Modelling of Consumer Goods Markets: An Agent Based Computational Approach

Stephen E. Glavin
Unilever SEAC, Colworth Science Park, United Kingdom

Abhijit Sengupta
Essex Business School, University of Essex, United Kingdom

ABSTRACT

An agent based behavioral model incorporating utility based rational choice enhanced with psychological drivers is presented to study a typical consumer market. The psychological drivers incorporate purchase strategies of loyalty and change-of-pace, using agent specific memory of past purchases. Attribute specific preferences and prices drive the utility based choice function. Transactions data is used to calibrate and test the model. Results indicate that prediction accuracy at both macro and micro levels can be significantly improved with the incorporation of purchase strategies. Moreover, increased agent memory does not improve predictions in the model beyond a threshold, indicating that consumer memory of past shopping instances is finite and recent purchase history is more relevant to current decision making than the distant past. The article illustrates the use of agent based simulations to model changes or interventions in the market, such as new product introductions, for which no past history exists.

INTRODUCTION

Consumer behavior as a field of study is highly interdisciplinary in its approach, and that is evident in the amount of literature on this topic in multiple fields of study – whether in economics, psychology, sociology, computer science or even applied mathematics. Correspondingly, the traditional methods of analysis used by researchers in this field are numerous, and they range from quantitative (statistical and regression based) to qualitative (surveys, interviews, ethnographic studies etc.). However, over the last few years, new studies are being increasingly seen in the literature, which use modern computational techniques based on computer simulations, data mining, big data analysis etc., which mirror the changes and technological progress in societies and markets the world over. This chapter introduces one such method – agent based simulations, and links real world empirical data with models of behavior from multiple disciplines. Additionally, the chapter also provides an example of how such models can be used to "explore the future", with their ability to incorporate "what-if" scenario building techniques. This chapter is based on Sengupta and Glavin (2013) and Sengupta and Glavin (2010), and introduces the models and methods used in both, and extends them by illustrating how radical changes in the market (such as new product introductions) can be modeled robustly using computational methods.

Markets often exhibit noisy dynamics in the form of volatile movements in market shares (Jager, 2007). Frequent competitive interventions by manufacturers, such as introduction of new products, aggressive marketing policies such as multiple pricing and promotion strategies – is definitely one reason behind this widespread phenomenon (Ailawadi et. al., 2001; Blattenberg & Wisniewski, 1989). However, the presence of a wide variation in tastes and preferences amongst a reasonably large and demographically varied consumer population is also a key factor leading to the noisy character (Allenby & Rossi, 1998; Sengupta & Glavin, 2010; Sengupta & Glavin, 2012). Such markets do not lend themselves easily to traditional statistical and econometric analysis. Nor do markets where major interventions or events have occurred in the immediate past, which have moved these markets "out of equilibrium" (Reid & Brentani, 2004; Mathews, 2006), such as new product introductions, innovations etc. Additionally, the presence of potential non-linear interactions such as social networks, word-of-mouth influences etc. means that they may also exhibit a "complex" character – hence making traditional techniques further redundant. Not surprisingly, markets in general and consumer packaged goods (CPG) markets in particular, are increasingly being brought under the purview of "complex systems" analysis – whereby more modern "bottom up" methodologies such as agent based modelling are being used for analysis, inference and predictions (Gilbert et. al., 2007).

Systems which exhibit "emergent behavior" of some kind cannot be fully examined and analyzed by traditional "top-down" methodologies. Simulation based techniques – relying on agent based constructs – where constituents of the system (in this case, shoppers, firms etc.) are treated as individual modelling units (or agents) with the ability to follow independent rules of behavior and engagement have become increasingly popular and are widely advocated (Epstein & Axtell, 1996; Gilbert & Troitzsch, 2005; Tesfatsion, 2006). CPG markets have been extensively studied in the mainstream literature, but in spite of exhibiting many characteristics of a complex system, have only recently been brought within the purview of complex systems analysis (North et. al., 2009; Sengupta & Glavin, 2010; Rand & Rust, 2011). This paper builds a behavioral model of consumer choice, which is then incorporated within a multi-agent simulation framework to illustrate the accuracy and usefulness of such an approach in predicting market phenomena. It builds on earlier simpler models by the same authors, by incorporating crucial psychological factors into the choice model, modifying and extending the validation methodology and finally showing that both market and individual level predictions can be significantly improved by using enhanced choice models.

BACKGROUND

Our earlier work, specifically Sengupta and Glavin (2010) focused on developing a theoretical model of behavior, which took into account the heterogeneity in tastes within the consumer population and illustrated the link between this heterogeneity and the resultant volatility in overall market shares of brands and specific product characteristics. In order to carry out this analysis, the authors developed a rigorous methodology which focused on *out of sample* predictions of both macro level market share movements *and* micro level household level choices. Results showed a reasonably high degree of accuracy with which market share movements as well as choices made by individual households could be predicted out of sample. These results were compared with a benchmark model, where agent level heterogeneity was ignored, and it showed that the former far outscored the latter at both the macro and micro levels.

However, the behavioral model presented in Sengupta and Glavin (2010) was restricted at a simple level where agents made *rational* ordinal utility based choices – where the utility function depended solely on prices, product characteristics and individuals' preferences towards these characteristics. This restrictive utilitarian nature of the choice algorithm was intentional as the focus was establishing a validation methodology which is able to capture the volatile dynamics of the market and subsequently, illustrate the link with agent level heterogeneity. Sengupta and Glavin (2012) introduced an added layer of complexity into this framework by incorporating "purchase strategies". Purchase strategies referred to the psychological perspective in repeated shopping instances, where non-utilitarian aspects have an important effect on choices made by individuals. Such psychological drivers have been deemed to play an important role in determining shopping patterns of individuals (see Bettman (1986) and Maheswaran and Shavitt (2000) for some useful reviews). Janssen and Jager (2003) go further by using a set of experiments to demonstrate the role of psychological processes in the self organization of market. Sengupta and Glavin (2010) showed that this additional layer of complexity within the behavioral model actually enhances the predictability of the simulations for large sections of the population.

Psychological aspects of shopping are introduced through two crucial shopping strategies – “loyalty” and “change of pace” – potentially opposite in effect, but both key psychological drivers in the context of repeated purchase instances. The former refers to the explicit desire of consumers to stick to their historical purchases in future shopping events and the latter refers to their explicit desire to try out something different, neither being a direct consequence of their tastes and preferences. Oliver (1999) identifies consumer loyalty as a crucial component of decision making, quite separate from the “satisfaction” or utilitarian aspects of product use. Loyalty, in the model presented here, refers to inherent tendency of agents to choose the same product characteristic repeatedly, irrespective of utilitarian considerations. On the other hand, consumer choice models should incorporate a random component (Dellaert et. al., 1999), accounting for inconsistencies in choice from a behavioral perspective. Change of pace has been introduced as a strategy in the model to account for one such aspect – a typical consumer’s desire “to do something different” in any given shopping instance. These strategies complement the “rational” utilitarian approach developed in Sengupta and Glavin (2010), in the sense that these strategies affect the *choice set* of individual agents. Once the choice set is decided for an individual agent, the utility framework is then used to choose the appropriate item as before. The validation methodology is adapted to account for this additional complexity and associated parameters.

As part of these modifications, Sengupta and Glavin (2012) also introduced consumer memory within the modelling framework. Macdonald and Sharp (2000) have convincingly argued that in case of repeat purchases, brand “awareness” plays a strong role in determining brand choice. Going hand in hand with awareness, is the effect of retaining past experiences from prior shopping experiences in memory (Cowley and Mitchell, 2003). A large number of studies have been carried out with respect to short term memory in humans, and it has been well established in neuroscience literature that short term memory capacity is finite and is usually centered around 3 to 5 discrete memory chunks (see Cowan (2000) for an interesting and extensive study on this subject). The behavioral model presented here has been enhanced with the introduction of individual agent specific memory, which can affect the final choice of products by the agents in the simulation.

The model validation and verification methodology uses three distinct steps – *initialization*, *calibration*, *testing*, with each step being carried out on a distinct temporal partition of the data. While initialization is necessarily carried out at the level of agents only, the calibration and testing phases are done at both macro (market) and micro (agent) levels. We use Netlogo as the modelling framework used to carry out the simulation. The data used for the complete exercise is similar to the data used in the initial paper but much wider in scope. Three years worth of complete individual level transactions data from a single category is used for model building and validation exercise. The data is provided by a large online retailer in Europe, and the category of choice is fresh fruit juice.

The primary aim of the work is to illustrate the use of a model of consumer goods markets, combining both psychological and economic elements into a multi-agent system, with a high degree of predictive accuracy at macro and micro levels. The model is validated using individual level transactions data from a given category. The exercise is aimed to show how relatively simple behavioral models are able to capture the complex dynamic nature of consumer goods markets. Additionally, the modelling exercise illustrates the use of a rigorous validation methodology that is both flexible and accurate at all levels. The chapter will provide an illuminating guide on how to combine multiple theories, simulation and real life data seamlessly into a predictive modelling toolkit, which would be useful to both researchers and practitioners alike.

We also introduce an additional useful feature of computational models for consumer goods markets – i.e. the ability to use them for predicting the impact of making changes or intervening in the markets from the point of view of the products. Some changes or interventions, such as price changes, promotions etc. can be readily modeled using traditional statistical techniques. These interventions can be relatively easily modeled primarily due to the presence of historical data. However, some type of interventions in the market, such as introduction of a new product or product variants, is more difficult to model, due to the lack of related historical data. Prior research carried out on strategic aspects of new product launch has not been too effective in providing practical solutions to the problem of missing data from the past (Calantone & Montoya-Weiss, 1993; Chiu et.al., 2006). In fact, as pointed out by Schneider and Hall (2011), most new product launches in consumer goods markets fail in terms of reaching their sales targets in the first year of launch. This chapter shows that the use of agent based computational modelling techniques can actually mitigate this problem to a great degree, and introduces a method of testing the market with “hypothetical products” within the simulation framework. This method provides a robust testing ground for new product development, and provides indications to managers of the consequences of introducing a new product into the existing portfolio available to consumers.

The next two sections present the agent level behavioral model and the validation method in a general mathematical format (which can be adapted or modified for separate studies). The particular data set used to illustrate and validate this methodology is presented next, followed by detailed description of how the theoretical models are adapted within the simulation setup. Finally, the results of the analysis are discussed along with concluding remarks. Since the theoretical models are intensive in mathematical notation, a brief summary of key variables are provided in Table 5 at the end of this chapter for the reader’s convenience.

MODEL

In this section, we develop the underlying theoretical model which drives consumer agent behavior in the simulations. This model combines two aspects to the behavior – a rational utility based choice and an intrinsic psychological drive based choice behavior. The two are described in detail below.

Utility Based Choice

The model framework presented here is adapted directly from Sengupta and Glavin (2012). The key assumptions in the model are the following.

Assumption 1 *Consumers act rationally and are able to rank the available alternatives in a consistent manner given their preference.*

Assumption 2 *All products and product characteristics remain unchanged during the given time period under consideration.*

Assumption 3 *Consumers’ tastes and preferences remain unchanged in the given period.*

Consider an industry with K distinct products and a consumer base of size I. Each product is endowed with a set of N attributes, which makes it unique for a consumer. In order to define the preferences of consumers in such a framework, we borrow from traditional discrete choice theory in which a product is consumed, not for its own sake, but for the set of attributes it embodies (Lancaster, 1966; DeSoete et. al., 1986). Hence, the following assumption has been made in the behavioral model as well.

Assumption 4 *Each consumer ranks alternatives based on a subjective ordinal utility measure, which is a function of the product specific price and attributes, as well as consumer specific preferences.*

Hence, given a vector of characteristics or attributes in a product, we are able to place it in a *discrete N*-dimensional *characteristic space*, where each dimension refers to a single attribute. This vector is then called the *address* of this product. Each consumer's preference is defined using a complementary *ideal point*, which is a vector of characteristics that he would ideally like to see in a product. The closer this ideal point is to the actual mix of characteristics of a commodity, the higher the subjective utility of the consumer from purchasing it.

For any $k \in K$, let $X_k = (x_k^1, x_k^2, \ldots, x_k^N)$ be the address for product *k*. For any consumer $i \in I$, consider $\lambda_i = (\lambda_i^1, \lambda_i^2, \ldots, \lambda_i^N)$ as *i*'s ideal point. Let P_k be the *net* price of the product *k*. We define $D_k^i = \sum_{j=1}^{N} |x_k^j - \lambda_i^j|$ as the deviation of product *k* from consumer *i*'s ideal point. Consumer *i*'s utility from product *k* is characterized as,

$$U_i(k) = \omega^i d_k^i + \left(1 - \omega^i\right) p_k \tag{1}$$

where, $0 \leq \omega^i \leq 1$. Identity (1) refers to the parameterized utility function where, ω^i and $1 - \omega^i$ are the normalized weights placed by individual *i* on d_k^i and p_k respectively, where $d_k^i = -\frac{D_k^i}{\max_{j \in K} D_j^i}$ and $p_k = -\frac{P_k}{\max_{j \in K} P_j}$. Note that P_k is the *per unit* net price of product *k* and not the listed net price[1].

Psychology Based Choice

Two consumer specific psychological drivers are incorporated into the model – loyalty (*L*) and change of pace (*CoP*) – both being applied at the attribute level. Both strategies help define the *consideration set* $C_i(t)$, or the set of alternatives of products that a consumer agent will finally choose from at time *t*. Strategy *L* refers to an individual consumer's loyalty or bias towards a given attribute present in his purchase history, while strategy *CoP* refers to the tendency of *ignoring* a given attribute present in his purchase history (out of boredom or motivation to try out something different). As mentioned above, memory plays an important role in the way these strategies affect consumer choice. Each agent is assigned a globally constant fixed memory length *M*, which is basically the number of his own past attribute purchases he can recall.

An agent can choose either strategy *L* or *CoP* at any given time step *t.* As the consumer makes purchases over time, their attribute specific memories fill up, recording the purchases they make in the last *M* shopping instances. Let $mem_i^n(t)$ be a vector of purchases of attribute $n \in N$ by consumer agent *i* in time *t*. Let Y_n be the set of values possible within attribute *n*. Hence,

$$mem_i^n(t) = \{y_i^n(-1),\ y_i^n(-2), \ldots, y_i^n(-M)\},$$

where $y_i^n(-s) \in Y_n$ is *i*'s purchase of attribute $n \in N$ in the s^{th} past shopping instance where, $s \leq M$. As an agent's number of shopping instances exceeds *M*, older purchases are forgotten as new ones are added.

In a given shopping instance in time *t*, an agent using strategy *L* on attribute $n \in N$, will consider only those products whose n^{th} attribute match the most frequently occurring or *modal* value of $mem_i^n(t)$. That is,

$$C_i(t) = \left\{k \in K \colon x_k^n = mode\left(mem_i^n(t)\right)\right\}. \tag{2a}$$

Hence the implementation of strategy *L* implies that the consumer agent purchases will consider those values of attribute *n*, which he has bought most frequently in his available memory.

The *CoP* strategy, on the other hand, will cause consumer agents to ignore products matching any brand or flavor which make up a given proportion of their memory. Let $L_i(t) \leq M$ be the filled *length* of consumer *i*'s memory vector and $0 < \alpha_i^n \leq 1$ be an agent specific constant associated with the *CoP* strategy for every attribute *n*. If the *CoP* strategy is used, the consideration set is constructed as,

$$C_i(t) = \{k \in K : x_k^n = y \; \& \; freq(y, mem_i^n(t)) \leq \alpha_i^n L_i(t)\}, \tag{2b}$$

where, $freq(y, mem_i^n(t))$ refers to the frequency of occurrence (within agent *i*'s memory vector) of property $y \in Y_n$ in attribute *n*. Hence, the implementation of strategy *CoP* implies that the consumer agent will consider those values of attribute *n* which he has bought less frequently in his available memory.

As mentioned above, the strategies *L* and *CoP* are implemented at the level of attributes and not products and the choice of either strategy is independent across all attributes. The parameter α_i^n can also be interpreted as the *probability* of choosing strategy *CoP* by agent *i* for attribute *n*. Conversely, $(1 - \alpha_i^n)$ can be considered as the probability of choosing strategy *L* for attribute *n*. At every shopping instance, an agent will choose either *CoP* or *L* for every attribute *n*, based on the probabilities specified in the *N*-dimensional vector $\{\alpha_i^1, \alpha_i^2, \ldots, \alpha_i^N\}$.

Choice of Product

The final choice of product at a given shopping instance in time *t*, uses both utility and psychology based choice rules as specified in (1), (2a) and (2b). Note that the cardinality of set $C_i(t)$ can be zero or positive. Consumer *i*'s final choice of the product $k_i^*(t)$ in time step *t* is made using the following rule.

$$k_i^*(t) = \begin{cases} arg \max_k U_i(k), & if \; |C_i(t)| > 0 \\ random \; choice, & if \; |C_i(t)| = 0 \end{cases} \tag{3}$$

This implies that once the consideration set $C_i(t)$ is defined and it contains at least one product, the utility function defined in (1) is then used to make the final choice. However, if $C_i(t)$ is empty, the consumer falls back on the original set of products and chooses one at random. The complete process of a typical agent choosing a product at any given time is represented in Figure 1.

VALIDATION METHODOLOGY

The validation strategy specified in Sengupta and Glavin (2010) uses three temporal partitions of the data. This work uses a similar overall validation strategy, but suitably modified to reflect the enhancements in the underlying behavioral model. Partition 1 is used to *initialize* the product addresses, agent specific ideal points and individual propensities for loyalty and change of pace, partition 2 is used to *calibrate* agent specific utility weights ω^i and partition 3 is used to test out of sample predictions using the initialized and calibrated model. Model validation is a crucial step in any simulation based model, but is especially relevant when the simulation is aimed primarily at accurate predictions (Rand & Rust, 2011; Mosler, 2006; Epstein, 2008; Thompson & Derr, 2009). According to Fagiolo et. al. (2005) and Windrum et. al. (2006), validated models should possess a satisfactory range of accuracy matching the real world with the results of the simulation. Moreover, validation strategies should aim at accuracy at multiple levels – at the macro levels capturing the emergent properties of the system under study, as well as at micro levels, capturing the behavior of agents and their interactions.

Initialization

The characteristics space is defined as the subset $[0,1]^3 \in \mathbb{R}^3$ (as shown below, the data requires three dimensions in the characteristic space), so that the maximum and minimum values attached to any one dimension are 1 and 0 respectively. For each characteristic, the unique categories are assigned a value based on their total sales volume in the first year (2007) of the dataset, normalized by the maximum within that dimension. For instance, within the dimension representing brand, the one with the highest total sales volume is assigned the value 1 while the one with the lowest is assigned 0. All other brands were placed equi-spaced within $(0,1)$, with each brand's position proportional to the relative sales volume. All ties are resolved randomly.

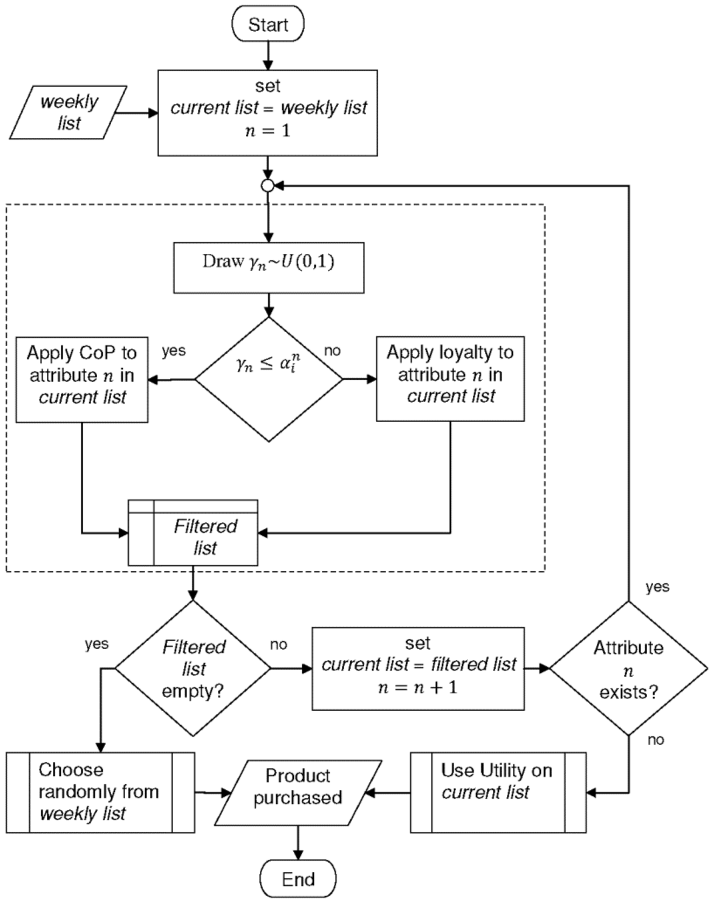

Figure 1: Consumer Agent Decision Making Process

Next, a proxy for the ideal point of each agent is calculated using the transactions history from the first year of the specific household which the agent represents. For a given characteristic dimension and for a given household, we calculate the weighted average of all the categories purchased, with the purchase frequency used as the weight. This is repeated for each characteristic dimension. A three dimensional vector, representing its ideal point, in the characteristics space can thus be constructed for each household, based on its transactions history. After ideal points are set they remain static throughout the rest of the analysis.

Finally we set the agents' individual propensities for loyalty and change of pace and seed the agents' memories for use in the calibration stage. As mentioned previously, the constant α_i^n, can be thought to represent the probability that agent i will apply a CoP strategy associated with attribute n ∈ N. This is done by considering the products purchased by an individual household on each of their purchasing occasions within the initialization partition of the data. We essentially set α_i^n as the proportion of times agent i switched, in terms of attribute n, between each purchasing occasion . To see this more clearly, consider an agent who made purchases on T occasions and thus consider the sequence of T sets, $(\mathrm{A}_1, \mathrm{A}_2, \ldots, \mathrm{A}_\mathrm{T})$ each of which contain the products they purchased on those separate occasions. We define $\mathrm{A}_\mathrm{t}^\mathrm{n}$ to be the set of values of attribute n for products in set A_t. The switching constant for attribute n is therefore defined as

$$\alpha^n = \frac{1}{T-1}\sum_{t=1}^{T-1}\Phi_t^n \tag{4}$$

with

$$\Phi_t^n = \begin{cases} 1, & \text{if } A_{t+1}^n \cap A_t^n = \emptyset \\ 0, & \text{otherwise} \end{cases}$$

where subscripts i have been dropped for convenience. Carrying out this procedure for each attribute, we may obtain an N-dimensional vector for each agent.

Also at this stage we need to seed the agents' attribute specific memories for use in the calibration stage. This is done by considering the M final purchasing occasions within the first year of data. Note that some households may not have as many as M previous purchasing occasions and so their seeded memories will not be full. We call the filled number of memory slots of agent i their memory length and denote it $L_i(t)$, as stated previously. Each entry of memory is filled with a randomly selected value of the attribute under consideration from the agent's corresponding set of purchased attributes so that

$$\begin{gathered} y^n(-1) \in A_T^n \\ y^n(-2) \in A_{T-1}^n \\ \vdots \\ y^n(-L(t)) \in A_{T-L(t)+1}^n \end{gathered}$$

where subscript i has been dropped for convenience and $L(t) \leq M$.

Calibration

The calibration stage is used to set the utility function weightings, ω^i, for each of the agents within the simulation and is carried out at both a micro and macro level separately. For the macro level, we aim to fit the evolution of market shares of product groups to the actual data, while at the micro level, the corresponding aim is to fit household specific choice of SKU and product characteristics. Because each level requires a different fitness metric over which agent specific parameters are calibrated, we carry out the calibration once for each metric. We use binary matching of simulated versus real take-up of SKUs to calibrate at market share (macro) level and use the city-block metric to calibrate take-up at the household (micro) level.

The simulation is run over the second temporal partition of the dataset (year 2008) with a time-step of one week. The parameter space is discretized into 25 equi-spaced points, $0 = \omega^0 < \omega^1 < \cdots < \omega^{24} = 1$, where we have dropped subscripts for convenience, and the simulation is run 25 times with each run corresponding to one value of ω. Two sets of optimum parameter subsets are constructed per agent i, Ω_b^i using binary matching and Ω_c^i using the city-block metric, defined in detail below.

Binary matching

The binary matching calibration is used to optimize each agent's utility function for optimum SKU choice. For each agent and each parameter value, the simulated and actual data are compared. In each purchase week a binary matching score is found in the following way. Let K be the set of all products and let the simulated choice by agent i at week t be $S_t^i \in K$. That is, for any given week t,

$$S_t^i = \arg\max_{k \in K} \{U_i(k)\}$$

is the product chosen by agent i after a comparison of utility from all available products (SKUs). The binary matching score for agent i at week t is then defined as

$$b_t^i(S_t^i) = \frac{1}{Q_t^i} \sum_{k=1}^{K} q_{t,k}^i \delta(k, S_t^i) \tag{5}$$

where

$$\delta(k, S_t^i) = \begin{cases} 1, & \text{if } k = S_t^i \\ 0, & \text{otherwise} \end{cases}$$

In the above definition, $q_{t,k}^i$ is the quantity of product k purchased by household i in week t and Q_t^i is the total number of purchases by the household in that week, so that $Q_t^i = \sum_{k=1}^{K} q_{t,k}^i$. For agent i, and for each parameter value ω^p, $p = 0,1,2, \ldots ,24$ (the subscript i is suppressed), the binary matching score is averaged over all of the purchase weeks to obtain an overall score,

$$B^i(\omega^p) = \frac{1}{|T^i|} \sum_{t \in T^i} b_t^i\left(S_t^i(\omega^p)\right) \tag{6}$$

where $T^i$is the set of weeks where household i made a purchase and $\left|T^i\right|$ is the cardinality of set T^i. The set of best parameter values for agent i, Ω_b^i is then constructed as,

$$\Omega_b^i = \left\{\omega^p \middle| B^i(\omega^p) = \max_p \left(B^i(\omega^p)\right), p = 0,1, \ldots 24\right\} \tag{7}$$

These identities spell out the macro level calibration strategy. For any agent i, every match with the corresponding household's purchase in the real data is given a score of 1 for the particular parameter value ω^p and 0 otherwise. The total score per week is summed up and normalized for the total number of purchases made that week (to account for multiple purchase instances within the week). Weekly scores are averaged for all weeks per parameter value and the parameter value(s) with the highest score added to the set Ω_b^i.

City-block metric

We use a city-block metric in order to calibrate the utility function of agents for optimum choice of product attributes. This calibration is based on matching agent based simulated choice and actual choice at the level of characteristics in the N-dimensional characteristic space, where the dimensions are orthogonal. This implies that in every individual transaction in the data set under consideration, the calibration algorithm should award a match in *any* dimension and penalize a *mis*-match in any other. The simplest algorithm to achieve this is the city-block metric, loosely based on the city-block *distance* measure (or the L1 norm), commonly used in geometry (see Fichet (1987) for details on general theoretical results on the L1 norm).

Once again, let $S_t^i \in K$ be the simulated choice made by agent i in week t. Following previous notation let x_k^n be the n^{th} characteristic. A city-block metric assigns a score of 0 per characteristic where it matches with the real data and a 1 for a mismatch. The characteristic matching score for agent i in the week t is then defined as,

$$c_t^i(S_t^i) = \sum_{n=1}^{N} \min_{\substack{k \in K \\ q_{t,k}^i \neq 0}} \Delta^n(k, S_t^i) \tag{8}$$

where

$$\Delta^n(k, S_t^i) = \begin{cases} 0, & \text{if } x_k^n = x_{S_t^i}^n \\ 1, & \text{otherwise} \end{cases}$$

Here $0 \leq c_t^i \leq N$. Since $N = 3$ for our data, a score of 0 would mean that each of the characteristics of the simulated product S_t^i matched with a corresponding characteristic in household i's actual product purchases in week t; a score of 1 would mean that any two out of three were matched (i.e. 1 mismatch); 2 that only one could be matched (2 mismatches) and 3 that none of s_t^i's characteristics matched with any of the actual purchases.

The set of best parameter values for agent i, Ω_c^i is found in the following manner. First, we define a binary variable $\Psi_t^i(\omega^p)$, which is used to judge whether a characteristic matching score is "good" enough, which in turn is used to determine the optimal ω^p.

Definition 1 *The parameter value ω^p is considered optimal for agent i in week t, if it is less than the mean city-block distance over all parameter values in week t, i.e.*

$$\Psi_t^i(\omega^p) = \begin{cases} 1, & \text{if } c_t^i\left(S_t^i(\omega^p)\right) \leq \dfrac{1}{25}\displaystyle\sum_{p=0}^{24} c_t^i\left(S_t^i(\omega^p)\right) \\ 0, & \text{otherwise} \end{cases}$$

As before, an overall score $C^i(\omega^p)$ is obtained for each parameter setting by averaging over all purchase weeks:

$$C^i(\omega^p) = \frac{1}{|T^i|} \sum_{t \in T^i} \Psi_t^i(\omega^p)$$

Finally, the set of optimal parameter values for agent i is defined as,

$$\Omega_c^i = \left\{ \omega^p \middle| C^i(\omega^p) = \max_p \left(C^i(\omega^p) \right), p = 0,1,\dots,24 \right\} \tag{9}$$

Once the Ω_b^i and Ω_c^i sets have been identified for each agent/household, we then need to seed the agent memories using the actual purchases from the calibration set in order to proceed with the final stage in the analysis, that of testing our model with out of sample data. Note that after the calibration and initialization stages it is likely that there will be enough transaction data for each household so that their memories will be full when entering the testing phase.

Testing

The model testing exercise is carried out on out of sample transactions data covering the final year. For each agent i, we now use both sets Ω_b^i and Ω_c^i in multiple runs. The use of either set Ω_b^i or Ω_c^i is dependent on the type of validation being carried out, i.e. whether macro or micro level respectively. By definition, $\Omega_b^i \subseteq \Omega_c^i$ and we could have just used the former for prediction. However, given that the objective of micro-validation is to parameterize the utility function on the basis of individual preferences, ignoring elements in $\Omega_c^i - \Omega_b^i$ is essentially loss of agent specific information.

It is very likely that the cardinality of these sets is greater than one and so for each agent one parameter value $\omega \in \Omega_{b,c}^i$ is selected at random. As before, at each purchase week the agent makes a choice from that week's product choice set using the parameterized utility function. The variability that is introduced through random choice of suitable parameter values necessitates the need for a Monte-Carlo type analysis and so the simulation is run 100 times for each agent. We select the modal value, or the product that is purchased the maximum number of times in a week among the 100 runs, as the week's predicted choice. Quantity of choice, $\bar{q}$ is not a part of the behavioral model, however it is calculated based on an average liquid volume (pack size) per transaction of the household vol^i. This is obtained from the initialization dataset. We define

$$\bar{q} = \min\left\{\chi \in \mathbb{Z} \middle| \chi \geq \frac{vol^i}{x_S^1}\right\}$$

Where x_S^1 is the pack size of product S (the simulated purchase of i), implying that the purchase for that week satisfies the household's average liquid volume purchase per transaction. In the characteristic matching (micro-level) validation, purchase quantities need not be considered.

Benchmarking

In a previous paper (Sengupta & Glavin, 2010), benchmarking was carried out against a random choice model in which all agents choose products purely on a probability distribution taken from sales volumes in the initialization dataset. The behavioral model was seen to outperform the probability model. So in this analysis, the strategy based models are compared to the strategy free behavioral model of the earlier paper. In order to see the effects of memory length the simulation is run for memory lengths of M = 4 and M = 8.

DATA

The data used to build and test the models presented here comes from the consumer packaged goods sector. We use 3 years (2007 - 2009) of individual transactions data from one category within a supermarket – with the category we focus on being fruit juices[2]. Each transaction features the following information on the purchase – a unique anonymized shopper ID, unique product ID (Stock Keeping Unit or SKU[3]), the price paid, the discount if any and the date of purchase. The product ID can be mapped to a product database, which provides the following product characteristics as well for each SKU – brand, flavor and pack size. Weeks 115 and 116 (in Year 3) could not be used as abnormal values indicated recording error.

We consider only those households within the model who made at least 24 purchases with the final two years of the data (calibration and prediction data-sets), we also only consider those with at least one purchase in the first year. Products which do not appear in the final two years were removed, leaving a total of 171. The product database provided information on the brand, flavor and pack size of all SKUs listed, but some SKUs had additional attributes associated with them as well. But given that only those attributes, which were available for all products could be used in the models, some SKUs had to be merged into product groups – essentially, unique products which agents in the simulation could distinguish between, given the three attributes of brand, flavor and pack size. In all, there were 57 unique products (or product groups), encompassing 25 unique brands, 11 unique flavors and a continuum of pack sizes.

SIMULATION SETUP

In this Netlogo based simulation each agent represents a unique household in the data. There are a total of 9379 agents which we break down into four groups based on the number of purchasing occasions within the initialization dataset. Boundaries between sets are based on quartiles. In the simulation Group 1 has the least number of purchasing occasions (less than 3) and consists of 1271 agents, Group 2 has three or more, but less than six purchasing occasions and consists of 2692 agents, Group 3 has 2361 agents whose number of purchasing occasions were less than ten but greater than or equal to six, finally Group 4 consists of 3055 agents with greater than nine purchasing occasions. The idea behind grouping agents in this way was to more clearly show the importance of the initialization and strategies. The differences in number of purchasing occasions within the initialization dataset potentially affects two aspects of the strategy simulation. Firstly, those with few purchasing occasions will have less accurate values for their switching constants α^n (in order to account a little for this, those who have only one purchasing occasion have their switching constants set to a small value $\epsilon = 0.005$). The second problem involves the initial memories of agents. Those with few purchases are likely to have their memories filled to far less than capacity, as such, the benefits of the strategies is less likely to be seen early on. It is hoped that there will be a clear benefit to using strategies seen by comparison of results between non-strategy and strategy simulations for groups with higher numbers.

Each agent has a parameter set $(\lambda_i, \omega^i, \alpha_i)$, where $\lambda_i = (\lambda_i^1, \lambda_i^2, \lambda_i^3)$ and $\alpha_i = \left(\alpha_i^1, \alpha_i^2, \alpha_i^3\right)$. We have that $\lambda_i, \alpha_i \in [0,1]^3$ and $\omega^i \in [0,1]$. The parameters λ_i and α_i are initialized using the first partition while the parameter ω^i is calibrated in the second partition. Along with these parameters each agent also has attribute specific memories. For the calibration stage these are pre-allocated by use of the initialization dataset and for the testing phase they are pre-allocated using the calibration dataset. While the simulation is running, these memories fill with the values of attributes for the new purchases made by the agent with past purchases being “forgotten”. Two separate memory lengths – $M = 4$ and $M = 8$ are used in the simulations, where 4 is the “magical number” according to Cowan (2000).

At every time step in the simulation, agents choose one product from a subset of the 57 total products. We do not make the assumption that all products were available at all weeks covering the simulation but use the transactions table to determine the available choices.

Definition 2 *A product* k ∈ K *is considered to be within the available set in any week t, if there is at least one transaction in the real data involving k in week t.*

Moreover, the only information that we have from the transaction table is which product was bought at what price/discount in a certain week by a customer, but not what the prices and discounts were of all alternatives that were available to him. Given that this information is required to evaluate the relative utility of all products by the representative agent, we once again compute this from the transactions table.

Definition 3 *For any product k, the price (discount) listed in time step t is the average price (discount) corresponding to all the transactions involving k in that particular week t.*

An agent, when facing the set of available choices, simply looks up the corresponding prices and discounts from a table for that particular week. If a household made a purchase within a given week, then the agent representing the household is provided with a choice set of products which includes those that were purchased in reality plus all the “available” alternatives. The agent will then use the strategies applied for each attribute in turn as sequential filters on the list of available products following the method outlined in equations (2a) and (2b).

Given the final list of products the agent evaluates the prices, discounts, product characteristics and selects the one product which maximizes its utility. In the circumstance that the final list is empty following the agent’s use of the strategies we say that the agent’s selection method has failed and she then resorts to random selection from the original list. The following rule is also implemented in the simulation:

• If the household made no purchases that week then no purchase is made by the agent.

Once the purchase decision has been made by an agent in a given week, the simulation then progresses by one time-step to the next week. The purchase made by the agent is recorded at each time-step with the relevant attribute being recorded to the attribute memories.

RESULTS

Three main experimental settings were used – two incorporating the two strategies L and *CoP* and the third without incorporating the strategies (i.e. based on rational utility based choice) which acts as the benchmark against which the other two are compared. We report a whole year of out-of-sample predictions at both macro (market) and micro (household) levels with two separate memory lengths of agents – $M = 4$, $M = 8$ and when no strategy is present. This implies predictions over week 105 to week 156, but excluding weeks 115 and 116, where anomalous numbers for all products indicate recording errors in the data. The benchmark for comparison are the results obtained using the simpler model presented in Sengupta and Glavin (2010) (where consumers do not use any strategy, but only the utility function) and which by itself provides noticeable accuracy and improvement over probabilistic models. At the macro level, we predict the simulated market shares of each brand and each flavor and compare them against the actual values. At the micro level, we test the model's ability to predict weekly household level choice of SKU, brand and flavor. All results are provided for each population groups 1, 2, 3 and 4 individually as well as combined for the whole population.

Macro Level Predictions

The macro level predictions are measured at two levels – (i) the average *root mean square error* (RMS) in the weekly market share prediction against the actual of each brand and flavor, averaged over 50 weeks; (ii) the week to week correlation between the predicted and actual market shares of each brand and flavor. The RMS per brand/flavor x is computed as,

$$RMS(x) = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(\frac{(\hat{S}_t^x - \hat{A}_t^x)}{\hat{A}_t^x}\right)^2}$$

where $\hat{S}_t^x$ and $\hat{A}_t^x$ are the simulated and actual market shares respectively. The correlation between the two is measured using the Pearson's correlation coefficient (cc). Both of these together provide a comparative indication of how well the models perform in predicting the dynamics of market shares at the brand and flavor levels.

The main results in the macro level tests are presented in Tables 1 and 2 for all three candidate models. Each table represents the average RMS and average cc from the market share predictions over weeks 105 – 156, within each model given a population group, across *all* brands (Table 1) and flavors (Table 2). Note that the combined population results are presented as well (ALL), apart from those of each group (GP1,…, GP4). Two results emerge from the tables.

First of all, both models ($M = 4$ and $M = 8$) perform significantly better than the benchmark (No Strategy) in all groups and in the combined population – more significantly in terms of magnitude than in direction. Overall, introduction of strategies do make a positive impact on the predictability of the models (except in the cases of groups 1 and 2 for flavor prediction). In the current scenario, GP4 is the group with the best historical data and hence, clearly outperforms the benchmark by a good margin. Note that the correlations in models with strategies do not always outperform the benchmark.

Table 1: *Average RMS and cc in market share predictions across all brands.*

Population	M = 4		M = 8		No Strategy	
	RMS	cc	RMS	cc	RMS	cc
ALL	0.74	0.34	0.71	0.36	1.30	0.3
GP1	1.54	0.32	1.50	0.31	1.59	0.28
GP2	1.72	0.39	1.41	0.38	2.22	0.23
GP3	1.13	0.27	1.14	0.29	1.16	0.18
GP4	0.90	0.29	0.95	0.30	1.95	0.32

Table 2: *Average RMS and cc in market share predictions across all flavors.*

Population	M = 4		M = 8		No Strategy	
	RMS	cc	RMS	cc	RMS	cc
ALL	0.51	0.18	0.50	0.18	0.7	0.3
GP1	0.87	0.25	0.88	0.24	0.83	0.22
GP2	0.81	0.26	0.86	0.24	0.79	0.1
GP3	0.71	0.15	0.79	0.14	0.72	0.24
GP4	0.45	0.20	0.47	0.24	0.73	0.28

Secondly, we see an interesting result in terms of memory length and predictability. A longer memory length is *not* synonymous with better predictability in all cases. Only when the number of past observations is limited (such as in groups 1 and 2), it seems to be the case that a longer memory length will provide more accurate out of sample predictions. In fact, if sufficiently large purchase history is present in terms of number of observations, using a lower memory length to define the strategies seem to be better in terms of predicting the future. This is indeed an interesting and significant result, which seems to suggest that households either do not retain memories of past purchases apart from a very few past instances, or even if they do, discount their effects severely when making current choices. This effect disappears if past history is sparse, in which case, the memory of past purchases seem to assist the decision making process, and hence provides more accurate predictions.

Micro Level Predictions

In order to validate a predictive agent based simulation, one needs to examine the micro level predictions as well. In order to measure the accuracy of the models, we use the following information from the SKU level predictions per household per week – the SKU id, the corresponding brand, flavor and pack size. Following Sengupta and Glavin (2010 and 2012), both the independent dimension specific prediction results as well as the joint prediction results across subsets of dimensions are reported. Given a specific dimension out of the three (brand, flavor or pack size), the former measures the proportion of times the model correctly predicts a household's choice along that dimension. The latter measures a joint score per household per transaction as defined in (8) – i.e. the number of times the model correctly respectively predicts 3, 2, 1 and 0 characteristics – irrespective of what those characteristics are. Score 0 indicates 3 characteristics were predicted correctly, score 1 indicates 2 characteristics were predicted correctly, score 2 indicates that only 1 characteristic out of 3 was predicted correctly and finally, a score of 3 indicates that the predictions failed for all characteristics in that particular transaction.

Table 3 shows the overall results of household level predictions for the whole population. Part 1 shows that models with strategies perform significantly better than the one without, when predicting individual household's choice of SKUs per week. Memory lengths of 4 and 8 improve percentage prediction from 19.92 to 27.81 and 28.56, respectively. However, it also shows that increased memory lengths from 4 to 8 do not bring about any significant increase in predictability. The distribution of prediction accuracy also varies across households and transactions, and Part 2 of Table 3 indicates the same. This distribution describes the break-up of households in terms of accuracy of predictions jointly across subsets of characteristic dimensions. Each column indicates the distribution of households under a particular score. For instance, 15.63% of the households had all products characteristics from *all* their transactions correctly predicted with $M = 4$, whereas the equivalent value for the model with $M = 8$ was 15.99%. In contrast, the model without strategies could achieve only 12.07% in this regard. Finally, Part 3 of Table 3 summarizes the accuracy of household level predictions within each dimension individually. Once again, the models with the strategies outscore the benchmark significantly, whereas, insignificant difference can be seen by increasing the memory length from 4 to 8.

Note that the proportion of times when all characteristics are predicted correctly in a transaction is higher than the SKU level prediction accuracy. This is a result of the fact that in a number of cases, two or more SKUs may have the same values for all three dimensions, yet remain distinct SKUs – due to factors (or characteristic dimensions) unaccounted for. In such cases, all three characteristics may have been correctly predicted, but the SKUs did not match – leading to lower prediction accuracy.

Table 3: *Household Level Predictions*

1. SKU Predictions	Strategy, M = 4				Strategy, M = 8				No Strategy			
Exact matches	23542				24175				16860			
No. of transactions	84655				84655				84655			
SKU matches (%)	27.81				28.56				19.92			
2. *Characteristic Predictions - Joint across dimensions*												
Percentage Accuracy (%)	*Sc. 0*	*Sc. 1*	*Sc. 2*	*Sc. 3*	*Sc. 0*	*Sc. 1*	*Sc. 2*	*Sc. 3*	*Sc. 0*	*Sc. 1*	*Sc. 2*	*Sc. 3*
100	15.63	2.97	2.26	4.91	15.99	3.14	2.40	5.42	12.07	3.15	5.31	8.27
75-100	10.09	3.46	3.35	6.16	10.40	3.52	3.01	5.93	7.11	3.20	5.76	7.96
50-75	13.51	7.99	10.20	13.65	13.63	8.06	10.13	13.10	9.75	7.32	12.52	14.15
25-75	12.89	12.80	15.85	17.55	13.00	12.15	15.02	17.32	9.02	11.62	15.28	16.41
0-25	47.88	72.78	68.34	57.74	46.98	73.13	69.44	58.23	62.06	74.71	61.13	53.21
Mean across households	38.22	17.07	18.55	26.16	38.79	17.06	18.08	26.08	28.27	16.11	24.78	30.85
SD across households	38.41	26.37	25.69	30.53	38.60	26.73	25.70	30.87	37.25	26.36	30.88	33.84
3. *Characteristic Predictions - Independent per dimension*												
	Brand	*Flavor*	*Pack Size*		*Brand*	*Flavor*	*Pack Size*		*Brand*	*Flavor*	*Pack Size*	
Mean across households	52.45	64.04	50.86		53.05	64.57	50.93		40.45	54.74	46.60	
SD across households	37.98	34.89	37.99		38.11	34.92	38.15		39.76	38.45	38.94	

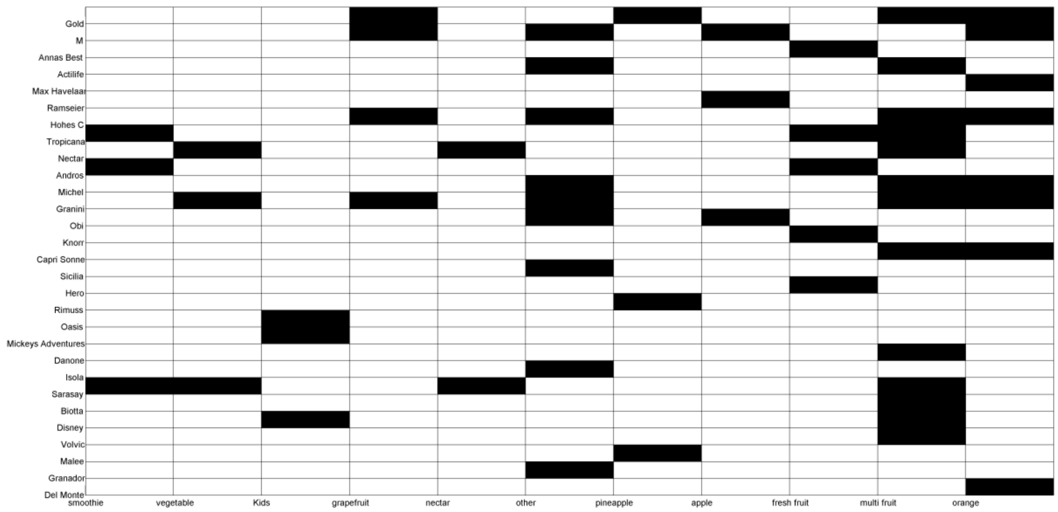

Figure 2: Gaps in the market. Brand flavor combinations in existing SKUs. Black – at least one SKU present with this combination; White – no existing SKU with this combination.

Table 4: Incremental market shares of Gold and Apple for new product introductions

New product variant	Brand GOLD	Flavor APPLE
Small, low price	18.3	40.6
Small, medium price	11.1	18.0
Small, high price	13.3	14.7
Large, low price	6.4	2.7
Large, medium price	2.5	0.9
Large, high price	1.3	1.1

New Product Introduction

As mentioned earlier, modelling changes or interventions in the market, which has limited or no prior history is difficult and so far has remained a challenge in the marketing literature. However, a multi-agent simulation framework coupled with a behavioral model as presented above, provides a way to test the consequences of interventions to a far greater degree than is possible using traditional analytical techniques. Here we present an illustration of a new product introduction, where an existing brand introduces a new flavor in multiple pack sizes. The new flavor that is being introduced was previously absent from the brand’s own portfolio, but was being supplied by other brands in the market. While this is not a completely new product, i.e. the product characteristics existed in the market prior to its launch individually in other products; they never existed together in a single SKU. From the perspective of a brand manager, the aim is to identify the consequences of this introduction on all products in his brand, under different price settings for the product.

Figure 2 shows a map of existing products, as brand flavor combinations. We selected the combination of brand Gold, flavor Apple as the new introduction into this market. The nearest competitor of Gold, in terms of overall market share is M, which already has Apple flavored products in the market. As can be seen, there are only three other brands which have Apple flavored products, and as such, is a good opportunity for Gold. Our aim is to see what the consequence of this introduction is going to be on Gold's market share and the Apple juice market overall. The new product offerings are launched (introduced into the simulation) at the beginning of testing phase of the simulation (week 105 or the beginning of year 3). All agent specific parameters are set at the calibrated levels, and the characteristics of all existing products are kept fixed as before.

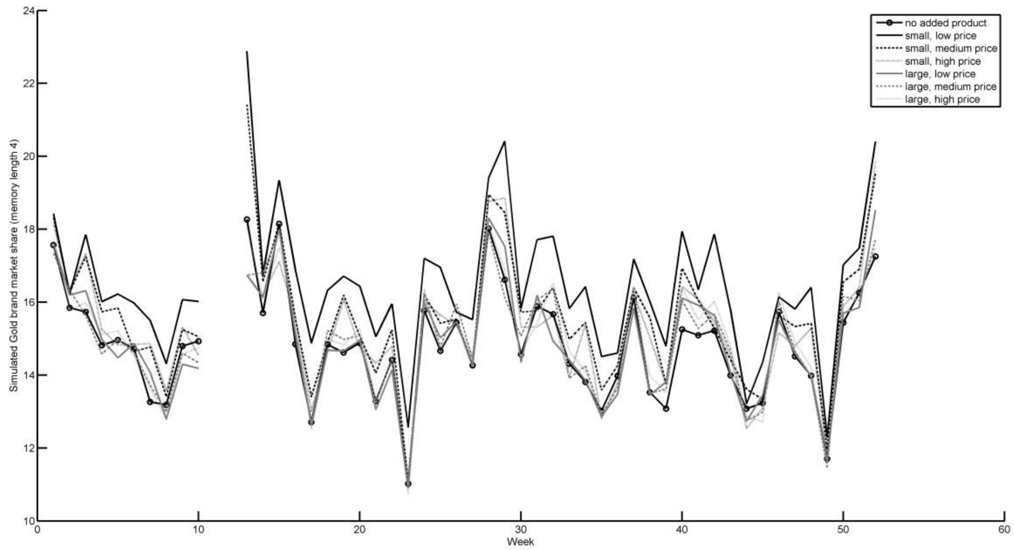

Figure 3: Market share dynamics of brand Gold in Yr 3 – actual, simulated (original), simulated (with new products). Excludes weeks 11 and 12.

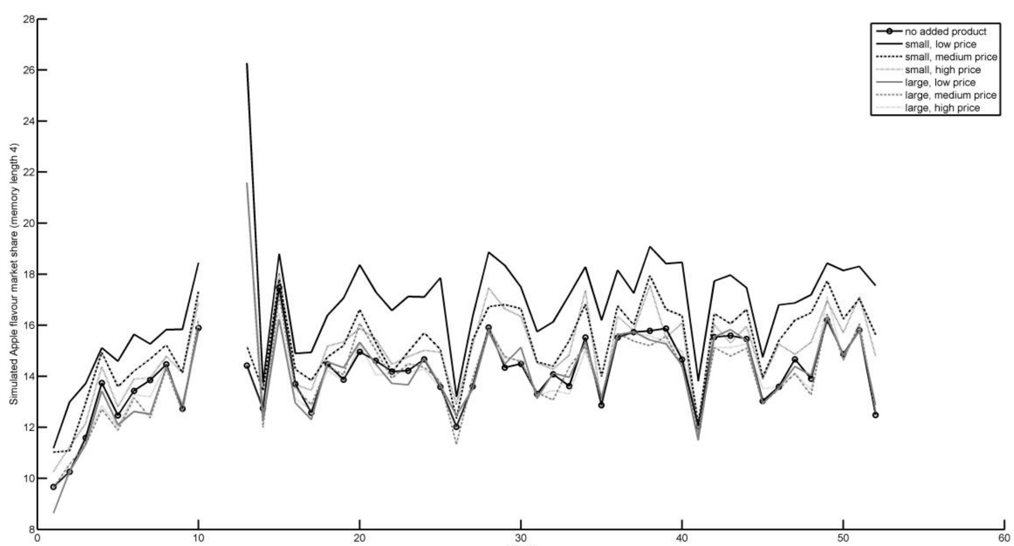

Figure 4: Market share dynamics of flavor Apple in Yr 3 – actual, simulated (original), simulated (with new products). Excludes weeks 115 and 116.

The new product is assumed to be available at all relevant weeks in Year 3, and the selected price for the product is fixed for the entire time. Three price points are tested for each product, high, medium and low. The high price coincides with the maximum per unit price charged by any of Gold's SKUs within weeks 105 to 156, the low price is the minimum of the same, and the medium is the average of all prices charged by Gold in the same period. Additionally, the new product is tested in two separate pack sizes – small and large. The small pack size coincides with the modal pack size less than the mean of all pack sizes in the market, while the large pack size coincides with the mode of all those larger than the mean. Note that, in theory, any product characteristic along any dimension can be incorporated into the new product(s) and potentially any price can be charged for it (including varying the price in different weeks). Also, we have tested each price pack size combination independently in the simulations, but more than one combination can be introduced simultaneously as well.

Figures 3 and 4 plot results of the introduction of the new product on the weekly brand Gold and flavor Apple market share movements – for each price pack size combination separately. The simulations indicate that the new product introductions will all lead to an increase in week by week market shares of both the brand and the flavor (conditional on everything else in the market being the same as it was in whole of Year 3). Table 4 provides the predicted percentage changes to overall market share of both for the whole of Year 3, and we can see that all are predicted to increase market share, with the small pack sized low priced product providing the largest increment (18% for the brand and 40% for the flavor approximately). These figures are based on consumer preferences and past purchase patterns of all product attributes and price in the first two years. If needed, one can even factor in potential competitor reactions to the new introduction and/or promotion strategies etc. and rerun the simulations to see what the impact of these might be on one's own products and competitor products.

CONCLUSION

Consumer goods markets have been established as complex in character, and hence inherently difficult to predict, even when good reliable micro level historical data is present. The complexity in such markets, in the form of heterogeneous tastes and preferences, inter-relationships, as well as frequent interventions by the firms themselves are difficult to handle analytically using traditional statistical and econometric techniques. Sengupta and Glavin (2010) provided the first glimpse of tackling the issue of predictability using a computational agent based technique based on a simple rational choice model, and showed that good predictions of a volatile market may be possible. This chapter, along with Sengupta and Glavin (2012) extends the above model, to incorporate crucial behavioral and psychological drivers within the consumer choice model and shows that predictions can indeed be improved significantly. This chapter also provides insight into how such a method can be used to model interventions in the market without directly relevant historical, using straightforward simulations.

We show that particular purchase strategies – *loyalty* and *change of pace* – can be particularly influential in consumer decision making in a repeat purchase framework. A simple utilitarian model (based on preferences over attributes and prices) is enhanced with the above strategies, with the strategies affecting the *choice set* of individual agents according to agent specific parameters. The important concept of consumer memory is used to define these strategies. We use a three year transactions database from a supermarket to test and validate this model, with the original rational choice model as the benchmark of comparison. Predictions show a marked improvement against the benchmark at both macro and micro levels. Additionally, we also show that though consumer memory plays an important role in improving the model, the effect is limited by a finite memory capacity of individual agents. Higher memory capacity attributed to agents seems to distort choices at the micro-level, supporting earlier findings from neuroscience literature. Just attributing additional memory to agents is not enough to improve the explanatory power of the model. Households do use the recent past purchases as a guide to driving current behavior, but discount the effect of purchases made in the distant past.

Consumer decision making and its effect on markets and market dynamics is a vast area of study, where many potential avenues of future research is possible – especially, under the purview of agent based computational economics. Effect of word-of-mouth through social networks in consumer markets, and its corresponding impact through psychological processes is one such avenue where research has only just started. Additionally, effects of various media campaigns – especially through modern electronic channels, is an additional area worth investigating. Especially difficult to predict currently, impact of new product launches, shifting technologies etc. can also be incorporated into the analysis.

All in all, the model presented above along with the validation methodology, is an especially flexible framework to study complex consumer goods markets. In fact, the model can be transferred to other markets with modifications to the behavioral model, which are not too difficult. As we have shown above, volatile consumer markets become increasingly tractable in terms of predictions, when analyzed using well designed computational models.

Table 5: Description of key variables used in the models and simulation setup

Variable	Description
Utility Model	
$U_i(k)$	Utility of product k to agent i
ω^i	Relative weight on distance from ideal point by for agent i
D_k^i and d_k^i	Absolute and Relative distance of product k from i's ideal point
λ_i^k	Ideal level of characteristic j for agent i
P_k and p_k	Absolute and relative *net* prices of product *k* respectively
Psychology Model	
M	Maximum memory length or memory capacity of agents
$mem_i^n(t)$	Memory vector for attribute n time t for agent i. Stores the choices made in the past M purchase instances
$y_i^n(-s)$	Purchase made of attribute n in the s^{th} past shopping instance by agent i
α_i^n	The probability of choosing the Change of Pace strategy on attribute n by agent i
$C_i(t)$	Consideration set of products in time t for agent i
Validation – initialization	
A_t^n	The set of values of attribute n purchased by an agent in time t
A_t	Set of products bought by an agent in time t
Φ_t^n	Binary variable representing whether a switch has been made in attribute n at time $t+1$ compared to t
Validation – calibration	
S_t^i	Simulated choice in time t by agent i

$b_t^i(S_t^i)$	Weighted average binary matching score per agent in time t, where the weights are quantities bought by the real household
$B^i(\omega^p)$	Overall binary matching score, given ω^p , for agent i
ω^p	p^{th} candidate value of ω within range of 0 to 1
Ω_b^i	Optimal parameter set for binary matching
$c_t^i(S_t^i)$	Characteristic matching score per agent in time t
$\Psi_t^i(\omega^p)$	Binary function indicating optimal value of ω^p in week t
$C^i(\omega^p)$	Overall characteristic matching score, given ω^p for agent i

REFERENCES

Ailawadi, K., Lehmann D., & Neslin, S. (2001). Market response to a major policy change in the marketing mix: Learning from procter and gamble's value pricing strategy. *Journal of Marketing 65*(January), 44–61.

Allenby, G. M., & Rossi, P. E. (1998). Marketing models of consumer heterogeneity. *Journal of Econometrics 89*(1-2), 57–78.

Bettman, J. R. (1986). Consumer psychology. *Annual Review of Psychology 37*, 257–289.

Blattenberg, R., & Wisniewski, K. J. (1989). Price induced patterns of competition. *Marketing Science 8*(4), 291–309.

Calantone, R.G., & Montoya-Weiss, M.M. (1993). Product launch and follow on. In: Souder, W.M.E., Sherman, J.D. (Eds.), *Managing New Technology Development*, 217–248.

Chiu Y-C, Chen B., Shyu, J.Z., & Tseng, G-H. (2006). An evaluation model of new product launch strategy. *Technovation* 26, 1244-1256.

Cowan, N. (2000). The magical number is 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences 24*, 87–185.

Cowley, E., & Mitchell, A. A. (2003, December). The moderating effect of product knowledge on the learning and organization of product information. *Journal of Consumer Research 30*(3), 443–454.

Dellaert, B. G., Brazell, J. D., & Louviere J. J. (1999). The effect of attribute variation on consumer choice consistency. *Marketing Letters 10*(2), 139–147.

Epstein, J., & Axtell, R. (1996). *Growing Artificial Societies – Social Science from the Bottom UP*. Cambridge, MA: MIT Press.

Epstein, J. M. (2008). Why model? *Journal of Artificial Societies and Social Simulation 11*(412).

Fagiolo, G., Moneta A., & Windrum, P. (2005). Empirical validation of agent-based models: A critical survey. In *International Workshop on Agent-Based Modeling and Economic Policy Design*, University of Bielefeld, Germany.

Fichet, B. (1987). The role played by L1 in data analysis. In Y. Dodge (Ed.) *Statistical Data Analysis based on the L1-norm and related methods*. Amsterdam:North-Holland.

Garcia, R., Rummel P., & Hauser, J. (2007). Validating agent-based marketing models through conjoint analysis. *Journal of Business Research 60,* 848–857.

Gilbert, N., Jager W., Deffuant, G., & Adjali, I. (2007). Complexities in markets: Introduction to the special issue. *Journal of Business Research 60*(8), 813–815.

Gilbert, N., & Troitzsch, K. G. (2005). *Simulation for the Social Scientist* (Second ed.). Maidenhead, UK: Open University Press.

Jager, W. (2007). The four p's in social simulation, a perspective on how marketing could benefit from the use of social simulation. *Journal of Business Research 60*, 868–875.

Janssen, M. A., & Jager, W. (2003). Simulating market dynamics: Interactions between consumer psychology and social networks. *Artificial Life 9*, 343–356.

Macdonald, E. K., & Sharp, B. M. (2000). Brand awareness effects on consumer decision making for a common, repeat purchase product: A replication. *Journal of Business Research 48*(1), 5–15.

Maheswaran, D., & Shavitt, S. (2000). Issues and new directions in global consumer psychology. *Journal of Consumer Psychology 9*(2), 59–66.

Mathews, J.A. (2006). Strategizing, disequilibrium & profit. Stanford University Press, California.

Mosler, H.-J. (2006). Better be convincing or better be stylish? A theory based multi-agent simulation to explain minority influence in groups via arguments or via peripheral cues. *Journal of Artificial Societies and Social Simulation 9*(3), 16–32.

Moss, S. (2008). Alternative approaches to the empirical validation of agent-based models. *Journal of Artificial Societies and Social Simulation 11*(1), 5.

North, M. J., Macal, C. M., Aubin, J. S., Thimmapuram, P., Bragen, M., Hahn, J.,…, Hampton, D. (2010). Mulitscale agent-based consumer market modelling. *Complexity 15*(5), 37–47.

Oliver, R. L. (1999). Whence consumer loyalty? *Journal of Marketing 63*, 33–44.

Rand, W., & Rust, R.T. (2011). Agent based modeling in marketing: Guidelines for rigor. *International Journal for Research in Marketing 28*(3), 181-193.

Reid, S.E., & Brentani, U.D. (2004). The fuzzy front end of new product development for discontinuous innovations: A theoretical model. *Journal of Product Innovation Management 21(3)*, 170-184.

Sallans, B., Pfister A., & Dorffner, G. (2003). Simulation and validation of an integrated markets model. *Journal of Artificial Societies and Social Simulation 6*(4).

Schneider, J., & Hall, J. (2011). Why most product launches fail. *Harvard Business Review*, April.

Sengupta, A., & Glavin, S.E., (2010). Volatility in a consumer packaged goods market: A simulation based study. *Advances in Complex Systems 13*(4), 579–605.

Sengupta, A., & Glavin, S.E., (2012). Predicting volatile consumer markets using multi-agent methods: Theory and validation. In Alexandrova-Kabadjova, B., Martinez-Jaramillo, S., Garcia-Almanza, A.L., & Tsang, E. (Eds.), *Simulation in Computational Finance and Economics: Tools and Emerging Application* (pp. 339-358). IGI Global.

Tesfatsion, L. (2006). Agent based computational economics: A constructive approach to economic theory. In Tesfatsion, L., & Judd, K. L. (Eds.), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*. Amsterdam: North-Holland.

Thompson, N., & Derr, P. (2009). Contra Epstein, good explanations predict. *Journal of Artificial Societies and Social Simulation 12*(9).

Windrum, P., Fagiolo, G., & Moneta, A. (2007). Empirical validation of agent-based models: Alternatives and prospects. *Journal of Artificial Societies and Social Simulation 10*(2), 8.

KEY TERMS AND DEFINITIONS

Complex System: A system which is usually composed of large number of possibly heterogeneous interacting agents, which are seen to exhibit *emergent* behavior. Emergence implies that system level behavior (macro level) cannot be inferred from observation of individual level behavior of its constituents (micro level). This absence of explicit links between the micro and macro levels makes complex systems especially difficult to analyze using traditional statistical and analytical techniques to study the dynamics of behavior. One typically requires the use of bottom up simulation based methods to study such systems. Complex systems are ubiquitous – markets, societies, social networks, the Internet, weather, ecosystems, are just a few examples.

Agent Based Modelling: Refers to a computer simulation methodology which is commonly used for modeling complex social systems. The system is constructed *bottom up*, that is, individual constituent units of the system are programmed as independent autonomous units (agents), following simple to complicated behavioral rules, and with the ability to interact with their environment and other agents in the system. The models are seeded with an initial set of values for its parameters, and the simulation is set to run for a given length of time. Behavior of individual agents as well as the system as a whole is recorded and analyzed once the simulation has run its course.

Model Validation: Validation in the context of agent based modeling refers to the process of ensuring the behavior exhibited by a simulated system matches the real target system which it is trying to mimic. This involves a systematic exploration and calibration of model parameters, ensuring robustness of the model, eliminating logical errors in the model code and finally, ensuring that the model replicates the target system with a high degree of accuracy both at the macro level and at the micro level.

Agent Memory: In the context of agent based models, this refers to the far back any individual agent in the model is able to recall past actions, behavior, parameter values etc., which may be relevant for its current state. The memory could refer to its own variables or variables of other agents or global and environmental variables, and this depends on the context and nature of the model being considered.

Initialization: The values of parameters that any simulation experiment starts with. These parameters could be agent specific or global.

Calibration: The process of exploring the right range of values for parameters in the model, which ensures realistic and accurate behavior at the micro and macro levels in any computer simulation. Involves re-running simulations with varying parameter values to see which provide the best match to real data.

Testing: Involves using the calibrated simulation to make predictions on a particular time period (or cross section) from which the data has not been used for calibrations. These predictions are then compared with the corresponding data to ensure the goodness of fit of simulated data with real data. Note that the data being used here is not the same as the one used for calibration. The testing exercise provides further validity of the models and ensures a degree of robustness for practical use.

Endnotes

[1] The utility based choice model is kept simple at this stage, where utility is a function of product characteristics and price only. A more complicated utility function is possible where other factors are seen to influence choice directly through the utility function. This would imply multiple weighting parameters – depending on the number of additional factors which affect utility. For instance, if one more additional factor is present apart from characteristics and price, we would require 2 parameters $\{ \omega_1^i, \omega_2^i \}$ in all to characterize the utility function.

[2] The fruit juice category is used in the analysis for two main reasons. Firstly, because of the nature of the category, the data is inherently more reliable as far as product descriptions and attributes are concerned.

Hence it is comparatively easier to mine the transactions data for product characteristics and use these for defining preferences, ideal points etc. Secondly, the earlier study of Sengupta and Glavin (2010) used the same category to establish the methodology. It would seem reasonable to use a similar (if not the same) data set to compare the results established here against the ones earlier, in order to establish the validity of psychological drives in consumer choice models.

[3] SKU or Stock Keeping Unit is a number code used to identify each unique product or item for sale in supermarkets or other stores.