Complexity Modelling in Marketing: A Look "Under the Hood" of the NBD

Cullen Habel and Steve Goodman

This paper demonstrates how agent based modelling can be utilised to model a complex marketing system. The Negative Binomial Distribution (NBD), a well validated, widely accepted marketing science model is used as the foundation to demonstrate the approach. Using the agent based approach individual shopper (agent) rules are defined that comply with the underlying assumptions of the NBD model. Using these rules and progressing from a simple to a more comprehensive picture achieves two objectives. Firstly the agent based approach's contribution to marketing is demonstrated and secondly it is possible to 'see under the hood' of the NBD model, namely what drives it and how it works. This is designed to assist in general terms, through demonstrating 'how' the NBD actually works, as well as demonstrating an approach that might lend itself to further exploration. By using the NBD as a foundation, the opportunity exists to move beyond the constraints of stochastic market modelling and incorporate new agent rules based on other theories (or even conjectures), allowing a means to develop falsifiable hypotheses where previously these theories and conjectures may have been untestable.

Keywords: NBD Model, Agent based modelling

Introduction

Markets exist as complex systems with almost infinite combination of products, outlets, advertising messages, price offerings, and consumer characteristics. An overarching principle of marketing and economics is that individuals endeavour to satisfy their own set of needs. Whether their goal is to maximise their outcomes or to satisfice (Simon 1978), consumers tend to follow their own sets of rules as semi-autonomous agents within a marketplace.

This paper introduces the concept of agent based modelling, then uses the distributional assumptions from the broadly validated Negative Binomial Distribution (NBD) model of repeat purchase (Ehrenberg 1959, Morrison 1969, Morrison and Schmittlein 1988) as the two agent rules for a basic agent based simulation. This simulation method is then applied to two markets (chocolate bars and detergents) using likelihood estimates for the markets' NBD parameters. This allows for validation of the method, which is done by the comparison of observed frequency distributions with against a set of five outputs from the agent based simulation.

From this point of validation, extensions are proposed beyond the strict foundations of the NBD framework. It is argued that this is where substantial promise exists for Agent Based Modelling as a predictive tool and as a method for testing theory.

Agent Based Modelling

Agent based simulation considers the perspective of individual actors within a system, as outlined in Tesfatsion's (2003) comprehensive review. Agent based modeling (ABM) 'is a bottom up approach to understanding systems which provides a powerful tool for analyzing complex non-linear systems' (Twomey and Cadman 2002). In this emerging research area Tesfatsion (2003) cites several studies of equilibrium pricing (Marks 1992), financial markets (Tay and Linn 2001) and electricity auction markets (Bower and Bunn 2001).

In agent based modeling, each actor has a degree of "agency" (Schieritz 2003) that is consistent with the level of effect the actor has on events within the system – hence the description of these actors as "agents" within the ABM literature. Each of the various agents works to a set of decision rules from their own perspective and their interactions determine the macro picture that emerges.

Agent based modelling provides an understanding of these self-organising systems by defining a set of individual level "agent rules". The method involves the simulation of the system; first by defining the set of agent rules and then by allowing the system to generate an overall output. Figure 1 gives the simplest possible example (the physical length of a line of ten people waiting for a bus). If each individual adheres to a rule whereby their personal space is equal to a half a metre, the length of the line will be 5m. For a different individual rule – say a personal space of 0.7m – the length of the line would be 7m and so on, allowing us to extrapolate individual rules into system level outputs. Thus, a self-organising system has led to a stable system output.

Figure 1: A simple self-organising system (bus queue)

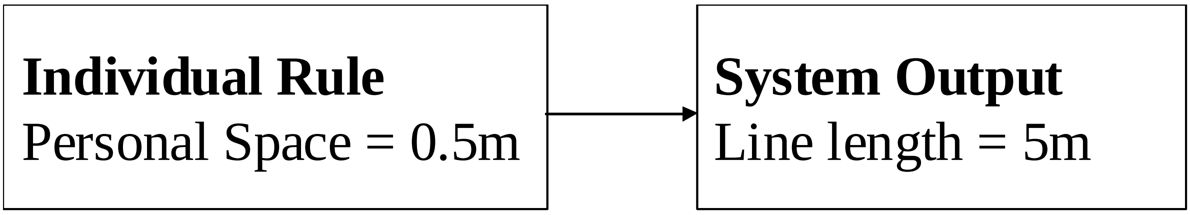

One might imagine how this approach could apply to the behaviour of a population of shoppers. As a brand manager of "A1 Cola" one may think of ten individuals and their purchase behaviour of "A1 Cola" over a week. In this simple scenario, each of the individuals may have a purchase rate of 3 purchases, which would lead to 30 unit sales of brand A. If the individual rule were changed to 7 purchases in the one-week period, total sales of "A1 Cola" would be 70. Figure 2 shows the application of this approach.

Figure 2: Applied to Brand Sales

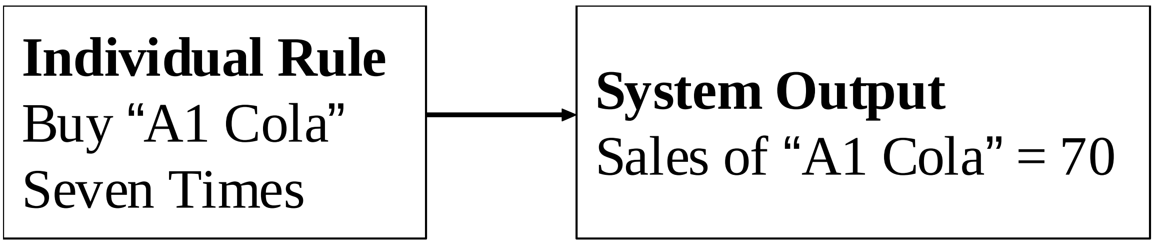

Representing the real world

The obvious shortcoming of this model is its simplicity. A marketing manager can only wish for their market to be as straightforward as that shown for A1 Cola. Some “Real life” considerations include:

Heterogeneity as not all shoppers will have a purchase rate that is the same value.

Purchase timing as an individual’s purchase rate is likely to be subject to periodic fluctuations.

For instance if a shopper has an underlying purchasing rate of 3.2 cans of cola per week, it is conceivable that, in some weeks, the shopper will purchase as few as none, while in other week the shopper may purchase substantially more than three.

In order to extend this approach to the real world, a set of agent rules must be defined that accommodate such considerations. A simple agent based model needs to do two things. It needs to accommodate the heterogeneity in buying rates across the population and be able to represent the short-term fluctuations that occur – even within the same shopper – over short periods.

Model Building: The NBD as a Set of Agent Rules

There is a model whose assumptions may provide a set of agent rules for a simple agent based model. The Negative Binomial Distribution (Ehrenberg 1959, Morrison 1969, Morrison and Schmittlein 1988) has been found to be a valid model for describing the distribution purchase rates for a single brand (or product category) across a population of shoppers. The distribution offers closed form estimates for a number of marketing metrics, such as penetration, average purchase rates and average purchase frequency, and the NBD model can be used to address questions about light and heavy buyers (Schmittlein, Cooper and Morrison 1993), helping to understand the Pareto effect (e.g. Habel, Driesener, Rungie and Jarvis 2003, Habel and Reibe 2004). The NBD offers much in demonstrating agent-based modelling as it is a rigorous and robust model that has been tested across many markets and categories for nearly 50 years. In ‘going under the hood’ to look at the mathematics that drive the model and transform the assumptions into consumer (agent) rules we are using well accepted science (at a market level) to derive individual consumer rules, which allows us to create a theoretical macro picture that can then be compared with the observed market.

The agent-based approach to complex marketing modelling is therefore confirmatory in this paper. Marketing scientists who are familiar with the NBD will recognise the 50 year old model, and this is not an accident. In this paper it is argued that an agent based version of the NBD provides a solid grounding for the testing of other agent level rules that may lie outside of the assumptions of the NBD. A more detailed discussion of these possibilities is addressed in a later section “future research – method applications”.

A Non-Deterministic Approach

The success of the Negative Binomial Distribution (NBD) model is despite (and possibly as a result of) its lack of covariates – its lack of "cause and effect" thinking. Each shopper in the population is considered to be an individual with their own reasons for purchasing a brand (or product category) at a given rate. These reasons are not considered in the traditional NBD studies. Rather, the purchase rate of each individual is treated as if it were a discrete and non-negative random variable. The probability density function that describes the distribution is a compound Gamma-Poisson, which is also known as a negative binomial (Johnson, Kotz and Balakrishnan 1997).

Rather than being a cause for concern, this lack of attention to each individual's reasoning is a highly attractive property of the NBD. Stochastic (probability based) models in have been widely used in marketing for many years (see Massy, Montgomery and Morrison 1970). By treating each individual's purchase as if it were random, progress can be made in obtaining accurate predictions of marketing metrics across a range of conditions.

Technically, the NBD is a Poisson distribution mixed by a gamma distribution. This technical description has substantial implications for the way in which the assumptions of the NBD can be used as a set of agent rules.

Agent Rule 1: The Poisson Distribution

The first agent rule attempts to address the period-to-period fluctuations around a shoppers average long run purchasing rate. A Poisson process (Schmittlein, Bemmaor et al. 1985) represents each shopper's observable purchases. Equation 1 gives the probability of a shopper buying k units if they have a known (long run) average purchasing frequency of λ.

$$f(k;\lambda) = \frac{\lambda^k e^{-\lambda}}{k!}$$ **Equation 1**

If the long run average purchasing rate (λ) is known, the probability of a shopper buying *k* units in a given time period can be calculated. For example, if the average number of purchases of "A1 Cola" an individual made in a week was 3.2 (i.e. λ=3.2), the Poisson probabilities that the shopper will buy at a given rate (*k*) are provided in Table 1.

Table 1: Poisson probabilities for an average purchase rate λ of

k	Probability of k
0	0.04
1	0.13
2	0.21
3	0.22
4	0.18
5	0.11
6	0.06
7	0.03
8	0.01
9	0.00

Consequently, it is possible to build an agent-based model that allows weekly fluctuations in an individual's purchase rate for a specified model average of 3.2 weekly purchases. Table 2 shows the results of a random number generation process that allocates such purchase rates according to the Poisson probabilities shown in table 1.

Table 2: Simulated purchases for a single shopper over 10 weeks

	Model average	Week 1	Week 2	Week 3	Week 4	Week 5	Week 6	Week 7	Week 8	Week 9	Week 10	Total	Observed Average
Shopper 1	**3.2**	3	0	3	4	3	5	2	3	4	5	32	3.2

This application of the Poisson assumption can be seen to be a simple agent based model that represents period-to-period fluctuations around an average long run purchasing rate (λ) of 3.2. Only ten weeks of behaviour are modelled, so it is coincidental that the observed average, which is shown in the last column, is equal to the "Model average", which is shown in the first column.

This model can also be applied to multiple shoppers. Table 3 shows a simulation of ten shoppers, each with identical long run purchasing rates (λ). The model has specified a homogeneous long run purchasing rate (specified at λ=3.2 for all shoppers). Over the relatively short 10 week time frame, the Poisson randomness generated some variability in the observed averages for each shopper, ranging from 2.2 to 3.6. This is consistent with the conceptualisation of the Poisson process as taking account of a population's "short term randomness".

Table 3: The Poisson Process Applied to a Set of Ten Homogeneous Shoppers

	Model average	Week 1	Week 2	Week 3	Week 4	Week 5	Week 6	Week 7	Week 8	Week 9	Week 10	Total	Observed Average
Shopper 1	**3.2**	3	0	3	4	3	5	2	3	4	5	32	3.2
Shopper 2	**3.2**	3	2	3	2	2	4	2	3	1	0	22	2.2
Shopper 3	**3.2**	3	6	4	3	3	3	2	3	5	4	36	3.6
Shopper 4	**3.2**	3	1	6	2	1	2	2	4	3	1	25	2.5
Shopper 5	**3.2**	2	1	3	5	2	5	1	4	2	2	27	2.7
Shopper 6	**3.2**	1	1	7	4	3	5	1	9	2	1	34	3.4
Shopper 7	**3.2**	2	3	7	0	1	3	5	2	3	3	29	2.9
Shopper 8	**3.2**	3	6	3	4	3	2	3	4	0	6	34	3.4
Shopper 9	**3.2**	4	1	5	1	2	1	4	4	4	5	31	3.1
Shopper 10	**3.2**	4	1	6	3	1	1	1	1	3	2	23	2.3

The observed averages would be expected to converge to 3.2 as the number of weeks modelled increased. The model is now beginning to represent reality a little more as it has multiple shoppers and allows for weekly fluctuations about an average (for each individual). However, a key shortcoming remains as it presumes the long run purchasing rate (λ) is homogeneous across the shopper population. This needs to be corrected as it would be a strange market indeed where every shopper bought, on average, 3.2 cans a week.

Agent Rule 2: The Gamma Distribution

The "average" column in table 3 would be expected to show some variation and a valid agent based model needs to specify a rule that represents such variability across shoppers. Under the NBD model, average (long run) purchasing rates are considered to be a random, continuous variable that is gamma distributed across a population of shoppers (Schmittlein, Bemmaor et al. 1985, Morrison and Schmittlein 1988). This assumption provides the second "agent rule" for a NBD agent based model. For shape parameter γ (gamma) and scale parameter β (beta) the probability of a shopper having a long run purchasing rate x is given by:

$$f(x;\gamma,\beta) = x^{\gamma-1}\frac{e^{-x/\beta}}{\beta^{\gamma}\Gamma(\gamma)} \qquad \textbf{Equation 2}$$

Where Γ is the gamma function (as shown in the appendix).

Under most conditions, the gamma distribution reflects the fact that many shoppers have a low – close to zero – long run average purchasing rate (Ehrenberg 1988). Table 4 represents both agent rules for our "A1 cola" example. The "average" column shows gamma heterogeneity (parameters γ=0.5 β=6.4) across shoppers and, for each shopper, there is weekly randomness around that average.

Table 4 is an agent based model of ten shoppers' purchasing over ten weeks, using the assumptions of the Negative Binomial Distribution as a pair of agent rules for the behaviour of each shopper.

Table 4: A model showing heterogeneity amongst shoppers and periodic randomness

	Model average	Week 1	Week 2	Week 3	Week 4	Week 5	Week 6	Week 7	Week 8	Week 9	Week 10	Total	Observed Average
Shopper 1	**5.2**	4	4	6	3	10	4	5	4	5	8	53	5.3
Shopper 2	**4.7**	6	9	2	6	10	2	8	2	4	4	53	5.3
Shopper 3	**3.5**	4	6	2	3	2	5	9	3	1	3	38	3.8
Shopper 4	**0.4**	0	2	1	3	0	0	0	0	0	0	6	0.6
Shopper 5	**1.1**	0	2	1	1	0	0	1	1	1	1	8	0.8
Shopper 6	**4.5**	3	3	2	5	4	8	3	7	5	4	44	4.4
Shopper 7	**1.5**	2	3	3	0	2	1	1	1	1	0	14	1.4
Shopper 8	**0.6**	1	0	0	1	1	0	0	0	0	1	4	0.4
Shopper 9	**0.7**	0	1	1	0	1	1	1	0	0	0	5	0.5
Shopper 10	**3.0**	2	4	1	4	3	4	4	7	4	4	37	3.7

The outputs of this model appear to offer a good representation of reality:

- The observed average number of purchases varies substantially across shoppers (Gamma Heterogeneity)
- The weekly purchase amounts for each shopper vary around that shopper's observed average (Poisson Purchase Timing)

It can be seen that shoppers 4, 8, and 9 rarely drink A1 cola, while shoppers 1, 2, or 6 are relatively heavy consumers of A1 cola. Further, as expected, each shopper does not purchase the same amount each week as there is variation.

It is prudent, at this point, to note that we do not know anything about the ten shoppers other than their average long run purchasing rate. Such is the nature of stochastic modelling. This variable is considered to be "as-if" random.

So far, a method has been described for representing reality according to a set of estimated model parameters. The parameters used for our simple agent based model were not based on prior estimation from observed data. The parameters used for the model shown in table 4 allowed for an easily understandable set of numbers.

It is unlikely that the cola market would look the way it was represented. It is more likely there would a majority of very low long run average purchasing rates (λ), as is common in many FMCG product categories (Ehrenberg 1988). The following section examines how well the model reflects reality, as well as examining its internal consistency.

Validation

The NBD agent based model can be applied to create a frequency distribution, as represented by a histogram. Consider week 1 from table 4. There were a total of three non-buyers, one shopper who bought once, two who bought twice and so on. Figure 3 shows these frequencies plotted as a histogram. The histogram tends towards a reverse j-curve shape, with high numbers of light (and zero) buyers and an ever-decreasing tail of heavier buyers, which is one of the recognised shapes of the NBD histogram, the other being a skewed normal distribution (Habel, Rungie, Lockshin and Spawton 2003).

Figure 3: A histogram of the simple agent based model

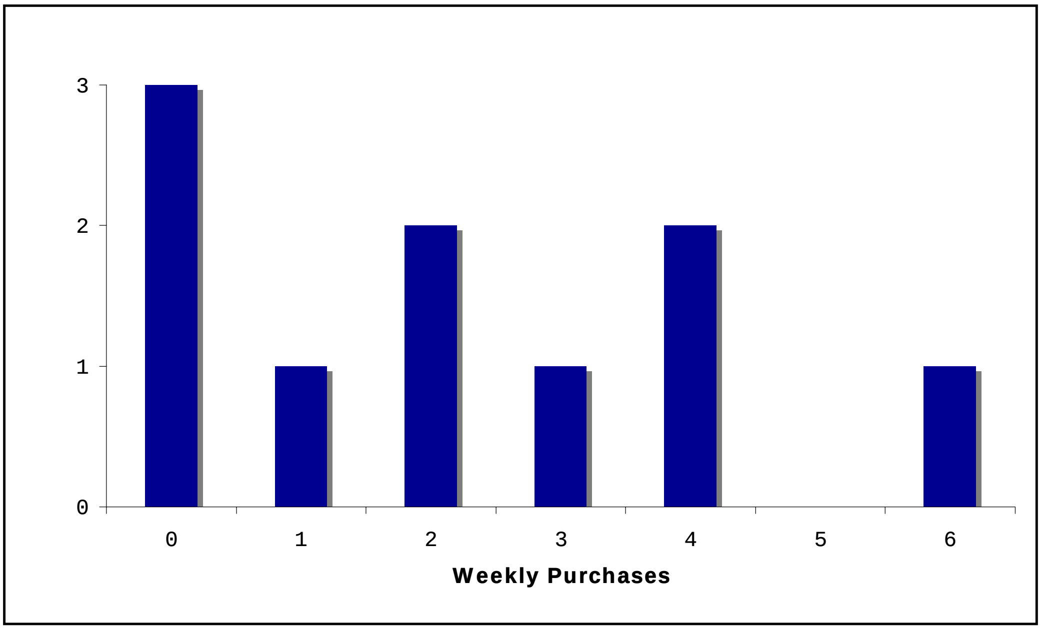

Two product categories that represent the ends of this continuum (chocolate bars and detergents) are now modelled. A market research company (GfK) collected the data in France during 1998 and 1999 as part of its MarketingScan service. The data were obtained from the combination of a questionnaire that records demographic and other variables and scanner data on all of the purchases from the seven major supermarket stores in Angers, a major city in France.

For both product categories the observed histogram is plotted first and the results of an agent-based simulation are plotted alongside. Figure 4 shows the histogram of observed purchase rates of chocolate bars. The pronounced reverse j-curve pattern is evident, with approximately 1100 non-buyers, 400 shoppers who bought once and an ever-decreasing tail of heavier buyers.

The NBD γ and β parameters (often referred to as K and A) were estimated using a maximum likelihood method (Eliason 1993). The likelihood estimates were found to be $\gamma = 0.37$ and $\beta = 18.4$. The NBD agent based model based on these parameters was run five times, and the results are shown in figure 5.

Figure 4: Chocolate bars, 12 months observed

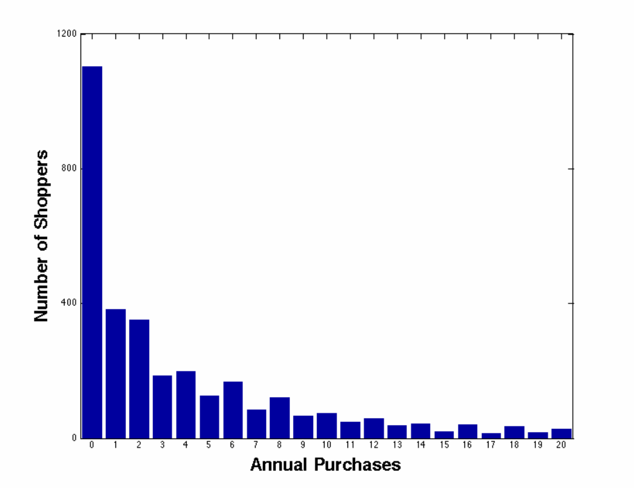

Figure 5: Five sets of estimates for chocolate bars

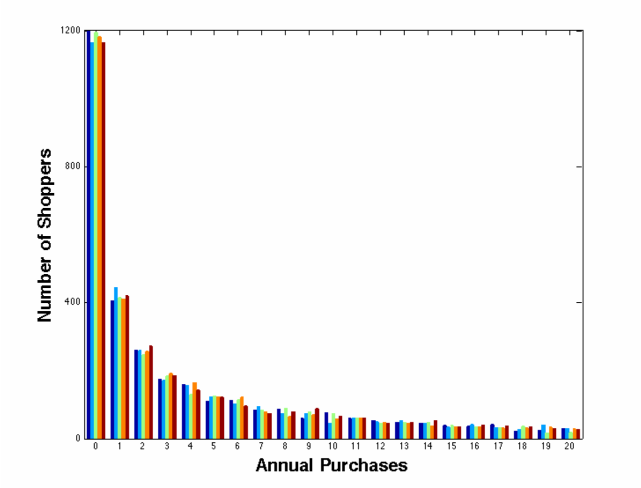

While each purchase class (0, 1, 2, 3 etc) in figure 5 may appear to be a solid bar, it is actually a set of five smaller bars (ranging from dark blue to red) that represent one run of the NBD agent based model. The first point is that the shapes are similar, suggesting a general fit to the NBD. Secondly, the top surface of each purchase class is not flat, suggesting Poisson randomness.

The NBD can also have a skewed normal distribution. This distribution can be seen in the detergent category, which is shown in figure 6. There is mode at 2, with a lower number of 1 only and zero buyers. This fits intuition about necessity product such as detergents over a 12 month period as most people buy such products a number of times, while fewer people buy once or not at all and there is an ever-reducing tail of heavier buyers. Figure 7 shows the results of the five simulations of the NBD agent based model as a grouped histogram. The general similarity in shape and Poisson randomness within each purchase class is again evident.

Figure 6: Detergents, 12 months observed

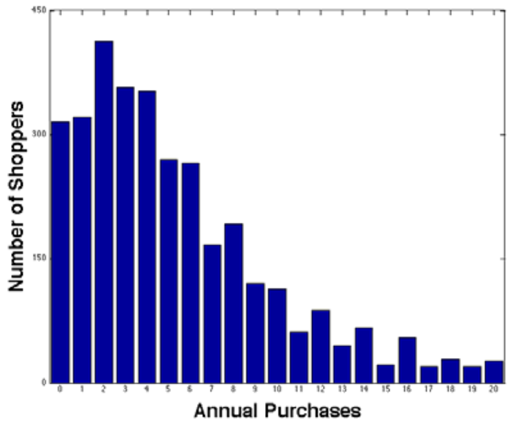

Figure 7: Five sets of estimates for detergents

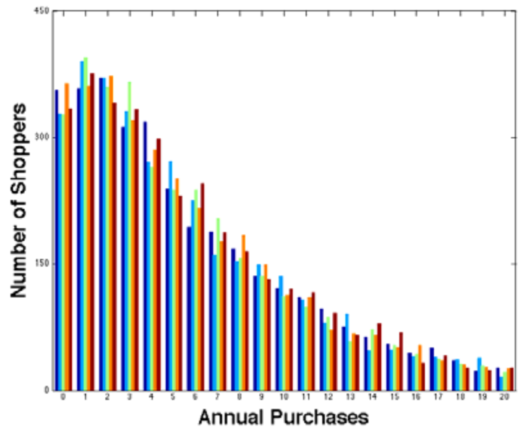

The NBD agent based model again provides consistent representations of the population of shoppers. The agent based model that generated Figure 6 and Figure 7 effectively simulated a dataset of five years' purchases of 3600 respondents. While each of the five runs of the model had some randomness, so too do real markets.

Future Research - Method applications

Given the NBD's use over nearly 50 years (Schmittlein, Bemmaor et al. 1985, Bass 1995, Uncles, Ehrenberg and Hammond 1995) we know the NBD-Dirichlet model is rigorous and robust. Prior to its use in this paper, the applications of the NBD have been many and varied:

- Simulated data can be used

- Out of sample predictions are routinely made - to shorter/longer time periods or providing conditional expectations for second period purchase rates (Goodhardt and Ehrenberg 1967, Morrison 1969)

- "What if" analyses are feasible, particularly within the benchmarking tradition of Andrew Ehrenberg (e.g. Ehrenberg 1988)

Part of the scope of this paper was to recast the familiar NBD model in the context of Agent Based Modelling and validate the ABM process. The following sections give an example of an application of the recast NBD and suggest an application that takes the ABM method beyond the two agent rules provided by the Negative Binomial Distribution.

An application within the NBD framework

Using this method – even with the foundation NBD agent rules – a researcher is able to provide two period estimates of purchasing for the same set of consumers (i.e. the same set of long run purchasing rates). This enables explicit predictions to be made for acquisition and defection of customers, much in the tradition of "Conditional Trend Analysis" (Goodhardt and Ehrenberg 1967, Morrison 1969).

Table 5 shows how such an application may be made, using a two week extract from the earlier example of an NBD agent based model.

Table 5: The simple agent based model, used to show acquisition and defection

	Model average	Week 1	Week 2	Customer Status
Shopper 1	5.2	4	4	
Shopper 2	4.7	6	9	
Shopper 3	3.5	4	6	
Shopper 4	0.4	0	2	Acquired
Shopper 5	1.1	0	2	Acquired
Shopper 6	4.5	3	3	
Shopper 7	1.5	2	3	
Shopper 8	0.6	1	0	Defected
Shopper 9	0.7	0	1	Acquired
Shopper 10	3.0	2	4	

Note that an explicit indication of acquisition and defection is provided by this simple two period analysis. Shoppers 4, 5 and 9 were predicted to have been acquired, with shopper 8 predicted to have defected. From simulations of this type, one may calculate metrics of expected customer acquisition, defection and churn in steady state markets all based on the solid foundation of the NBD.

Beyond the NBD – Testing a "Propensity Hypothesis"

In the above example the analyst can clearly see the model estimates of which customers they may lose and which they may gain, but it is nothing that could not be done with the existing NBD model. A benefit of the Agent Based Modelling approach, however, is that it is possible to explore conjectures that lay outside the strict Gamma-Poisson constraints of the NBD. For instance, a "propensity hypothesis" may be that a given marketing intervention – such as advertising or broader distribution – might increase the each shopper's underlying propensity of product purchase by 50%. This could be built into the foundation model by increasing each figure in the "model average" column of Table 5 by 50% and proceeding with agent rule 2; that of Poisson purchase timing. The outputs of such an agent based model could then be compared against observations of the market as a means of testing this "propensity hypothesis".

The validity of the above hypothesis may be unclear, but it is argued here that Agent Based Modelling provides a tool to investigate the hypothesis further. Furthermore, an investigation of the propensity hypothesis is not possible with the foundation NBD model. A multiplication of a set of gamma distributed variables by some constant does not produce a set of gamma distributed variables (Johnson, Kotz and Balakrishnan 1994), so the resultant agent based model has no counterpart in the NBD world. Unlike the acquisition/defection example, there are no closed form solutions available, and simulating data through the application of agent based modelling appears to be a sensible alternative.

Thus, an agent-based modelling approach may become a tool for developing falsifiable propositions in contexts where theoretical assertions (such as an expected 50% increase in purchase propensity) have largely remained untestable.

Agent rules may be developed to account for seasonality and promotional activity, allowing an analyst to compare before and after 'actuals' with theoretically derived results. It may be possible to move beyond stochastic agent rules and test deterministic approaches, allowing one to consider "cause and effect" questions.

Conclusion

The present paper has presented agent based modelling as a tool to generate falsifiable hypotheses from theory. As an example, assumptions of a well-established theory of marketing (NBD theory) were used to simulate individuals' purchase behaviour. The simulations were good representations of the behaviour of the population for two product categories, but the real benefit lies in where the method can be taken. Rather than be constrained by established probability density functions and closed form calculations, the suggested method makes it possible to decide on other agent rules (whether stochastic or deterministic) that can be used to simulate behaviour, create a new set of estimates and to compare these estimates with observed behaviour.

References

Bass FM (1995). Empirical generalizations and marketing science: A personal view. *Marketing Science, 14,* (No. 3, Part 2 of 2), G6-G18.

Bower J and Bunn D (2001). Experimental analysis of the efficiency of uniform-price versus discriminatory auctions in the England and Wales electricity market. *Journal of Economic Dynamics and Control*, *25*, 561-592.

Ehrenberg ASC (1959). The pattern of consumer purchases. *Applied Statistics, 8,* (1), 26-41.

Ehrenberg ASC (1988). *Repeat-Buying: Facts, Theory and Applications*. London, Oxford University Press.

Eliason SR (1993). *Maximum Likelihood Estimation - Logic and Practice*. California, Sage Publications, Inc.

Goodhardt GJ and Ehrenberg ASC (1967). Conditional Trend Analysis: A breakdown by initial purchasing level. *Journal of Marketing Research*, *4,* (May), 155-161.

Habel C, Driesener C, Rungie C and Jarvis W (2003). Pareto and Parameters: How Customer Concentration Varies with Nbd Parameters. *ANZMAC*, Adelaide.

Habel CA and Reibe E (2004). An investigation into Australian wine wholesaling using the Negative Binomial Distribution. *European Marketing Academy Conference*, Murcia, Spain.

Habel CA, Rungie CM, Lockshin L and Spawton AL (2003). The Pareto Effect (80:20 Rule) in consumption of liquor: A preliminary discussion. *Australian Wine Marketing Colloquium*, Adelaide.

Johnson NL, Kotz S and Balakrishnan N (1994). *Continuous Univariate Distributions*. New York, John Wiley and Sons.

Johnson NL, Kotz S and Balakrishnan N (1997). *Discrete Multivariate Distributions*. New York, John Wiley and Sons.

Marks RE (1992). Breeding hybrid strategies: Optimal behaviour for oligopolists. *Journal of Evolutionary Economics*, *2*, 17-38.

Massy WF, Montgomery DB and Morrison DG (1970). *Stochastic Models of Buying Behaviour*. Cambridge, The M.I.T. Press.

Morrison DG (1969). Conditional Trend Analysis: A model that allows for nonusers. *Journal of Marketing Research*, *VI*, (August), 342-6.

Morrison DG and Schmittlein DC (1988). Generalizing the NBD Model for customer purchases: What are the implications and is it worth the effort? *Journal of Business & Economic Statistics*, *6*, (2), 145-159.

Schieritz NaM, Peter M (2003). Modeling the forest or Modeling the trees: A comparison of system dynamics and agent-based simulation. *21st International Conference of the System Dynamics Society*, New York.

Schmittlein DC, Bemmaor AC and Morrison DG (1985). Why does the NBD Model work? Robustness in representing product purchases, brand purchases and imperfectly recorded purchases. *Marketing Science*, *4*, (No. 3, Summer), 255-266.

Schmittlein DC, Cooper LG and Morrison DG (1993). Truth in concentration in the land of (80/20) laws. *Marketing Science*, *12*, (2),

Simon HA (1978). Rationality as a process and product of thought. *American Economic Review*, *68*, 1-16.

Tay NSP and Linn SC (2001). Fuzzy Inductive Reasoning, expectation formation and the behaviour of security prices. *Journal of Economic Dynamics and Control*, *25*, 321-361.

Tesfatsion L (2003). Agent based computational economics. *Iowa State University Working Paper*, *1*,

Twomey P and Cadman R (2002). Agent-based modeling of customer behaviour in the telecom and media markets. *The Journal of Policy, Regulation and Strategy for Telecommunications*, *4*, (1), 56-63.

Uncles M, Ehrenberg A and Hammond K (1995). Patterns of buyer behavior: Regularities, models, and extensions. *Marketing Science*, *14*, (3, Part 2 of 2), G61-G70.

Cullen Habel is lecturer in Marketing and Steve Goodman is Senior Lecturer in Marketing at The University of Adelaide Business School, South Australia.

Appendix - The gamma function

A straightforward function that is embedded within many probability density functions is that of the gamma function. The notation Γ represents the function that is:

$$\Gamma(z) = \int_0^\infty t^{(z-1)} e^{(-t)} dt$$

Source: Milton Abramowitz and Irene A. Stegun, eds. Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables. New York: Dover, 1972.

This function is rarely calculated in an explicit fashion. All statistical software packages (and even Excel) carry some form of automatic gamma function calculator.