knowledge that the number of children an individual has depends directly on age and marital status but does not depend directly on gender or occupation. Occupation, by contrast, depends on age, number of children and gender. Ownership of Product A (perhaps a mortgage product) is described as depending on home status and income.

Data mining would be used to estimate the conditional probabilities in this network. Where suitable data are unavailable, we could use our

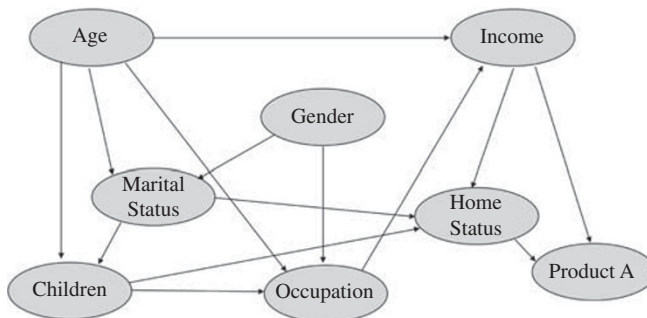


Figure 8: Bayesian network example

best-guess estimates initially (and these would be refined over time when results can be compared with actual outcomes). Once the probabilities have been determined, the network can be used for probabilistic inference — for example, to compute the probability that Product A=yes, given evidence about some or all of the other variables. For those customers who do not already hold Product A, this probability might then be used as a propensity-to-buy score for targeting purposes.

The prior knowledge built into the network in the form of the causal structure reduces enormously the number of parameters that would otherwise need to be determined by the data mining. If, for example, we had to assume the most general form of joint probability distribution for the variables, then the number of parameters to be estimated could be very large.

A particular advantage of the Bayesian Network model over a more conventional predictive model is that we can use the model to compute the conditional probabilities under any conditions of evidence — including situations in which there is no evidence for some variables. This is useful where customer data may be missing, or in situations where data are collected in real time and decisions have to be made quickly with as little data as possible — for example, in determining what to offer a customer in an interactive web session, or in an inbound telesales situation.

For complex networks, the inference process (updating probabilities in the light of new evidence) can be very difficult, with the algebra associated with the multiple application of Bayes's Theorem intractable. Approximations using MCS are then necessary. A popular variant of MCS used in this context is 'Gibbs Sampling', also the 'Markov Chain Monte Carlo'.<sup>3</sup> This use of MCS is highly technical and relies on the mathematical properties of a class of stochastic processes, known as Markov Chains, which can be used to estimate the conditional probabilities in a Bayesian Network. It is not possible to give an intuitive description for how the MCS operates in this case, unlike our other cases studies. The likely wider adoption of Bayesian methods in customer modelling in the digital world will, however, inevitably lead to greater use of MCS.

### Software tools

Figure 9 gives a brief summary of some of the tools available, which can help with application of the MCA. They fall into three main categories:

- *General purpose tools*: Basically any tool that has a random number generator, a scripting capability and the ability to handle arrays can be used to program an MCS. But beware potential difficulties with verification particularly the housekeeping associated with entity histories.
- *Specialized simulation modelling tools*: These have extensive housekeeping and model visualization capabilities that can

**Bayesian methods are particularly suited to modelling behaviour in a way which takes account of the digital data trail left by customers**

**MCS provides a technical solution to the difficult problem of updating a Bayesian network in the light of new evidence**

**Specialized MCS tools can dramatically reduce the time needed to implement a useful model**



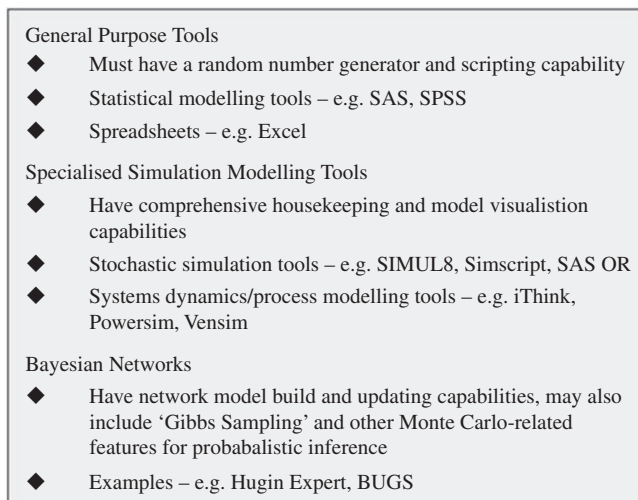


Figure 9: Software tools for MCS

dramatically reduce the time needed to implement a useful model. The downside is that they can be expensive and may not always have the flexibility to handle the particular situation one is dealing with. They may also be slower to run than a bespoke solution built with a general-purpose tool.

- *Bayesian Network tools*: This is still a very immature market. There are some excellent tools available, but one may expect rapid developments over the next few years, especially in the application of MCS for probabilistic inference with complex networks.

### Conclusion

MCS provides a way for analysing complicated systems in which chance plays a key role. In marketing analytics, it has a range of important applications, including the integration of customer behaviour models, the modelling of customer service processes and the evaluation and selection of analytical software tools.

There are a number of MCS software tools to support the analyst but beware of issues such as ‘verification’ and ‘validation’ and the computational burden of MCS — it is advisable to get expert advice before embarking on a simulation development.

Looking into the future, developments in the digital world will soon force more widespread adoption of MCS.

### Acknowledgments

This paper is based on a presentation given by the author at the seminar ‘Marketing Analytics 2010 — How to make the most effective use of Advanced Analytics Techniques’, organized by Henry Stewart Conference Studies and held on Friday, 19 November 2010 at the Mayfair Conference Centre, London W2 2ES.

**It is advisable to get expert advice before embarking on a simulation development**

**References**

1. Rhodes, R. (1995) *Dark Sun — The Making of the Hydrogen Bomb*, Touchstone/Simon & Schuster, New York.
2. Vose, D. (2008) *Risk Analysis: A Quantitative Guide*, John Wiley and Sons Ltd, Chichester.
3. Gamerman, D. (1997) *Markov Chain Monte Carlo: Stochastic Simulation of Bayesian Inference*, Chapman & Hall Texts in Statistical Science, London.
4. Jensen, F.V. (1996) *An Introduction to Bayesian Networks*, UCL Press Ltd., London.