

VOLATILITY IN THE CONSUMER PACKAGED GOODS INDUSTRY — A SIMULATION BASED STUDY

ABHIJIT SENGUPTA

*Unilever Research & Development,
Sharnbrook, MK44 1LQ, United Kingdom
abhijit.sengupta@unilever.com*

STEPHEN E. GLAVIN

*Department of Mathematics, UCL,
London, United Kingdom
glavin@math.ucl.ac.uk*

Received 11 March 2010

Revised 4 June 2010

The volatility in a CPG market is modeled using a bottom-up simulation approach and validated against disaggregated supermarket transactions data. The simulation uses independent agents, each agent representing unique households in the data. A simple behavioral model incorporates household preferences for product attributes and prices. Our validation strategy tests the model predictions at both macro and micro levels and benchmarks the performance in each against a random choice model. The model significantly outperforms the benchmark at both levels. At the macro level, choices made by heterogeneous agents accurately captures the volatility in market shares over time. This accuracy at the macro level is driven by the accuracy of predictions at the micro household level SKU and attribute choice.

Keywords: Complex system; validation; market dynamics; social simulation.

1. Introduction

Software based simulation techniques have become a popular tool in the analysis of populations and societies in recent years. Social scientists have repeatedly pointed out that human societies, groups and organizations usually comprise of multiple heterogeneous agents who often interact with each other as well as with the environment they are placed in. Such “complex systems” do not lend themselves to analysis using traditional top-down analytical/quantitative techniques. These systems typically exhibit *emergent behavior* of some kind, and simulations seem to be one of the best tools to use for study [9, 16, 27]. Consequently, bottom-up approaches such as agent-based modeling/agent-based computational economics (ABM/ACE) have increasingly been used in analyzing such systems. This paper uses such a bottom-up approach in building a simple but effective behavioral model of a typical consumer

packaged goods (CPG) market, whose volatile and heterogenous nature makes it very suitable for a simulation-based study.

Consumer markets are usually characterized by noisy dynamics and instability [18]. Such volatility may arise due to multiple factors, for instance, large variations in tastes/preferences [2] and/or intense competitive market interventions by firms such as discounts and pricing campaigns, competitive packaging, advertising, etc. [1, 5]. Moreover, potential lateral effects such as word-of-mouth can become important as well in consumers' decision making and as a result lead to nonlinearities in the system. Hence CPG markets possess many key characteristics of a complex system and consequently, lend themselves readily to agent-based modeling techniques [15].

This paper takes the first step in building and validating a model of CPG markets using a combination of theory, simulation and to real life data. We model the fresh fruit juices category and the choices made therein by a set of supermarket consumers facing a large variety of competing products. We focus on consumers' tastes and preferences for various attributes inherent in a product, as well as their reactions towards pricing and promotions in order to build a behavioral model of choice. This behavioral model is subsequently used as the template for running simulations where we attempt to replicate the choices made by real life consumers in an online retail supermarket. This is a multi-agent model analyzing a dynamic system from the bottom-up. Independent virtual agents, each of which represents an individual unit in the real population, are incorporated into the simulation. However, at this stage of the analysis, we do not incorporate agent-to-agent interactions and social networks. This model should be treated as the *first step* towards a fully developed ABM incorporating potential social interactions, and for use as the benchmark for further analysis. The main aim is to demonstrate that a simple enough model incorporating multiple heterogenous agents can predict a volatile and noisy CPG market with a high degree of accuracy. Once the underlying theoretical model of choice has been suitably validated, further compounding factors can be introduced in the model for more theoretical and/or empirical experimentation.

This paper's primary aim is to address the twin issues of individual heterogeneity and volatility, and to illustrate the relationship between the two. To demonstrate this link, a model based on random choice is used as a benchmark, which does not incorporate agent-level heterogeneity in tastes and preferences. The choice probabilities in the random model are generated from the data. The data consists of customer transactions within the fresh juice category from LeShop (www.leshop.ch), an online supermarket based in Switzerland. The simulations within the main model and the corresponding validation strategy results in a good fit of the out-of-sample predictions in comparison to the benchmark, particularly the direction and frequency of changes in market shares of product groups (where products are grouped on the basis of brands and flavors). At the micro-level, a remarkably good out-of-sample fit is achieved for individual households, which drives the

accuracy of the model at the macro-level. The model significantly outperforms the benchmark at all levels.

Over and above the stated aim of demonstrating the link between volatility and heterogeneity in consumer goods markets, this paper has two overarching objectives. The first is to establish a simulation-based approach which is not only able to estimate market share movements over time, but also predicts individual choices of products and features with a reasonable degree of accuracy. The second is to provide a validated benchmark model which can be used to explore abstract scenarios, carry out “what if” exercises and more generally, explore theoretical constructs which can in turn explain a variety of complex phenomena. Examples of potential directions of theoretical exploration are numerous, with multiple motivations [10] and are discussed later. Such normative exploration necessitates the existence of a benchmark which has been tested and validated, against which reasonable comparisons can be made. This is precisely what a model like the one presented here may be used for.

The time period considered here is one calendar year, but both the simulation methodology and the behavioral model are flexible enough to be used for examining the short-term and medium to long-term dynamics. The latter, especially the medium term, can be studied with small modifications to the current model, and a brief discussion on this is provided in the concluding section of this paper. Tackling the longer term using a similar methodology might be complicated, given that the environment and consumers might “drift” away from the calibrated model; this needs to be incorporated into the main model itself or addressed separately. Household-level choice within a CPG context also involves additional aspects which are outside the scope — timing of purchase, quantity bought, product portfolio, etc. being a few of them. In spite of these shortcomings, the model stands out from the rest for its simplicity, and hence implementation and extensions are easy.

1.1. *Background*

It is important to note that there exists a substantial amount of prior literature on quantitative analysis of CPG markets, especially with relevance to the effect of pricing and promotions on the evolution of market shares. Multinomial logit, nested and mixed logit and probit models have been heavily used in this area (see [3, 13] for details). Factors such as brand choice, purchase quantity and category incidence have usually been the focus of these studies (see for instance [4, 6, 17]). Most of the analysis has been carried out almost exclusively using a static framework, where the temporal element was ignored. These models performed well in picking out short-term responses to marketing mix policies. However, not as widespread and not as successful, are a few papers dealing with the issue of the long-term evolution of markets in response to multiple marketing strategies [7, 21, 22]. A handful of

papers have used dynamic programming as well to analyze the impact of promotions [25, 26], citing that discrete choice models overestimate promotion effects.

In spite of all the research done in this area, some questions remain unanswered which researchers and practitioners alike find very hard to tackle. For instance, what are the key differentiating factors which lead to variation in behavior within the population and how does this heterogeneity affect market dynamics? Is it possible to model the evolution of a market over time? Additionally, they lack the ability to take into account a number of compounding factors which may have important implications with regard to market behavior. For instance, potential nonlinearities (such as social networks), disruptive changes (paradigm shifting innovations in products and technologies), dynamic nature of market fundamentals (shifts in consumer preferences), changing nature of how firms and consumers interact in today's world (effect of the internet and viral marketing techniques for instance) are just a few examples of such factors. This is where a simulation-based model such as the one presented here is useful [18], forming the basis of further experimentation and analysis.

Model validation using real life data is an important component of any simulation-based study. Although various definitions of validation exist in the literature, there is one common theme among all: A validated model will possess a satisfactory range of accuracy matching the simulated model to the real world phenomenon [11, 14, 29]. Most authors also stress that in typical models which study large complex entities such as markets, validation should be carried out at multiple levels. For instance, with a model of a market such as the one presented in this paper, not only should the simulation model mimic the macro-level dynamics of the real market (macro-level validation), individual agents should also suitably mimic the behavior of the real households or consumers they represent (micro-level validation).

One important aspect of validation is the calibration of model parameters in order to match the simulation output with real data. The Fagiolo and Windrum survey [11] proposed three alternative calibration strategies — the indirect calibration approach, the Werker–Brenner approach and the history-friendly approach. While the first two involve more rigorous quantitative techniques, the third encompasses a more general qualitative method. For more details on all three, please see the following [12, 28] for indirect calibration and Werker–Brenner approaches respectively; [14, 20] for the history-friendly approach. This paper does not focus on any of these approaches exclusively, although the model specification lends itself more towards a direct quantitative approach rather than a qualitative one.

The choice model presented here has been developed specifically with the CPG industry in mind — where consumers make frequent purchases, tastes and preferences have an important role to play and firms engage in frequent price/promotion led marketing strategies. Additionally, it lends itself to a novel validation technique. Since our methodology validates the model both at the macro and the micro levels, two different fitness metrics are used for the two separate levels. At the macro-level,

the performance is judged by how closely market shares of product groups evolve over time in the simulation vis-a-vis the real data. At the micro-level, we focus on household level preferences relating to individual products and product attributes. We divide the data into *three* sets — the first used for initialization of agents, the second for direct calibration and the third for testing out-of-sample predictions.

The rest of the paper is organized as follows. Section 2 characterizes the underlying model which forms the basis of the agent-level behavior in the simulation. The data used in validating the simulation results is described in Sec. 3 while the actual validation methodology is described in Sec. 4. Next we present the results from the modeling and validation exercise in Sec. 6 while the paper concludes in Sec. 7 with a brief discussion and directions for future research.

2. The Market Model

A theoretical ABM framework is developed in [18], which incorporates the “4 P”s of marketing mix (product, price, placement and promotion) within a simulation framework. Our approach is similar to the extent that we start at the level of an individual shopper — a linear ordinal utility-based behavioral model is used to characterize the choice mechanism. However, we deviate in the model specification itself. Following are the key assumptions made in our model.

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 N . Each product is endowed with a set of M 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 or characteristics it embodies [8, 19, 24]. Hence, we make the following final assumption in our model of consumer choice:

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

The product specific characteristics are quantified as a M -dimensional vector belonging to a *characteristic space*, where each dimension represents one attribute. This vector is then called the *address* of the product. Each consumer’s preference is defined using a complementary *ideal point*, 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 $k \in K$, let $X_k = (x_k^1, x_k^2, \dots, x_k^M)$ be the product address for k . For any consumer $i \in N$, consider $\lambda_i = (\lambda_i^1, \lambda_i^2, \dots, \lambda_i^M)$ as i 's ideal point. Let P_k be the price of the product k . In order to characterize the “distance” of consumer i 's ideal point with a given product k , we use the 1-norm distance^a measure, defined as,

$$D_k^i = \sum_{j=1}^M |x_k^j - \lambda_i^j|.$$

This is simply the sum of *absolute distances*, along each dimension of the characteristic space, between the two vectors. Consumer i 's subjective utility from product k is hence characterized as,

$$U_i(k) = \omega_1^i d_k^i + \omega_2^i p_k \quad \text{where } \omega_2^i = 1 - \omega_1^i, \quad 0 \leq \omega_1^i \leq 1. \quad (1)$$

Identity 1 refers to the parametrized utility function where,

$$d_k^i = -\frac{D_k^i}{\max_{j \in K} D_j^i} \quad \text{and} \quad p_k = -\frac{P_k}{\max_{j \in K} P_j},$$

and ω_1^i , ω_2^i are the weights place by individual i on d_k^i and p_k respectively, in determining the utility from product k .^b Note that P_k is the *per unit* price of product k and not the listed price.

Next we consider price-based promotions, where prices of specific products are reduced (discounted) for individual or groups of consumers. Consequently, we need to differentiate between the *gross* and *net* prices, i.e. P_k^{gross} and P_k^{net} respectively where the discount on the product k is, $\Delta P_k = P_k^{gross} - P_k^{net}$. Identity 1 is redefined as,

$$U_i^{net}(k) = \omega_1^i d_k^i + \omega_2^i p_k^{net} \quad \text{where } \omega_2^i = 1 - \omega_1^i, \quad 0 \leq \omega_1^i \leq 1, \quad (2)$$

in order to incorporate promotions, where p_k^{net} is the relative price defined as above, but now in net terms. Note that identity 2 is the simplest possible characterization of an utility function incorporating promotions implicitly. A more complete characterization should explicitly take into account the incidence of promotion itself as well as the quantity bought as a result.

^aAny distance norm applicable to the Euclidean space \mathbb{R}^n can be used here, without any change in the results. Our use of this particular distance measure was entirely based on simplicity of form.

^bThe relative measures are used in order to normalize each component between 0 and 1, so that no single one dominates the function *numerically* although either one of them might dominate *functionally*, for instance if ω_1^i is equal to or very close to 0 or 1. In such a case, the individual agent's choice is based purely on either price or characteristics, and not on a convex combination of both.

3. Data

The data used in the current analysis consists all transactions within the fresh fruit juice category from LeShop, for the period of January 2006 to December 2006. Given the nature of the fruit juice category and available data, we consider three dimensions of product characteristics for the analysis — brand, flavor and pack size. Table 1 gives a break up of the products in terms of the identified brands and flavors. The third characteristic dimension — that of pack size — is expressed in grams, which ranges from 280 gm to 12,552 gm.

Each recorded transaction contains the following items: a household ID, product ID(s) of the product(s) purchased, the week number (indexed from 1 to 52) when the transaction was made, the net price paid, the discount applied if any, and finally the quantity of the product purchased. The timeline of transactions is indexed by weeks. An associated data table provides information about each product sold in that year and which has the following items: the product ID, the pack size, the flavor of the product, and the brand selling the product.

Each household ID represents a unique user of the LeShop website, and forms the basis of the agents in the simulation. Each product ID, similarly represents a unique product (Stock Keeping Unit or SKU), and hence is part of the set of alternatives available to the consumer. A few households had to be eliminated from the analysis given that they were infrequent purchasers. As a result, the only households considered were those who had at least three or more transaction points over the whole time period, with at least two of those transactions within the weeks 25–52. The final data comprised of 55 unique SKUs, 2,435 households and 28,179 transactions for all of the 52 weeks.

Table 1. Product break up by brand and flavors.

Brand	No. of products	Flavor	No. of products
Isola	2	Kid's Drink	3
Danone	2	Grape fruit	2
Mickey's Adventures	2	Nectar	7
Oasis	1	Pineapple	2
Capri Sonne	2	Apple	11
Granini	5	Multi-fruit	12
Michel	5	Orange	12
Obi	2	Other fruit	6
Nectar	8		
Hohes C	4		
Ramseier	5		
Max Havelaar	1		
Actilife	3		
M	8		
Gold	5		
Total	55		55

4. Validation

For the purpose of validation, the full year's transactions data is split into three sets along the temporal dimension — weeks 1 to 24, 25 to 38 and 38 to 52. The first set is used for initialization of individual ideal points in the characteristics space, the second for calibration of the choice model, and the third for testing the predictions made by the simulation.

- (1) *Initialization.* The *first* partition is used for initialization of agent-specific taste and preferences in the simulation. We initialized the product specific characteristic vectors and the agent-specific ideal points using the transaction history covering weeks 1 to 24 in the data set.
- (2) *Calibration.* The parameter space is partitioned in a suitable manner. Simulations are run for all partitions of the parameter space. The results are recorded for every individual agent for every time step. Using the *second* partition of the data set, we calibrate each agent individually, i.e. obtain the intersection of parameters which provide the best in-sample predictions for that particular agent. The optimum parameter set is then recorded for each agent. The optimization method is described in more detail below.
- (3) *Testing.* The final step involves the use of the parametrized agents to make out-of-sample predictions. We use the *third* data partition for this purpose. A Monte-Carlo type of method is used, where multiple runs of the simulation are made where each run corresponds to a random draw from the optimized parameter set of each agent. The results are collated and statistically compared against the data under consideration.

4.1. Initialization

The characteristics space is defined as the subset $[0, 1]^3 \in \mathbb{R}^3$, i.e. numerically, and the maximum and minimum values attached to any one dimension are 1 and 0 respectively. For each characteristic, the unique categories were assigned a value based on their total sales volume in the weeks 1–24, normalized by the maximum within that dimension. For instance, within the dimension representing brand, the one with the highest total sales volume (Gold) is assigned the value 1 while the one with the lowest (Isola) is assigned 0. All other brands were placed equi-spaced within $(0, 1)$, with each brand's position proportional to the relative sales volume (for e.g., *M* has the second highest sales volume and hence is placed at 0.86 in the brand dimension). All ties were resolved randomly. The relative positions of brands and flavors in the characteristics space for the given data, correspond to the ordering in Table 1.

Next we estimated a proxy for the ideal point of each agent using the transactions history within weeks 1 to 24 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 categories purchased, with the purchase frequency used as the weight. This is repeated for every dimension. Given the household and

its transaction history, we generate a three-dimensional vector in the characteristics space, which represents its ideal point.^c

This method does have a drawback when we consider households whose purchases concentrate on the two ends of the scale, in which case the weighted average falls somewhere in the middle, which is *not* representative of its preferences and happens only for a minority of cases. This can be seen from the spread of distances of agents' ideal points to their actual purchases (within partition-1 of the data), the mean and standard deviation of the spread over all agents is 0.285 and 0.149 respectively. Figure 4 shows this distribution in a histogram. The ideal points remain static for the rest of the analysis (given in Assumption 3).

4.2. Calibration

Calibration is done on both the macro and micro levels separately. For the macro-level, we aim to fit the evolution of *market shares of product groups* to the actual data, while for the micro-level validation, the corresponding aim is to fit *household-specific choice of SKU and product characteristics*. As each level requires a different fitness metric over which agent-specific parameters are calibrated, we carry out the calibration exercise *twice* — once for each level of validation. We use *binary matching* of simulated versus real take-up of SKUs to calibrate at the market share (macro) level and use the *city-block metric* to calibrate characteristic take-up at the household (micro) level.

For calibrating agent i 's specific ω_1 (the superscript i is dropped for notational convenience), the simulation is run for time-steps (weeks) 25 to 38. The parameter space is discretized into 25 equi-spaced points, $0 = \omega_1^0 < \omega_1^1 < \omega_1^2 < \dots < \omega_1^{24} = 1$ and the simulation is run 25 times for each agent, with each run corresponding to one value of ω_1 within the discretized space. Since there are no nonlinearities involved in terms of social influences and feedback in the model, we can run the simulation sequentially for each agent in turn without having to carry it out simultaneously for all agents. Moreover, given that we have considered deterministic rational choice here, the number of runs for each agent for a specific parameter value is limited to one. 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. The results of the calibration exercise is provided in Fig. 8.

4.2.1. Binary matching

The binary-matching calibration is done to optimize each virtual agent's utility function for optimum *SKU* choice. For each agent and each parameter value, the

^cFor example, suppose household i purchases the brands M and Gold two and three times respectively in the first half of the year. M 's and Gold's position in the $[0, 1]$ interval along the brand dimension of the characteristics space is given by 0.86 and 1 respectively. Hence, household i 's ideal point coordinate along the brand dimension is $\frac{2 \cdot 0.86 + 3 \cdot 1}{5} = 0.94$. The coordinates along the remaining two dimensions are computed similarly.

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 give week t ,

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

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,

$$b_t^i(S_t^i) = \frac{1}{N_t^i} \sum_{k=1}^K n_{t,k}^i \delta(k, S_t^i), \quad \text{where } \delta(k, S_t^i) = \begin{cases} 1, & \text{if } k = S_t^i \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

In the above definition, $n_{t,k}^i$ is the quantity of product k purchased by household i in week t and N_t^i is the total number of purchases by the household in that week, so that $N_t^i = \sum_{k=1}^K n_{t,k}^i$. For agent i , and for each parameter value ω_1^p , $p = 0, 1, 2, \dots, 24$ (the superscript i is suppressed), the binary-matching score is averaged over all of the purchase weeks to obtain an overall score,

$$B^i(\omega_1^p) = \frac{1}{|T^i|} \sum_{t \in T^i} b_t^i(S_t^i(\omega_1^p)), \quad (4)$$

where T^i is the set of weeks where household i made a purchase and $|T^i|$ is the cardinality of set T^i .

The set of best parameter values for agent i , Ω_b^i is then constructed as,

$$\Omega_b^i = \left\{ \omega_1^p \mid B^i(\omega_1^p) = \max_p (B^i(\omega_1^p)), p = 0, 1, \dots, 24 \right\}. \quad (5)$$

Identities 3, 4 and 5 spell out the macro-level calibration strategy. For any agent i , every match with corresponding household's purchase in the real data is given a score of 1 for the particular parameter value ω_1^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 set Ω_b^i . See Fig. 8(a) for a distribution of calibrated ω_1 across all agents in the binary matching case.

4.2.2. City-block metric

We use the city-block metric^d in order to calibrate the utility function of agents for optimum choice of *product characteristics*. Once again, let $S_t^i \in K$ be the simulated

^dThe name derives from the commonly used concept of city-block distance which measures the absolute distance between two points in Euclidean space, by measuring along each dimension independently. In the calibration exercise presented here, each dimension over the three-dimensional characteristic space is checked independently for a match and the results added up for the overall score — similar to the distance measure — and hence the name.

choice made by agent i in week t . Following Sec. 2, let x_k^m be the m th characteristic of product k , where M is the total number of relevant characteristics. A city-block metric assigns a score of 0 per characteristic where it matches with the real data and 1 for a mismatch. The characteristic matching score for agent i in week t is then defined as,

$$c_t^i(S_t^i) = \min_{\substack{k \in K \\ n_{t,k}^i \neq 0}} \sum_{m=1}^M \Delta^m(k, S_t^i) \quad \text{where } \Delta^m(k, S_t^i) = \begin{cases} 0, & \text{if } x_k^m = x_{S_t^i}^m \\ 1, & \text{otherwise.} \end{cases} \quad (6)$$

Here $0 \leq c_t^i \leq M$. Since $M = 3$ 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. Note that a score of 0 in most cases imply an exact SKU match as well, except when multiple products are purchased within a week.

The set of best parameter values for agent i , Ω_c^i is found in the following manner. First, we define a binary variable $\psi_t^i(\omega_1^p)$, which is used to judge whether a characteristic matching score is “good” enough, which in turn is used to determine the optimal ω_1^p .

Definition 1. The parameter value ω_1^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.

$$\psi_t^i(\omega_1^p) = \begin{cases} 1, & \text{if } c_t^i(S_t^i(\omega_1^p)) \leq \frac{1}{25} \sum_{p=0}^{24} c_t^i(S_t^i(\omega_1^p)) \\ 0, & \text{otherwise.} \end{cases}$$

As before, an overall score $C^i(\omega_1^p)$ is obtained for each parameter setting by averaging over all purchase weeks:

$$C^i(\omega_1^p) = \frac{1}{|T^i|} \sum_{t \in T^i} \psi_t^i(\omega_1^p). \quad (7)$$

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

$$\Omega_c^i = \left\{ \omega_1^p \mid C^i(\omega_1^p) = \max_p (C^i(\omega_1^p)), p = 0, 1, \dots, 24 \right\}. \quad (8)$$

See Fig. 8(b) for a distribution of calibrated ω_1 across all agents for the city-block metric case. Once the Ω_b^i and Ω_c^i sets have been identified for each agent/household, we proceed with the final stage in the analysis — that of testing our model with the out-of-sample data.

4.3. Testing

The model testing exercise is carried out on out-of-sample transactions data covering weeks 39 to 52. For each agent i , we now use both the sets Ω_b^i and Ω_c^i in multiple runs. The use of either set Ω_b^i or Ω_c^i is dependant 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 $\omega_1 \in \Omega_{b,c}^i$ is selected at random. As before, at each purchase week, the agent makes a choice from that week's product choice set using the parametrized 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 that week's predicted choice.

Since choice of quantity is not a part of the behavioral model, it is matched from the actual data in the testing phase. If the simulated choice is identical to the actual choice, the quantity bought by the agent is copied from the data. However, if the simulated and actual choices do not match, then the agent is made to purchase a quantity, q of the product. Quantity q is calculated using the average liquid volume (pack size) per transaction of the household vol^i , which is calculated on the first six months of data. Define,

$$q = \min \left\{ n \in Z \mid n \geq \frac{vol^i}{x_S^1} \right\}$$

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.

4.4. The benchmark

A random choice model is used as a benchmark to compare the results of the model described above. This random model is based on probability distributions generated from the data itself. All transactions from week 1 to week 38 are jointly used to generate the vector of probabilities $(P_1, \dots, P_j, \dots, P_{55})$, such that,

$$P_j = \frac{N_j}{N} \quad \forall j \in [1, 55]$$

where N_j is the number of transactions in which SKU j was purchased, and N is the total number of transactions. As defined above, P_j represents the *average market share* of SKU j within weeks 1 to 38. Note that this random model disregards the

effect of heterogeneity in individual consumer preferences as well as specific price effects. By definition, it captures the *observed* market level preferences, without considering any underlying factors which led to the observations.

As part of the experimental procedure, each agent is then made to choose from all SKUs based on the probability distribution generated above. The choice of purchased quantity by each agent is decided in the same manner as defined in Sec. 4.3. As before, the experiments are run a 100 times for each agent and the modal value considered as the simulated/predicted choice.

4.5. Notes on validation

A caveat is necessary to clarify some aspects of the validation methodology developed for the model.

First of all, a single data set has been used in the whole exercise to initialize, calibrate and test the model. The authors acknowledge the need to carry out some form of “cross-validation” with related data sets in order to test the robustness of this methodology. Several possibilities arise in this regard. The most straightforward one is to establish this method within a different product category from the same source (i.e. transactions from another group of products within LeShop in this case). The set of consumers would most likely be a different one, with some intersection between the two. It would also be possible to establish the same methodology using transactions data from a different source as long as the data structure is more or less similar (i.e. transactions carry information about individuals’ purchase of products, which link up with standardized product characteristic descriptions). An interesting deviation from the above would be to track consumer households across *different* shopping channels to test for the consistency of the model — assuming that most consumers are largely consistent in their shopping habits across channels. This would involve initialization and calibration of households within one data set and testing it in another.

All three methods of cross-validation are currently being tested, achieving comparable results in most product categories with a few specific exceptions (an interesting and open research question remains on why this is so). Unfortunately, current restrictions on the dissemination of the results based on some of this data prevents us from discussing them any further here.

Additionally, a simplifying procedure is followed with regard to the quantity chosen by agents in the testing phase of the model, although market share — the key component in macro-validation — is a function of quantity. This work specifically focuses on the choices of product and product characteristics made by households, and modeling quantity choice is beyond scope here. Predicting the quantity of products bought by any household involves a number of household specific factors (size, storage capacity, durability of the product, etc.) which are not available in the current data. More information regarding the households themselves and their shopping habits are required in order to incorporate this into the model.

5. Simulation Setup

The simulation is designed such that each virtual agent represents one consumer household in the data within a Matlab environment. The simulation consists of 2,435 agents whose behavior is defined by the model described above. Each agent represents an unique household in the data, with parameter set (λ_i, ω_1^i) , where $\lambda_i = (\lambda_i^1, \lambda_i^2, \lambda_i^3)$. Note that, $\lambda_i \in [0, 1]^3$ and $\omega_1^i \in [0, 1]$. As mentioned earlier, ideal point λ_i is initialized using the first partition of the data, which leaves ω_1^i as the only parameter to be calibrated using the second partition.

At every time-step in the simulation, agents choose *one* product from a subset of 55 total products. We do not make the assumption that all products were available at all weeks covering the simulation, but we use the transactions table to determine the *available* choices.

Definition 2. A product $k \in K$ 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 evaluates the prices, discounts, product characteristics and selects the one product which maximizes its utility. The following rules are implemented in the simulation:

- The quantity of fruit juice purchased by the household in the data determines the quantity purchased by the agent.
- If the household made no purchases that week, then no purchase is made by the agent.

These rules are essentially simplifying the model but as mentioned before, we are not modeling incidence or quantity bought, just the choice itself. 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.

6. Results

Out-of-sample (weeks 39 to 52) predictions and goodness-of-fit results are presented separately for the market (macro) and household (micro) levels. Results from the random model are used for comparison. In order to differentiate between the main model and the benchmark, we refer to them as the ABM and the Random model respectively.

6.1. Market level

At the level of the market as a whole, the models are tested on the basis of comparison between predicted and actual *market shares*. Given the large number of products, we group them on the basis of brands and flavors. Table 2 reports the degree of accuracy with which the simulated market shares match the actual ones in the data. We provide two measures — the average *relative difference* between the actual and simulated market shares per brand/flower ($r_{b/f}$) and the correlation coefficient ($c_{b/f}$) between the two. For each brand b , r_b is computed as:

$$r_b = \frac{1}{14} \sum_{t=1}^{14} \frac{|S_t^b - A_t^b|}{A_t^b}$$

while S_t^b and A_t^b are the simulated and actual market shares of b in week t . r_f is computed similarly for each flavor f . Both measures considered together indicate how well the models perform in terms of capturing the volatility in the market.

Table 2. Comparison of mean relative distance (r_b) and correlation coefficient (c_b) between predicted and actual market shares for ABM versus random model.

Brand	ABM		Random		Flavor	ABM		Random	
	r_b	c_b	r_b	c_b		r_f	c_f	r_f	c_f
Isola	0.56	0.75	0.82	0.37	Kid's	0.73	0.22	0.66	0.29
Danone	0.63	0.36	0.72	−0.28	Drink				
Mickey's	0.83	0	0.80	−0.23	Grape fruit	0.97	0	0.49	−0.20
Adventures					Nectar	0.32	0.50	0.69	0.14
Oasis	0.67	0.23	0.68	0.39	Pineapple	0.42	0.30	0.52	0.14
Capri Sonne	0.20	0.78	0.46	0.38	Apple	0.36	0.59	0.42	0.16
Granini	0.55	0.64	0.80	0.41	Multi-fruit	0.30	0.88	1.37	−0.33
Michel	0.31	0.36	0.57	0.57	Orange	0.12	0.59	0.47	−0.56
Obi	0.73	0.24	0.65	0.26	Other fruit	0.31	0.55	0.27	0.49
Nectar	0.41	0.46	0.52	0.20					
Hohes C	0.33	0.76	0.59	0.24					
Ramseier	0.23	0.86	0.71	0.36					
Max Havelaar	0.33	0.95	0.51	0.46					
Actilife	0.29	0.43	0.69	−0.41					
M	0.70	0.89	1.05	−0.23					
Gold	0.14	0.76	0.43	−0.06					
Average	0.53	0.57	0.67	0.16	Average	0.44	0.45	0.61	0.01

The key results from the macro-level validation presented in Table 2 can be summarized as follows. First, for most brands and flavors, the ABM predicts the *direction of change* of market shares very well. It performs reasonably well with regard to the *magnitude* of market shares, although not as accurately as the prediction of direction. Second, both the correlation coefficient as well as the relative difference is significantly better in the ABM for brands/flavors which have high market shares overall than those which are in the lower end. Generally, brands with market shares consistently less than 2% of the market are difficult to predict, whereas the others are well predictable either in direction or both in direction and magnitude. The pattern within flavors is similar as well (with the threshold around 8% of market share), but with one important difference — as compared to brands, the average prediction is better in magnitude (0.44 against 0.53) but worse in direction (0.45 against 0.57). Equivalent values for the Random model are far worse, for most brands and flavors (Mickey’s Adventure in brands and Kid’s Drink, Grape Fruit and Other fruit are the only exceptions). Overall, the ABM significantly outperforms the Random model in predicting the volatility in the market. Note that choice of quantity (number of SKU’s being purchased) is not endogenous in the model; it is not a decision variable for the agents although pack size is. Figures 1 and 2 show the predictions of the ABM versus actual market shares of a selected set of brands and flavors over weeks 39 through 52. The rest of the brands and flavors can be found in Figs. 5, 6 and 7 in the Appendix.

Table 3 presents the number of weeks a brand or a flavor has been on promotion within the 14 week period. As can be seen, incidences of promotions in the form of price discounts were quite rare for most brands, but not so when product groups are grouped as flavors. Yet actual market shares of both product groupings exhibit high volatility, *which is also picked up by the ABM simulation*, driven through SKU-level choices. Agents differ from one another in terms of how strongly they react to price changes and inherent preferences. Since different attribute combinations are under promotion at different times, varying responses to promotions by any one agent as well as across agents lead to the high volatility in the simulations. On the other hand, the Random model, which disregards individual level preferences and price responses, is unable to pick up on market share movements either in direction or magnitude, both at the level of attributes as well as on average. However, as mentioned before, examination of the household/micro-level choices made by individual households is necessary in order to establish the validity of the model.

6.2. Household level

The performance of the models at the micro-level provides the basis of their performance at the macro-level. A counting scheme similar to that of the city-block metric is used to test the model at the micro-level. The accuracy of predictions in a household’s choice of characteristics is measured, in each dimension of the

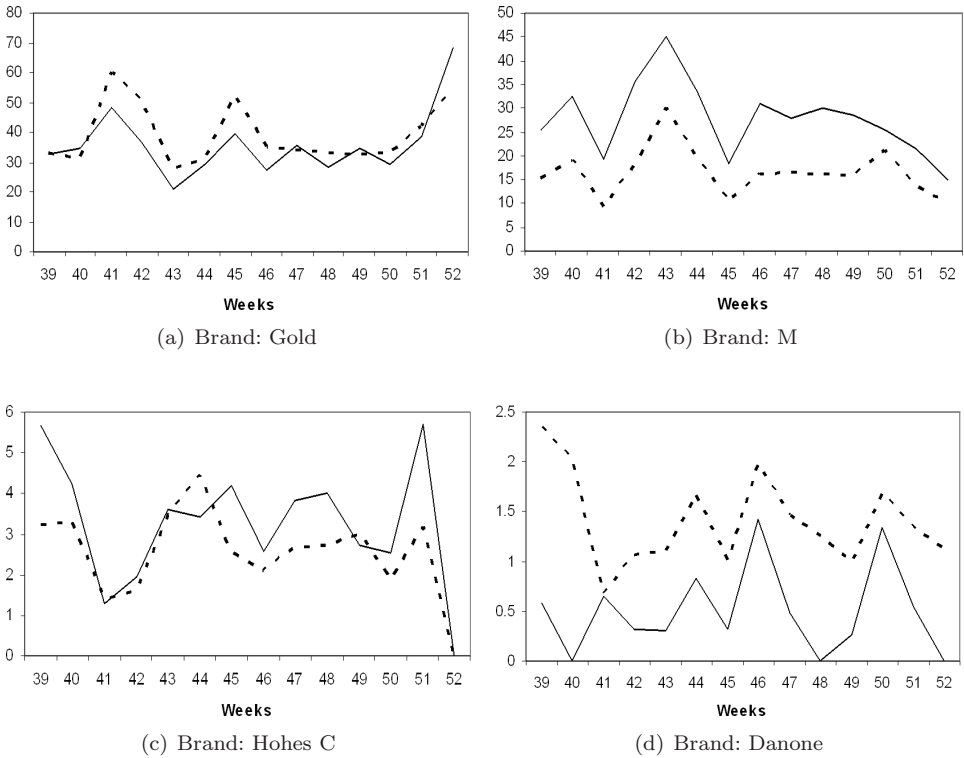




Fig. 1. Weekly market share predictions of ABM versus actuals of selected brands. Solid line — predicted; dotted — actual.

characteristic space independently, as well as jointly along subsets of dimensions. Given a particular dimension within the characteristic space and for any given household, the number of times a correct prediction is made along that dimension, is estimated. For measuring the joint predictions along all dimensions, the number of times the simulation respectively predicts 3, 2, 1 and 0 characteristics correctly is also estimated for every household, irrespective of what those characteristics are. The former implies finding the proportion of times the model correctly predicts the choice along a dimension, independent of other characteristic dimensions. The latter implies a joint prediction “score” per household per transaction (as defined in [6]). Note that a score 0 for agent i implies that, i ’s choices in the simulation matched the corresponding real household’s purchase in *all* dimensions, while a score of r , where $0 < r \leq 3$, implies only $3 - r$ characteristics matched correctly. Table 4 and Fig. 3 summarizes the results for the ABM and Random models.


Part 1 of Table 4 indicates that the household-specific SKU choice was predicted correctly 35.74% of the times in the ABM and 2.1% in the Random model, from a pool of 8,213 transactions in total. The total 8,213 transactions incorporate




(a) Flavor: Other



(b) Flavor: Multi-fruit



(c) Flavor: Nectar



(d) Flavor: Kid's Drink

Fig. 2. Weekly market share predictions of ABM versus actuals of selected flavors. Solid line — predicted; dotted — actual.

Table 3. Number of weeks within weeks 39 to 52, when brands and flavors are in promotion.

Brand	No. of weeks	Flavor	No. of weeks
Isola	0	Kid's Drink	0
Danone	0	Grape fruit	7
Mickey's Adventures	0	Nectar	1
Oasis	0	Pineapple	5
Capri Sonne	0	Apple	4
Granini	4	Multi-fruit	10
Michel	2	Orange	11
Obi	0	Other fruit	2
Nectar	1		
Hohes C	2		
Ramseier	1		
Max Havelaar	1		
Actilife	0		
M	6		
Gold	5		

Table 4. Prediction Results at the Household Level for the ABM model. Numbers in parenthesis indicate the corresponding values for the Random model.

1. SKU Predictions				
Exact Matches	2936			
Number of Transactions	8213			
Accuracy	35.74%	(2.1%)		
2. Characteristics Predictions				
<i>Joint across dimensions</i>				
Percentage Accuracy (%)	Score 0	Score 1	Score 2	Score 3
100	24.43%	11.21%	10.14%	8.25%
75–100 (excl.)	4.56%	1.89%	2.75%	1.97%
50–75 (excl.)	11.13%	9.20%	10.88%	11.66%
25–50 (excl.)	7.97%	8.13%	9.16%	10.06%
0–25 (excl.)	51.83%	69.53%	66.98%	67.97%
	~ 100%	~ 100%	~ 100%	~ 100%
Mean across hh.	37.18%	20.82%	21.94%	20.07%
Standard Deviation	21.21	17.17	17.01	15.94
3. Characteristics Predictions				
<i>Per dimension</i>				
	Brand	Flavor	Size	
Mean across hh.	61.68%	73.52%	39.81%	
	(5.9%)	(9.5%)	(3.7%)	
Standard Deviation	42.1425	35.6669	42.6331	

households for whom all 100% of the transactions were correctly predicted, as well as households for whom none were correct and those for whom predictions were only partially correct. This distribution is presented in Part 2 of Table 4 for ABM, which 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, and importantly as well, for 24.43% of the households, all product characteristics from all their transactions were correctly predicted. For the small percentage of cases where multiple brands/flavors and/or sizes are bought by the same household in the same transaction, we might be able to match all characteristics, but not the product. This explains why the accuracy in SKU prediction (35.74%) is slightly less than the mean of the proportion of Score 0 for all households (37.18%). Also, it is only a minority of households (8.25% as seen in the Score 3 column), that predictions are inaccurate in all dimensions for all transactions. Figure 3(a) summarizes mean and standard deviations graphically.

Part 3 of Table 4 and Fig. 3(b) summarize the accuracy of household level predictions within each characteristic dimension individually. Flavor is definitely the most predictable characteristic, followed by the brand and finally the size, for both the ABM and the Random models. However, the ABM significantly outperforms




Fig. 3. Household level prediction mean and variation in the characteristics space.

the Random model in all three. The low degree of predictability-of-size is understandable given that it is likely to be influenced by individual consumption rates, frequency of visits to the shop, the incidence and size of promotions, etc. — which have not been currently included in the behavioral model. Note that it is the low level of predictability-of-size which is pulling down the overall SKU prediction figure as well. If we discount product size, these results definitely indicate that ranking products based on attribute-specific preferences and prices is able to mimic household level choice behavior significantly — thus establishing the overall validity of the simulation model.

7. Conclusion

This paper used a bottom-up simulation based approach in order to model the volatile fresh juice market within the CPG industry, using checkout data from an online supermarket as the basis for empirical validation. Virtual agents were designed using a simple behavioral model and were initialized, calibrated and finally tested on the basis of predictions made out-of-sample. The behavioral model took account of both the heterogeneity in preferences for product attributes as well as the heterogeneity in individual responses to pricing. A random choice model was used as a benchmark to compare the performance of the multi-agent model.

The validation methodology tested the model and its predictions at both macro and micro levels. The volatility seen in the market at the macro-level, in terms of week to week market share movements of product groupings (brands and flavors), were captured by the model with high degree of accuracy — both in terms of direction of change and to a slightly lesser degree, the magnitude of change. The accuracy in predicting the volatile market at the macro-level is driven by the accuracy in the prediction of choices of SKUs and various product attributes by

individual households at the micro-level. On the other hand, once the heterogeneity in the behavioral model is removed, as in the case of the benchmark, predictions at the micro-level fall dramatically in accuracy and consequently, it is unable to capture the volatile market.

Although the approach used in the simulation is bottom-up and multi-agent based, it lacks one key feature of a full fledged agent-based model — agent-to-agent interactions through social networks. The data set used here does not allow us to go any further, for example, to test the presence of networking effects. This might be possible with a different data set within a different market, as is part of ongoing research. However, what the current model is effective for is in running *what if* scenarios — for instance, testing the efficacy of alternative pricing/promotion strategies, alternative product offerings and even imposing agent-to-agent interactions within the current model to analyze levels of deviation from current behavior.

The behavioral model itself can and should be enhanced in a number of ways. Instances of promotions can be incorporated explicitly into the analysis instead of being implicitly considered through the net price. Also, rationality is a strong assumption to impose on the agents; this can be relaxed in varying degrees to test for effects on market share evolution. The current model allows the researcher to explicitly incorporate cognitive elements within agent-level behavior as well (such as static or shifting attitudes towards specific characteristics or characteristic dimensions), in line with [23]. Additionally, allowing for the introduction of new products and exit of existing ones would be an interesting extension within the current framework. The last would be especially important when examining the longer term dynamics of markets. Within this context, it would be interesting to analyze the effects of changing preferences (instead of static ideal points), changing attitudes and even development of social norms with regard to shopping behavior.

All in all, the current model's simplicity, parsimony and ability to combine theory, simulation and real life data seamlessly should make it very appealing for researchers and practitioners alike. Moreover, the possibilities of enhancing this model in many directions makes it an exciting starting point and benchmark for future experimentation and analysis.

Acknowledgments

The authors would like to thank Frank Smith from UCL, Shail Patel, Iqbal Adjali, Benjamin Dias from Unilever and Vera Hazelwood from The Smith Institute for valuable comments at all stages of the project. We would also like to acknowledge the crucial role played by the Industrial Mathematics Knowledge Transfer Network in enabling this project. Funding was provided jointly by EPSRC and Unilever.

Appendix A. Figures




Fig. 4. Spread of distances of agents' ideal points from the purchases with which the ideal points are calculated.

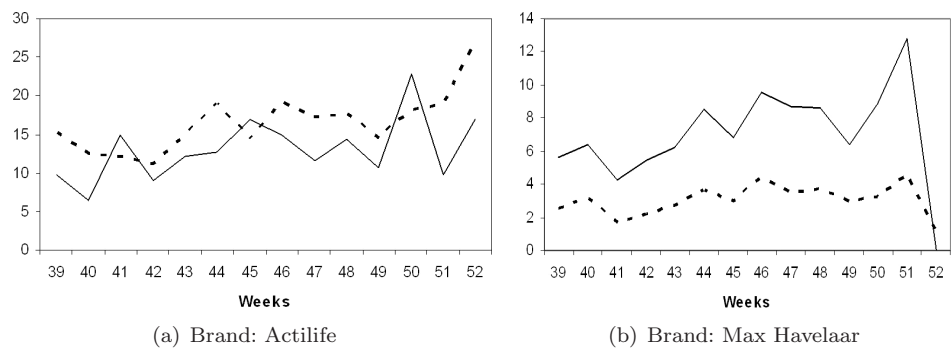





Fig. 5. Predicted and actual weekly market shares of brands. Solid line — predicted; dotted — actual.




(c) Brand: Ramseier



(d) Brand: Nectar




(e) Brand: Obi




(f) Brand: Michel

Fig. 5. (Continued)



(a) Brand: Granini



(b) Brand: Capri Sonne

Fig. 6. Predicted and actual weekly market shares of brands (Cont.). Solid line — predicted; dotted — actual.

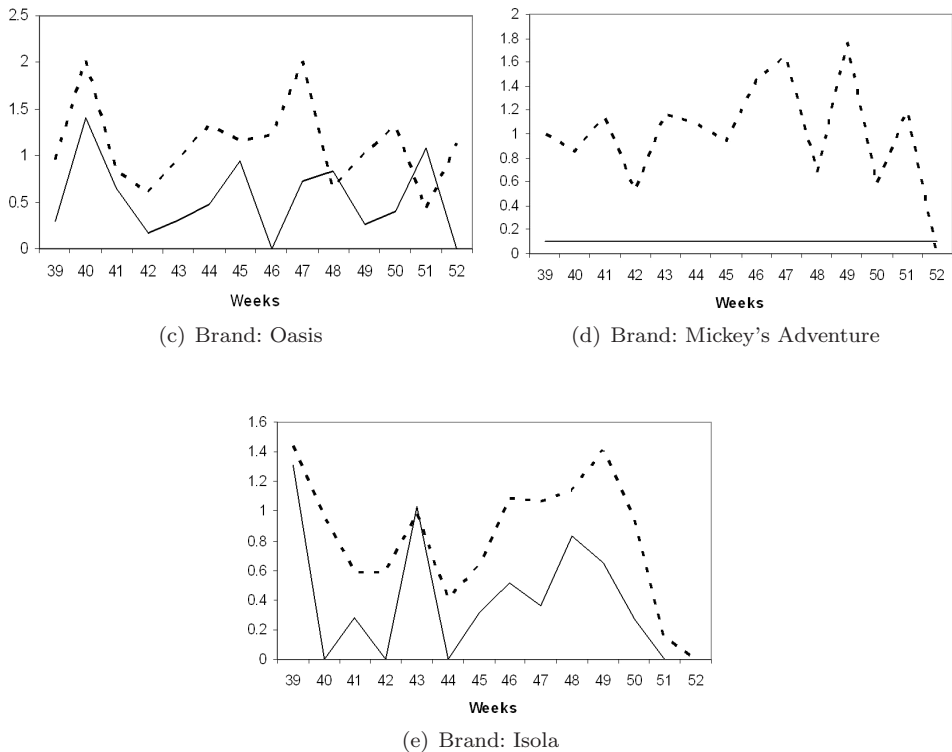


Fig. 6. (Continued)

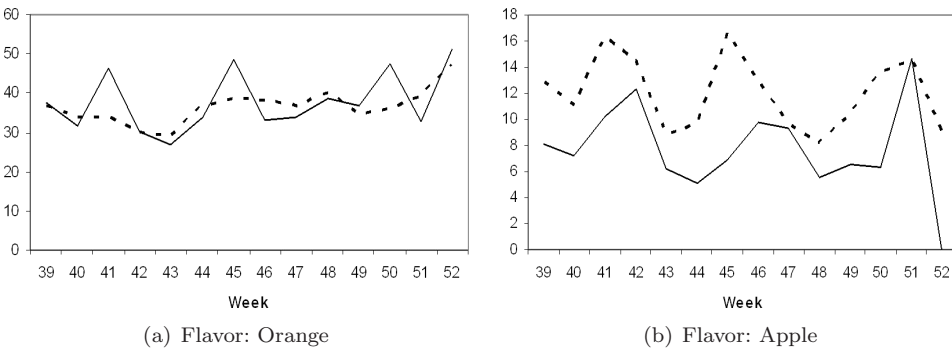




Fig. 7. Predicted and actual weekly market shares of flavors. Solid line — predicted; dotted — actual.




(c) Flavor: Pineapple




(d) Flavor: Grape Fruit

Fig. 7. (Continued)



(a) Binary Matching

Fig. 8. Distribution of optimal ω_1 across all agents for each calibration type. For agents with multiple optima, the average is considered.



(b) City Block Metric

Fig. 8. (Continued)

References

- [1] Ailawadi, K., Lehmann, D. and Neslin, S., Market response to a major policy change in the marketing mix: Learning from procter and gamble's value pricing strategy, *Journal of Marketing* **65** (2001) 44–61.
- [2] Allenby, G. M. and Rossi, P. E., Marketing models of consumer heterogeneity, *Journal of Econometrics* **89** (1998) 57–78.
- [3] Anderson, S. P., de Palma, A. and Thisse, J.-F., *Discrete Choice Theory of Product Differentiation* (Cambridge: The MIT Press, 1992).
- [4] Bell, D., Chiang, J. and Padmanabhan, V., The decomposition of promotional response: An empirical generalization, *Marketing Science* **18** (1999) 504–526.
- [5] Blattenberg, R. and Wisniewski, K. J., Price induced patterns of competition, *Marketing Science* **8** (1989) 291–309.
- [6] Chan, T., Narasimhan, C. and Zhang, Q., Decomposing promotional effects with a dynamic structural model of flexible consumption, *Journal of Marketing Research* **45** (2008) 487–498.
- [7] DeKimpe, M., Hanssens, D. and Silva-Risso, J., Long-run effects of price promotions in scanner markets, *Journal of Econometrics* **89** (1999) 269–291.
- [8] DeSoete, G., Carroll, J. and DeSarbo, W., The wandering ideal point model: A probabilistic multidimensional unfolding model for paired comparisons data, *Journal of Mathematical Psychology* **30** (1986) 28–41.
- [9] Epstein, J. and Axtell, R., *Growing Artificial Societies — Social Science from the Bottom UP* (MIT Press, Cambridge, MA, 1996).
- [10] Epstein, J. M., Why model? *Journal of Artificial Societies and Social Simulation* **11** (2008).

- [11] Fagiolo, G., Moneta, A. and Windrum, P., Empirical validation of agent-based models: A critical survey, in *International Workshop on Agent-Based Modeling and Economic Policy Design* (University of Bielefeld, Germany, 2005).
- [12] Fagiolo, G., Dosi, G. and Gabriele, R., Matching, bargaining and wage setting in an evolutionary model of labor market and output dynamics, *Advances in Complex Systems* **14** (2004) 237–273.
- [13] Franses, P. and Paap, R., *Quantitative Models in Marketing Research* (Cambridge University Press, 2001).
- [14] Garcia, R., Rummel, P. and Hauser, J., Validating agent-based marketing models through conjoint analysis, *Journal of Business Research* **60** (2007) 848–857.
- [15] Gilbert, N., Jager, W., Deffuant, G. and Adjali, I., Complexities in markets: Introduction to the special issue, *Journal of Business Research* **60** (2007) 813–815.
- [16] Gilbert, N. and Troitzsch, K. G., *Simulation for the Social Scientist*, 2nd edn. (Open University Press, Maidenhead, UK, 2005).
- [17] Heerde, H., Gupta, S. and Wittink, D., Is 75 percent of the sales promotion bump due to brand switching? No, only 33 percent is, *Journal of Marketing Research* **40** (2003) 481–491.
- [18] Jager, W., 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** (2007) 868–875.
- [19] Lancaster, K., A new approach to consumer theory, *Journal of Political Economy* **74** (1966) 132–157.
- [20] Malerba, F., Nelson, R., Orsenigo, L. and Winter, S., History-friendly models of industry evolution: The computer industry, *Industrial Corporate Change* **8** (1999) 3–40.
- [21] Mela, C. F., Gupta, S. and Lehmann, D. R., The long-term impact of promotion and advertising on consumer brand choice, *Journal of Marketing Research* **34** (1997) 248–261.
- [22] Pauwels, K., Hanssens, D. and Siddarth, S., The long-term effects of price promotions on category incidence, brand choice, and purchase quantity, *Journal of Marketing Research* **39** (2002) 421–439.
- [23] Richetin, J., Sengupta, A., Perugini, M., Adjali, I., Hurling, R., Greetham, D. and Spence, M., A micro-level simulation for the prediction of intention and behavior, *Cognitive Systems Research* **11** (2010) 181–193.
- [24] Sengupta, A., Greetham, D. and Spence, M., An evolutionary model of brand competition, in *Proceedings of IEEE SSCI 2007: Symposium on Artificial Life* (2007), pp. 100–107.
- [25] Sun, B., Promotion effect on endogenous consumption, *Marketing Science* **24** (2005) 430–443.
- [26] Sun, B., Neslin, S. and Srinivasan, K., Measuring the impact of promotions on brand switching when consumers are forward looking, *Journal of Marketing Research* **40** (2003) 389–405.
- [27] Tesfatsion, L., Agent based computational economics: A constructive approach to economic theory, in *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*, eds. Tesfatsion, L. and Judd, K. L. (North-Holland, Amsterdam, 2006).
- [28] Werker, C. and Brenner, T., Empirical calibration of simulation tools, in *Papers on Economics and Evolution*, Vol. 0410 (2004).
- [29] Windrum, P., Fagiolo, G. and Monnetta, A., Empirical validation of agent-based models: Alternatives and prospects, *Journal of Artificial Societies and Social Simulation* **10** (2007) 8.