A Multi-Product Individual-Level Model for New Product Sales: Forecasting and Insights

By

Carmen Liutec
Doctoral Candidate
University of Houston

June 2013

Dissertation Co-Chairs:

Rex Yuxing Du
Associate Professor – Hurley Professor of Marketing
University of Houston

Ed Blair
Professor of Marketing and Entrepreneurship and Department Chair
University of Houston

Committee Members:

Ye Hu
Associate Professor of Marketing
University of Houston

Arthur Francia
KPMG Peat Marwick Professor of Accounting
University of Houston

Wagner Kamakura
Ford Motor Company Professor of Global Marketing
Duke University

A Multi-Product Individual-Level Model for New Product Sales:

Forecasting and Insights

A Multi-Product Individual-Level Model for New Product Sales: Forecasting and Insights

ABSTRACT

We develop a novel approach for modeling new product trial and early repeat purchase behavior, and we apply this approach in the context of consumer packaged goods. Our approach takes advantage of the cross-individual, cross-product, cross-time data that is increasingly available from retail customer relationship management programs as well as research panels. It enables us to account for differences in consumers' intrinsic preferences for new products as well as for differences in their responsiveness to marketing variables, during both trial and early repeat purchases.

By leveraging these uncovered individual differences, we attempt to achieve three goals. First, we aim to improve the accuracy of post-launch sales forecasts based on data from a period that can be as short as two to three months, tapping into the fact that each individual trial or early repeat purchase observed during the post-launch period sends a different signal about the new product's sales potential. Second, we aim to provide more informative diagnostics for managers to act upon. (e.g., how to re-position the new product or target it for a particular consumer segment). Finally, we aim to investigate potential empirical regularities regarding consumer responses to new products (e.g., potential differences in responsiveness to marketing variables between early and late triers).

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1. INTRODUCTION

The number of new product introductions in the consumer packaged goods (CPG) industry has steadily increased in recent decades. According to the U. S. Census Bureau (2000), the number of new product introductions in the food category tripled in the 1980- 2000 period, while the number of new products in the beverage category increased more than five times during that period. More recently, the number of new CPGs introduced every year has increased at an average rate of 5.4% per year (productscan.com). However, according to Catalina Marketing, the failure rates for new CPGs range between 70% and 90%. Even for those products that survive past the first year, only about 25% have sales above $7.5 million during their first year, and less than 3% of new CPGs exceed first-year sales of $50 million, which is often considered the benchmark of a highly successful launch (Schneider and Hall 2011).

Given low success rates and high development and launch costs, there is need for a model of new CPG adoption that yields insights into how companies can market new products more effectively, provides accurate sales forecasts with limited data, and enables companies to take action early in the life of a product if it is not on track for success. Such a model should fulfill several requirements:

- First, it should be able to provide insights into both trial and repeat purchase patterns for the given product and yield managerially actionable recommendations. For example, identifying those consumers who are most likely to try the new product can be very useful for post launch targeting decisions. Likewise, identifying the reasons for trial can provide information about the long term likelihood of success of the product. For instance, some early adopters might try a new product only due to its low introductory price, others might purchase it because they have a strong preference for the product, while others might try it because of some intrinsic variety-seeking tendency. Different proportions of these consumer segments would have very different implications for repeat purchasing rates and recommended marketing actions.

- Second, a good model should be able to provide accurate sales forecasts based on only a short post launch period.
- Third, a useful model should be able to make good forecasts based on the behavior of a relatively small sample of consumers. Given the high costs of launching the new product nationally, manufacturers have strong incentives to launch a new product on a small scale, in a test market (such as BehaviorScan of IRI).

This study attempts to address these issues by taking a novel approach to new product early post-launch sales forecasting. While many existing forecasting models in the literature are applied to aggregated post-launch sales data (e.g., Fader, Hardie and Zeithammer 2003), the proposed model leverages the rich information contained in individual-level variation in trial and repeat purchase behavior. In addition, while the large majority of existing forecasting models only study one product at the time, we leverage information from multiple new product adoption occasions for the same consumer – which in turn allows disentangling behavioral elements that are specific to the consumer from those that are specific to the particular product analyzed.

This study contributes to the new CPG forecast literature on several dimensions. First, while there are many new product forecasting models in the literature, most of them are developed at the aggregate level. Even when heterogeneity is accounted for, it is treated as means to obtaining a better model fit, while ignoring potentially valuable information content of individual behavior. Second, there are few empirical studies that attempt to characterize new product trial behavior for CPGs and identify reasons for different adoption patterns based on actual purchase data (e.g., Du and Kamakura 2011). Rather, most studies that attempt to explain new product trial behavior are based on survey data (e.g., Wood and Swait 2002; Im et al. 2003). Third, little attention is given in the literature to early post-launch repeat purchase behavior. None of the studies that account for repeat purchases of new CPGs (e.g., Schweidel and Fader 2009, Fader et al. 2004) attempt to make inferences about systematic patterns in these purchases, such as potential differences between trial and repeat responsiveness to marketing variables, or potential changes in the perceived positioning of a new product after the consumption experience.

This study aims at achieving three goals. First, we aim to provide more informative diagnostics to help companies manage new products. Our results have implications for promotional budget allocations for the new product, as well as for potential perceptual repositioning which could improve new product success. Second, we investigate potential empirical regularities regarding consumer responses to new products. For instance, for the four product categories analyzed, our results suggest a higher responsiveness to marketing variables for early (vs. later) triers, which might indicate increased awareness of the in-store informational environment for these consumers. Third, we aim to improve the accuracy of early post-launch sales forecasts when limited sales data are available. Our results show good performance of the proposed model with only four weeks of data, and better accuracy than alternative forecasting models proposed in the literature.

The remaining of the paper is organized as follows. Section 2 gives a brief overview of the existing new product sales models and underlines the new elements brought by our proposed approach. Section 3 describes the proposed model, while Section 4 summarizes the steps involved in new product sales forecasting. Section 5 describes the data available for estimation, while Section 6 presents the results of the analysis. Section 7 concludes.

2. EXISTING MODELS FOR NEW PRODUCT SALES

Numerous new product sales models have been proposed in the literature. The following discussion briefly reviews the major types of models.

2.1. Single Product Models

Many of the earlier studies in this area have modeled the adoption of durable goods at a product category level. These studies typically use the Bass (1969) diffusion framework and relax its simplifying assumptions. For instance, unlike most previous studies which were concerned with modeling diffusion for the first few years after the introduction of a new product category, Kamakura and Balasubramanian (1988) take a long term view on new product adoption. They study the entire product life of several durables with long sales histories and investigate the role of price in the diffusion process. In their model, prices can have an effect on the diffusion process by impacting the market potential, the adoption probability, or both. For the products analyzed, their results suggest that price impacts the adoption probability of relatively higher price goods, while it does not have a significant effect on the diffusion process of lower priced durables. Bass, Krishnan and Jain (1994) attempt to reconcile the mixed patterns of results found in previous studies that investigated the effect of marketing variables on the Bass-type of diffusion models. They propose a generalization of the Bass model, which includes marketing variables (i.e., price and advertising), and reduces to the traditional Bass (1969) model under certain regularity conditions for the decision variables. Their empirical results for three new durable product categories support their contention that, if the period-to-period changes in the values of the marketing variables are relatively constant, the original Bass model provides a good fit with the advantage of parsimony, while for less regular evolution of marketing mix variables over time, the generalized Bass model has higher explanatory power.

Other studies focus on extending the Bass framework to include replacement purchases. For instance, Olson and Choi (1985) develop a more general version of the Bass model that covers replacement purchases, which relaxes the assumption of constant fraction of repeat buyers. More specifically, they assume that the timing of repeat (i.e., replacement) purchases depends on the new durable's product life, which is modeled as a stochastic process that can be estimated from the data. Kamakura and Balasubramanian (1987) further expand this framework. Their model allows other factors besides product's age to affect the duration of product life, incorporates exogenous variation in the total market potential and price, allows for wide ranging flexibility of the life curve specification and can be estimated on total sales data, without requiring information on the cumulative number of adopters.

For products with a short life and frequent number of repeat purchases, models estimated on aggregate data can be problematic, as the inferred trial and repeat patterns can have misleading implications (e.g., Fader and Hardie 2003). Individual-level information is needed to address this issue. One of the earliest new product models incorporating individual level information is the stochastic evolutionary adoption model (STEAM) of Massy (1969). Massy proposes a "depth of repeat" type of model, applied to forecasting the sales of a new non-durable product for a panel of consumers. The model allows for some degree of consumer heterogeneity, i.e., it posits that consumers' transitions between classes of repeat purchases can be impacted by both the time of purchase (relative to product launch) and the time since last purchase.

As more detailed consumer level information became available, new product trial and repeat forecasting moved towards more sophisticated modeling approaches than the early depth-of-repeat type of models. For instance, Sinha and Chandrashekaran (1992) propose a split hazard model for adoption which accounts for both the timing of adoption and the probability of eventual adoption (i.e., explicitly allows for a segment of potential non-adopters). Further, cross sectional heterogeneity can be incorporated through covariates. Chandrashekaran and Sinha (1995) develop a model that includes both repeat purchases and actual purchase volume. They apply the model to data on the adoption of personal computers by a sample of US firms, and they find a significant improvement in the forecasting accuracy over models which do not account for purchase volume.

Many of the most recent new product forecasting models have been developed in the context of the consumer packaged goods industry, where detailed individual level data has been increasingly available. While the success of a new CPG product depends mostly on the product's ability of attract repeat purchases, many different studies have focused exclusively on forecasting new product trial. Fader, Hardie and Wisniewski (1998) and Fader, Hardie and Zeithammer (2003) present an empirical comparison of the most commonly used trial models in the context of consumer packaged goods. They find that predictive accuracy is significantly higher for those models that account for consumer heterogeneity and include the effect of marketing covariates. Out of the numerous models investigated, the exponential-gamma with covariates emerges as the most accurate specification, with the lowest prediction error (Fader et al. 2003).

Also, while multiple models for purchase timing which accommodate repeat purchases have been developed in the marketing literature (e.g., Gupta 1991; Jain and Vilcassim 1991; Seetharaman 2004), few of these models have been applied to new product forecasting. A more recent trial and repeat purchase model that nests several of the multiple-event timing models previously proposed in the literature is the dynamic changepoint model of Fader, Hardie and

Huang (2004). Their specification allows for variation of the consumer-level purchase rates over time, being able to accommodate many different shapes of the new product sales curve. Heterogeneity is incorporated with a gamma distribution of buying rates in the population; change in behavior across purchase occasions is captured by assuming that consumers may receive a new “draw” of the buying rate on the next purchase occasion. They apply the model to data for one new beverage product and find very good forecasting accuracy. Schweidel and Fader (2009) extend this model to include the situation where consumers’ behavior may change with experiencing the product. More specifically, consumers’ purchase rates come from two different regimes, one for trial-type of purchases (the first few purchases, before enough experience with the product is gained), and one “steady-state” regime afterwards.

Some of the more recent studies have moved towards investigating the effect of other types of individual-level influences on new product adoption. For example, Iyengar, Van den Bulte and Valente (2011) question whether social contagion affects new product diffusion. They apply a hazard model in the context of a new prescription drug, combining individual level adoption data with survey measures and social network data; their findings indicate effects of social contagion on new product adoption, and moderating effects of recipients’ perception of their opinion leadership.

2.2. Multiproduct Models

All the studies discussed above rely on information observed for a single focal product (or product category). However, data available for other similar products can prove very useful for either improving forecasting accuracy, or making more general inferences about new product behavior. Part of the existing research in the new product literature has adopted this multiproduct perspective, for either descriptive or predictive purposes.

Some studies use information across multiple products for explanatory purposes. For instance, Golder and Tellis (1997) use aggregate level data for 31 product categories to investigate the timing of new product "takeoff " (i.e., the first dramatic increase in sales early after introduction). They specify a hazard model for the time of takeoff, with controls for price, year of introduction, market penetration, economic and product category variables, and find that price and economic conditions are the most important factors influencing the timing of the first drastic sales increase. Van den Bulte (2000) uses a hierarchical model calibrated on aggregated data for 31 durable product categories to find out whether diffusion speed varies over time. His results show strong evidence of diffusion speed acceleration; further, the acceleration can almost entirely be attributed to changes in economic and demographic factors, as well as to the changing nature of the products (e.g., the occurrence of competing standards). Steenkamp and Gielens (2003) combine multiple sources of data (i.e., consumer panel data, consumer surveys, retail data, advertising expenses and expert ratings) for 239 new CPGs to investigate the relative importance of several factors (i.e., consumer characteristics, product category characteristics, marketing strategy and marketing communications) on a new product's trial probability. Du and Kamakura (2011) investigate individual-level drivers of new CPG adoption, in the form of social contagion/network diffusion effects. They develop a multiproduct model that allows for spatial and temporal heterogeneity in contagion and controls for various confounds, applied to 67 new products. Their results document significant effects of social contagion on the adoption process for a large set of the products analyzed.

Other studies use information across multiple products to improve forecasting accuracy for a focal new product. One of the earliest models in this stream of research is the ASSESSOR model of Silk and Urban (1978), designed for pre-test-market forecasting. The model combines information from laboratory experiments (e.g., simulated product choice situations and laboratory advertising exposure) with survey and interview data for a sample of customers in the intended target market. Together with managerial input for the new product (e.g., positioning strategy and marketing plan) this information is used for calibrating a choice model that helps to predict the new product's share in the test market. Other studies use information on multiple historical new products to relate the parameters of some aggregate diffusion-type model to various product characteristics, which in turn allows approximating the parameters of a prelaunch or early post launch model for the new product. For instance, Rao and Yamada (1988) build on the diffusion-type modeling with repeat purchases literature and propose a prelaunch forecasting model which can be calibrated on past product introductions. They apply their model to the pharmaceutical industry, by first calibrating aggregate level diffusion type of models for past new drug introductions; next, survey data on physician perceptions about the drugs are used to explain the parameters of the aggregate level models, and to project prelaunch aggregate model parameters for the focal new drug. Hahn, Park, Krishnamurthi and Zoltners (1994) use a similar approach to forecasting, proposing an aggregate diffusion type model with four classes of repeat, which incorporates the effect of competitive marketing actions. They calibrate the model on a set of 21 pharmaceutical products, and they use a 2-stage approach to relate the parameter estimates to product, market and launch strategy characteristics.

2.3. Our Approach

Our proposed study fits in the area of multiproduct forecasting literature. However, unlike the studies above, our approach involves integrating both multiproduct and individual level information within a single step, with the goal of early forecasting of trial and repeat sales. Several papers in the marketing literature use a similar approach, but they are missing some of the above elements. Moe and Fader (2001) propose a model that incorporates information for multiple products, which they apply in the context of music sales. They account for consumer heterogeneity in the form of latent segments (imputed from aggregate data), and they build on the idea that similar products will draw similar proportions of sales out of each consumer segment. However, their model is a trial only model. Kamakura, Kossar and Wedel (2004) explicitly investigate the heterogeneity structure underlying trial in the context of pharmaceuticals, with the purpose of identifying the best targets for cross-selling new products. Our proposed model extends the Kamakura et al. (2004) framework to incorporate heterogeneity in responsiveness to marketing variables and to account for repeat purchases. Lee, Boatwright and Kamakura (2004) propose a prelaunch forecasting model, applied in the context of recorded music sales. Their study proposes improving forecasts by leveraging information from historical data on sales of past albums. Specifically, they develop a hierarchical Bayes model that incorporates multiple observed product attributes, and that allows for early sales projections when very few points of data for the focal product are available. While our proposed approach is conceptually similar, we augment potentially available information on product similarity[1] with

[1] Product similarity can be inferred in our model based on unobserved, latent product “dimensions”. However, information on the actual product attributes, etc. can be easily incorporated. The latent factor approach has the advantage of not requiring an apriori, arbitrary decision about which actual product attributes are relevant for new product adoption decisions.

historical data on individual-level purchase behavior. As such, our model should yield not only relatively precise forecasts, but also actionable recommendations in terms of targeting specific consumer segments.

Overall, our proposed model brings together three desirable features as depicted in Table 1. First, we develop an integrated model of trial and repeat purchase. Second, we account for individual-level heterogeneity and we allow the consumers' responsiveness to marketing variables, product preferences, etc. to vary across trial and repeat occasions. Third, our model incorporates historical data on past individual-level adoption behavior, leveraging information across multiple products. All these elements are expected to improve forecasting ability when only a short period of data is observed for the focal product. Moreover, the proposed model will provide a more complete picture of the individual and product level adoption patterns and will yield useful practical suggestions for marketers.

Table 1. Positioning in the Existing New Product Models Literature.

		Trial	**Trial + Repeat**
Aggregate Data	**Single Product**	<u>Post launch:</u> E.g., Bass (1969); Bass, Krishnan, and Jain (1994); Kamakura and Balasubramanian (1988).	<u>Post launch:</u> E.g., Olson and Choi (1985); Kamakura and Balasubramanian (1987).
	Multiple Products	<u>Prelaunch:</u> E.g., Lee, Boatwright and Kamakura (2004). <u>Post launch:</u> E.g., Moe and Fader (2001). <u>Explanatory:</u> E.g., Golder and Tellis (1997); Van den Bulte 2000.	<u>Prelaunch:</u> E.g., Hahn, Park, Krishnamurthi and Zoltners (1994); Rao and Yamada (1988).
		Trial	**Trial + Repeat**
Individual Data	**Single Product**	<u>Post launch:</u> E.g., Sinha and Chandrashekaran (1992); Fader, Hardie and Zeithammer (2003). <u>Explanatory:</u> Iyengar, Van den Bulte and Valente (2011).	<u>Prelaunch:</u> E.g., Massy (1969). <u>Post launch:</u> Chandrashekaran and Sinha (1995); Fader, Hardie and Huang (2004); Schweidel and Fader (2009).
	Multiple Products	<u>Prelaunch/explanatory:</u> E. g., Kamakura, Kossar and Wedel (2004). <u>Explanatory:</u> E.g., Du and Kamakura (2011); Steenkamp and Gielens (2003).	This study

3. MODEL DESCRIPTION

We develop an individual level hazard model of trial and repeat purchases for multiple new products. The model accounts for individual level heterogeneity and allows for potential changes in purchase behavior (in terms of product preferences, response to marketing variables, etc.) across first time (trial) purchase and subsequent repeat purchases.

Let TP_{ijt} represent the "trial purchase" hazard, i.e., the likelihood that consumer i, who has not yet tried product *j* at the beginning of the t^{th} post-launch period, will try it during period *t*. The log-hazard rate for trial is specified as:

$$(1)\qquad \ln\left(TP_{ijt}\right) = \alpha_{ij}^{TP} + \beta_{j1}^{TP} Usage_i + \beta_{j2}^{TP} H_i + \pi_{ij1}^{TP} Price_{jt} + \pi_{ij2}^{TP} Promo_{jt} + \beta_{j3}^{TP} TimeSLaunch_{jt}$$

where:

α_{ij}^{TP} = the baseline propensity for consumer *i* to make a trial purchase of product *j*;

β_{j3}^{TP} = linear time trend on the hazard rate, where $TimeSLaunch_{jt}$ is the number of weeks since the launch of product *j*;

$Usage_i$ = average weekly purchase volume (units) for consumer *i* in the product category. This variable is calculated based on an initialization period (i.e., a period of one year before the start of the estimation sample period) and reflects differences in usage rates across consumers.

H_i = measure of "variety seeking" behavior for consumer *i*, operationalized as the brand-level

Herfindahl index[2] (based on consumer *i*'s purchases during the same initialization period as above). Lower values of this variable indicate a tendency to seek variety.

$Price_{jt}$ = price of product *j* during week *t*;

$Promo_{jt}$ = dummy variable for in-store promotions, equal to *1* if product *j* was on display and/or featured during week *t*, and *0* otherwise.

The consumer-specific factors above might be expected to have an effect on the time of trial and likelihood of repeat purchase. All else equal, variety seeking might favor early trial regardless of the product or brand. Category level usage can also be expected to relate to adoption behavior; heavy users face the product choice decision more often than light users (i.e., consider buying in the product category more often) and hence they might be more likely to be early triers. Some existing research on adoption behavior for CPGs suggests that product usage is an important factor in determining trial patterns, with heavy users being among the first triers (e.g., Steenkamp and Gielens 2003, Taylor 1977, Morgan 1979).

During the periods following the trial purchase of a new product, the consumer may decide to make a repeat purchase. Let RP_{ijt} represent the "repeat purchase" hazard, i.e., the likelihood that consumer *i*, who has already tried product *j* before the beginning of period *t*, will make a repeat purchase during period *t*.

The log hazard rate for repeat purchase is specified as:

$$(2) \quad \ln\left(RP_{ijt}\right) = \alpha_{ij}^{RP} + \beta_{j1}^{RP} Usage_i + \beta_{j2}^{RP} H_i + \pi_{ij1}^{RP} Price_{jt} + \pi_{ij2}^{RP} Promo_{jt}$$
$$+\beta_{j3}^{RP} NRepeat_{ijt} + \beta_{j4}^{RP} (NRepeat_{ijt})^2$$

[2] This variable is computed as: $H_i = \sum_{m=1}^{M} s_{im}^2$ where: s_{im} = share of brand *m* in consumer *i*'s total purchases during the initialization period; *M* = number of brands in the product category.

where:

$NRepeat_{ijt}$ = number of repeat purchases of product *j* made by consumer *i* before the beginning of week *t*.

The number of purchases made in the past might be informative for consumer i's propensity to purchase the product in the future. For instance, having bought the product multiple times in the past might reflect a relatively strong preference towards the product or even a habit formation type of effect, and therefore increase the likelihood of another purchase; on the other hand, repeat purchases in the recent past might also result in a saturation effect, and therefore increase the time between successive purchases. The quadratic specification is flexible enough to allow for either of these potential effects.

Note that the effect of consumer and product variables is allowed to be different for trial and repeat purchases. Marketing mix elements, such as price and promotions, might have different effects across these two types of purchases. For instance, consumers might be more price-sensitive when deciding whether to try a new product. Similarly, category usage level and variety seeking might be more important in prompting early trial, but the actual positive or negative experience with the product might overweight these factors during the repeat purchase decision.

The responsiveness to marketing variables in the Trial and Repeat functions are allowed to have heterogeneous effects across consumers. All variables are allowed to have different effects across products.

With the exception of the baseline propensity terms { α_{ij}^{TP},α_{ij}^{RP}} which will be explained in detail later, all parameters follow a variance-components type of specification. More exactly, the individual and product level coefficient for any variable *k*, has the form:

(3) $$b_{ij}^{k} = b + \eta_j^k + \eta_i^k$$

where b captures the average (across individuals and products) effect of the variable k, η_j^k is a (zero-mean) random deviation from the average of the effect of variable k for product j, η_i^k captures the individual-specific deviation from the average, and η_j^k and η_i^k are independent from each other. The trial and repeat models described by equations (1) and (2) are connected in the parameter space, i.e., the trial and repeat parameters come from a common product and individual level distribution, respectively.

This type of specification is however, not suitable for the baseline propensity terms $\{\alpha_{ij}^{TP}$, $\alpha_{ij}^{RP}\}$. Specifying these terms in the form $\alpha_{ij} = \alpha + \eta_j^\alpha + \eta_i^\alpha$ has the unappealing implication that consumer *i*'s "preference" deviates from the "average preference" by the same quantity η_i^α for all products. The commonly used specification $\alpha_{ij} = \alpha + \eta_j^\alpha + \eta_{ij}^\alpha$ is also not quite useful: when the terms η_{ij}^α are allowed to correlate freely across products (and with the consumer-level η_i^k s for all other variables), the number of covariance parameters becomes extremely large.

An alternative is to impose a latent-factor structure on the baseline propensity terms (e.g., Du and Kamakura 2011, Kamakura et al. 2004). That is, we could use a specification of the form: $\alpha_{ij} = \alpha + \eta_j^\alpha + \Lambda_j Z_i$, where Z_i is a consumer-specific *P*-dimensional random vector with each element distributed i.i.d. *N (0, 1)* and Λ_j is a *P*-dimensional parameter vector to be estimated for each product *j*. This specification allows the baseline propensities to be correlated across products, while keeping the number of covariance parameters manageable.

As such, we specify the baseline propensities for Trial and Repeat as:

(4) $$\alpha_{ij}^{TP} = \alpha^{TP} + \eta_j^{TP} + \Lambda_{j1}^{TP} z_{i1} + \Lambda_{j2}^{TP} z_{i2} + z_{i3}^{TP}$$

$$\alpha_{ij}^{RP} = \alpha^{RP} + \eta_j^{RP} + \Lambda_{j1}^{RP} z_{i1} + \Lambda_{j2}^{RP} z_{i2} + z_{i3}^{RP}$$

where each individual-specific score z_{ip} *(p = 1,…, 3)* of the Trial and Repeat function is normally distributed with mean zero and unit variance.

The terms in equation (4) have an intuitively appealing interpretation. The product-level factor weights may be viewed as product-specific latent "attributes"[3] that pinpoint the product's location in a 2-dimensional trial and repeat perceptual space (i.e., Λ_{j1}^{TP}, Λ_{j2}^{TP} and Λ_{j1}^{RP}, Λ_{j2}^{RP}, respectively). Prior to trial, consumers' perceptions are mostly determined by the public information available about the product, such as brand, ingredients, packaging or messages communicated through advertising. Experience with the product through consumption might prompt an adjustment of the pre-trial assessment, resulting in a shift of the product in the perceptual space. For instance, the new sugar-free version of a soft drink might taste more similar to the original high-sugar content version than initially expected, hence shift closer to the original version in the repeat perceptual space. Together with the factor weights above, the factor scores z_{i1} and z_{i2} reflect consumers' heterogeneous preferences across products. The term z_{i3}^{TP} may be interpreted as an individual specific "early trier score"[4], which captures consumer's intrinsic propensity to try new products, after controlling for product preferences, category usage rates and variety seeking tendencies. All else equal, consumers with a higher z_{i3}^{TP} are more likely to be early triers. Similarly, consumers with a higher z_{i3}^{RP} are more likely to repeatedly purchase the new products tried (for instance, these consumers could be more discerning triers, i.e.,

---

[3] The latent factors could reflect, for example, a combination of both observable (e.g. ingredients or package size) and unobservable (e.g. brand reputation) attributes.

[4] The propensity to be an early trier is different from "intrinsic consumer innovativeness" as a generalized personality trait, although the two characteristics may be correlated to some degree. Existing studies find only a weak association between innovativeness and new product purchase behavior (e.g., Im, Bayus and Mason 2003); moreover, the strength of the association varies with the product categories analyzed (e.g., Foxall 1995).

consumers who are able to anticipate their preferences for the product even before the consumption experience).

The complete distribution of individual-level heterogeneity in the model can then be captured by the factor scores Z_i and by the all other random effects for the Trial and Repeat functions:

$$(5) \qquad E^i = \begin{bmatrix} \eta_i^{TP} \\ \eta_i^{RP} \end{bmatrix} \sim N[0, \Sigma^C] \, ; \qquad Z_i \sim N[0, I_4]$$

where each of the vectors η_i^{TP}, η_i^{RP} contain the stacked random deviates for consumer *i* for the corresponding parameters in the Trial and Repeat functions (i.e., the η_i^ks in equation 3) and I_4 is a *(4 x 4)* identity matrix.

The formulation of the model is completed by specifying the product-specific random effects for all variables (including the η_j quantities from the baseline propensity terms) to be jointly normally distributed with mean zero:

$$(6) \qquad E^j = \begin{bmatrix} \eta_j^{TP} \\ \eta_j^{RP} \end{bmatrix} \sim N \, [0, \Sigma^{UPC} \,],$$

with each of the vectors η_j^{TP}, η_j^{RP} containing the stacked random deviates for product *j* for the corresponding parameters in the Trial and Repeat functions (i.e., the η_j^ks in equation 3).

With a complementary log-log function for the Trial and Repeat rates, the conditional likelihood function for consumer *i* has the form:

$$(7)\quad L_i|\Theta = \prod_{j\in O_i}\left\{\prod_{\tau=1}^{t_{ij}-1}\exp(-TP_{ij\tau})\left[1-\exp(-TP_{ijt_{ij}})\right]\prod_{\tau=t_{ij}+1}^{T_j}[1 - \exp\left(-RP_{ij\tau}\right)]^{d_{ij\tau}}[\exp(-RP_{ij\tau})]^{1-d_{ij\tau}}\right\}$$

$$\prod_{j\notin O_i}\left\{\prod_{\tau=1}^{T_j}\exp(-TP_{ij\tau})\right\}$$

where:

$d_{ij\tau}$ = 1 if consumer *i* purchases product *j* at time τ, and 0 otherwise;

t_{ij} = adoption time for product *j* by consumer *i* (if product *j* is adopted before the end of the observation period);

T_j = end of observation period for product *j*;

O_i = set of products tried by consumer *i*.

Θ = set of all model parameters.

The model presented above allows making inferences about the individual level new product purchase behavior. For instance, we will be able to answer questions about the potential behavioral differences between "innovators" (i.e., people who consistently try newly launched products) and more "discerning" buyers, who only try a small number of new products. The model also allows uncovering potential behavioral regularities across trial and repeat purchases, such as systematic differences between responsiveness to marketing mix variables between trial and repeat.

The model is estimated through hierarchical Bayesian procedures (e.g., Train 2003), which are well equipped for handling large models. Details of the estimation procedure are provided in Appendix A.

4. IMPROVING NEW PRODUCT SALES FORECASTING

Bayesian methods provide a natural framework for incorporating information from various sources. What makes forecasting sales for new products difficult early during the post launch period is the scarcity of information, typically with only a few weeks of observed trial and repeat purchases.

Our proposed approach incorporates additional information, in the form of past adoption behavior for the target market. Trial and repeat purchases are available for many other products launched in the past. Incorporating the existing data on past adoption behavior into the calibration of the model for forecasting the sales of a new UPC helps reducing uncertainty in both the individual and product specific estimates.

More exactly, a large new product purchase history for a given individual yields quite precise estimates of the individuals' preferences and responsiveness to marketing variables. These estimates act as very informative priors when data for the newly launched product become available, resulting in more accurate predictions even with a very limited amount of data for the newly launched product. Intuitively, observing some consumers making trial purchases early during the post-launch period, would inform us that the new product is more similar to other products tried early by the same group of consumers, and that other consumers who have tried those products identified as similar are likely to also try the new product. Moreover, some approximation of repeat rates can also be obtained very early, by using the historical repeat behavior of the relevant groups of consumers as a prior.

Our approach is similar to the idea used in Lee et al. (2003), for their model of prelaunch forecasting of music sales. However, while their approach relies on inferred similarities across products, our proposed model builds forecasts starting at a micro level, relying on a good description of the behavior of individual consumers. As such, our model can yield actionable suggestions for manufacturers and retailers, by identifying the reasons behind observed post launch performances. In addition, the proposed model is also applicable to cases where the similarities across products are difficult to describe in terms of objective attributes (or to cases where intangible attributes are the main differentiating elements between seemingly similar products in terms of objective attributes).

The actual forecasting procedure has several steps. First, estimates for the historical set of products are obtained. Second, these estimates are used as priors for estimation of the "forecasting model", which is calibrated on data for both the focal product (first weeks available post launch, and incorporating more observations as they become available) and the historical adoption data. Third, posterior estimates for all (focal) product-level and individual level parameters are computed (e.g., Train 2003). Finally, the posterior estimates obtained above are used to simulate out-of-sample purchases for the focal product. More specifically, many out of sample "purchase paths" are simulated for each consumer, then averaged across the number of simulations and summed across consumers.

5. DATA DESCRIPTION

The data consists of weekly shopping records for panelists in one BehaviorScan[5] IRI test market. The data contain the shopping behavior of a consumer panel for a 6-year period. The dataset includes all shopping trips for five major grocery stores in the area. Store-level weekly data on pricing and in-store promotional activities (i.e., features and in-store displays) is also available for all products purchased by the panelists.

Data is available for purchases made in four food product categories (carbonated beverages, cereals, frozen dinners and salty snacks). Food product categories are good candidates for calibrating the proposed model, since they have a large number of new product introductions and have the widest potential target segment of buyers (unlike products such as diapers or beer).

We define "new product" as a newly introduced UPC. Some UPC level information is available for all product categories (short description, attributes, brand, etc.), which allows us to screen out temporary UPCs (e.g., those associated with special events, seasonal products, etc.). Further, we only include in the analysis those UPCs that were introduced in all stores in the sample, and were available for the entire period of analysis.

Within each product category, most new products are used for the model calibration, and several products are set aside for evaluating forecasting performance. More specifically, UPCs introduced during [Year 2, Week 1 – Year 5, Week 26] period are used for model calibration. UPCs introduced during the last 26 weeks of Year 5 are set aside for forecasting, which ensures

[5] *BehaviorScan* is an in-market testing service that utilizes small markets (towns with population under 200,000) for quantitative assessment of marketing activities, including "real life" testing of new CPGs prior to national launch. The market in our dataset has a population of 42,000, and the stores included in the sample account for over 90% grocery shopping revenue in that market.

that for every new UPC in the forecasting sample, at least 52 weeks of data are available for evaluating forecasting performance. Year 1 data are used for the initialization period (i.e., calculating the *Category Usage* and *H Index* variables discussed in Section 3[6]). Figure 1 provides a graphical representation of our sample period.

Figure 1. Sample Period.

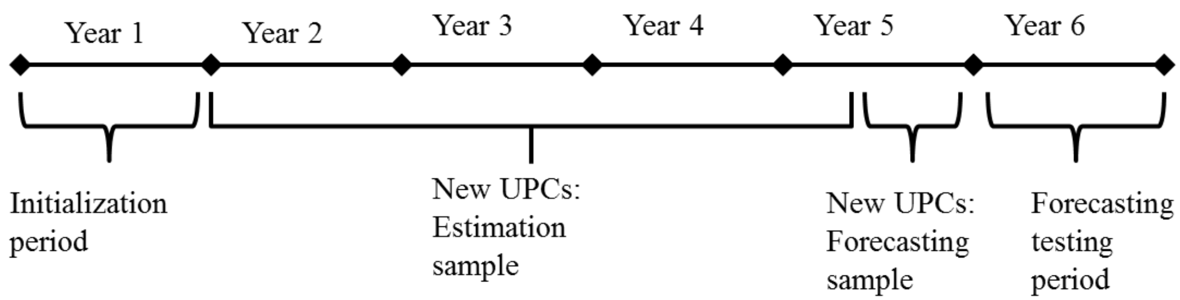

For each product category, we keep for analysis those panelists who: 1) make at least two purchases in the category during the initialization year[7] (i.e., non-users of the product category are excluded), and 2) remained in the panel for the entire 6-year sample period.

Table 2 shows the number of UPCs in the calibration/forecasting sample for each category. Tables 3 and 4, which show summary statistics for these samples, indicate a wide range of variation of trial and repeat rates across the products in the sample.

[6] Using a time window before the start of the estimation period for calculating *Category Usage* and *H Index* removes endogeneity concerns for these variables.

[7] The number of panelists included in the sample is quite insensitive to alternative usage rules based on 1, 2, 3 or 4 purchases in the category during the initialization period.

Table 2. Sample Size by Category.

	Number UPCs Estimation Sample	Number UPCs Forecasting Sample	Number Panelists
Beverages	55	8	1,149
Cereal	41	3	1,331
Frozen Dinner	47	8	883
Snacks	58	7	1,200
TOTAL	201	26	

Table 3. Descriptive Statistics: Estimation Sample.

		Penetration Rate[a] (%)	Average Repeat Volume[b] (units)	Price ($)	Promotion Frequency	Category Usage (units)	H Index
Beverages	Mean	12.06	1.70	2.78	0.29	1.30	0.37
	(St.dev)	(6.89)	(0.77)	(1.36)	(0.17)	(1.24)	(0.21)
Cereal	Mean	7.79	0.76	3.79	0.29	0.56	0.46
	(St.dev)	(4.74)	(0.43)	(0.54)	(0.05)	(0.52)	(0.19)
Frozen Dinner	Mean	4.89	0.62	3.23	0.14	0.42	0.61
	(St.dev)	(1.27)	(0.33)	(0.37)	(0.04)	(0.52)	(0.27)
Snacks	Mean	9.16	0.99	2.52	0.25	0.78	0.32
	(St.dev)	(9.52)	(0.68)	(0.59)	(0.17)	(0.75)	(0.19)

[a]*Penetration Rate* for a product is defined as the percentage of individuals making a trial purchase of the UPC.

[b]*Average Repeat Volume* for a product is defined as the total number of repeat purchases divided by the total number of trial purchases.

Table 4. Descriptive Statistics: Forecasting Sample.

		Penetration Rate[a] (%)	**Average Repeat Volume[b] (units)**	**Price ($)**	**Promotion Frequency**
Beverages	Mean (St.dev)	5.97 (0.85)	0.96 (0.40)	2.43 (1.44)	0.34 (0.09)
Cereal	Mean (St.dev)	5.61 (1.83)	0.47 (0.09)	4.03 (0.10)	0.17 (0.02)
Frozen Dinner	Mean (St.dev)	5.52 (1.65)	058 (0.23)	2.59 (0.84)	0.19 (0.06)
Snacks	Mean (St.dev)	19.39 (16.32)	1.49 (0.57)	2.49 (0.21)	0.43 (0.22)

[a]*Penetration Rate* for a product is defined as the percentage of individuals making a trial purchase of the UPC.
[b]*Average Repeat Volume* for a product is defined as the total number of repeat purchases divided by the total number of trial purchases.

We are using a 79-week period as the model calibration window. This period is long enough to capture trial/repeat behavior for the chosen categories, but also short enough so that the newly introduced UPCs are still available in the stores.

For assessing the model's predictive performance, we are using progressively longer windows of observations for the UPCs in the prediction sample. Most existing new product forecasting models used for predicting CPG success (e.g., Fader et al. 2003) tend to forecast quite well when 21-24 weeks of data are available. Hence our prediction window includes a maximum of 52 weeks post-launch, and we focus on forecasts when only shorter periods of data are available.

6. RESULTS

Next, we discuss the estimation results and evaluate the forecasting performance of the proposed model.

The parameter estimates for the four product categories are shown in Tables 5 and 6. The first column for each category shows the category-level means of the parameters, while the next two columns present the estimates of the product and individual-level heterogeneity. Qualitatively, the parameter estimates shown in Tables 5-6 are consistent across all product categories.

Table 5. Parameter Estimates: Beverages and Cereal

		Beverages			**Cereal**		
		Estimate	Product Level Variance	Individual Level Variance	Estimate	Product Level Variance	Individual Level Variance
Trial	***Intercept***	-7.499***	0.628		-8.216***	0.816	
	Categ.Usage	0.246***	0.293		0.353***	0.408	
	H index	-0.501***	0.283		-0.416***	0.373	
	Price	-0.275***	0.277	0.083	-0.157*	0.391	0.099
	Promo	0.240***	0.295	0.085	0.110	0.397	0.118
	TimeSLaunch	-0.384***	0.396		-0.399***	0.547	
Repeat	***Intercept***	-4.630***	0.610		-5.522***	0.790	
	Categ.Usage	0.095	0.312		0.190*	0.493	
	H index	0.030	0.280		-0.037	0.395	
	Price	-0.188***	0.289	0.079	-0.094	0.420	0.096
	Promo	0.258***	0.286	0.092	0.076	0.401	0.124
	NRepeat	0.254***	0.273		0.417***	0.538	
	NRepeat^2	-0.013	0.247		-0.114	0.386	
LogL		-96,090.87			-35,547.33		

+Product level and individual level variances are significant at 5% in all cases.

Table 6. Parameter Estimates: Frozen Dinner and Snacks

		Frozen Dinner			**Snacks**		
		Estimate	Product Level Variance	Individual Level Variance	Estimate	Product Level Variance	Individual Level Variance
Trial	***Intercept***	-8.944***	0.671		-8.133***	0.782	
	Categ.Usage	0.361***	0.382		0.425***	0.407	
	H index	-0.424***	0.324		-0.866***	0.264	
	Price	-0.360***	0.419	0.176	-0.089	0.291	0.063
	Promo	0.012	0.421	0.228	0.245***	0.340	0.093
	TimeSLaunch	-0.232***	0.368		-0.249***	0.328	
Repeat	***Intercept***	-5.925***	0.842		-5.361***	0.988	
	Categ.Usage	0.245***	0.471		0.264***	0.512	
	H index	0.139	0.391		-0.253**	0.342	
	Price	-0.282***	0.437	0.141	-0.064	0.306	0.065
	Promo	0.008	0.436	0.156	0.180***	0.305	0.093
	NRepeat	0.199	0.574		0.387***	0.456	
	$NRepeat^2$	-0.194**	0.382		-0.116*	0.269	
LogL		-19,081.63			-76,980.00		

+Product level and individual level variances are significant at 5% in all cases.

For the Trial function, *Category Usage* has a significant positive effect, suggesting that heavy users are most likely to be early triers. The *H index* has a negative impact, indicating that more “inertial” consumers, who habitually purchase only a small number of brands, are less likely to try the new products. The effect of *TimeSLaunch* is consistently negative, indicating that a consumer who has not already tried the new product becomes less likely to do so over time. Both *Price* and *Promotion* coefficients have the expected signs, but their significance varies across categories, suggesting that the importance of marketing mix variables during the trial decision might vary across types of products. For instance, for the Snack category, insignificant price effects and significant impact of in-store promotions might indicate that impulse is a strong determinant of trial purchases for this category. The opposite pattern can be expected for categories with more involved pre-choice deliberation and limited opportunities for in-store displays, such as Frozen Dinners.

For the Repeat function, the effect of observable consumer characteristics (i.e., *Category Usage* and *H Index*) becomes reduced, or even insignificant. Given the first purchase experience, the differences with respect to the likelihood of repeat between light and heavy users are less pronounced. Variety seeking/inertia is also a poor predictor of Repeat (likely due to both the trial purchase experience, and due to the fact that consumers tend to make trial purchases within their favorite set of brands to begin with). The number of past repeat purchases of the product seems to have an inverted-U type of effect, providing some support for both habit formation and saturation type of behaviors. Another difference between Trial and Repeat functions that holds for all four categories is the lower magnitude of the Price coefficient during Repeat, suggesting that while price discounts can be successful at prompting trial, they have lower returns during Repeat. Since the magnitudes of the coefficients are not necessarily directly comparable across Trial and Repeat, we explore this finding further by computing the elasticities with respect to changes in price for the three categories with significant price effects. The elasticities are calculated by using the posterior product and individual level coefficients, and simulating purchases under three pricing scenarios[8]: actual price; a 10% price reduction; a 20% price reduction. The Trial and Repeat elasticities for each category are obtained by averaging the corresponding product level elasticities over a total number of 1000 simulations. Figure 2 shows the relative magnitudes of the Trial and Repeat elasticities (i.e., the Repeat elasticities relative to Trial elasticities for each category).

[8] Different magnitudes of the price discount in the 5%-50% range yield qualitatively similar results.

Figure 2. Trial and Repeat Price Elasticities[+].

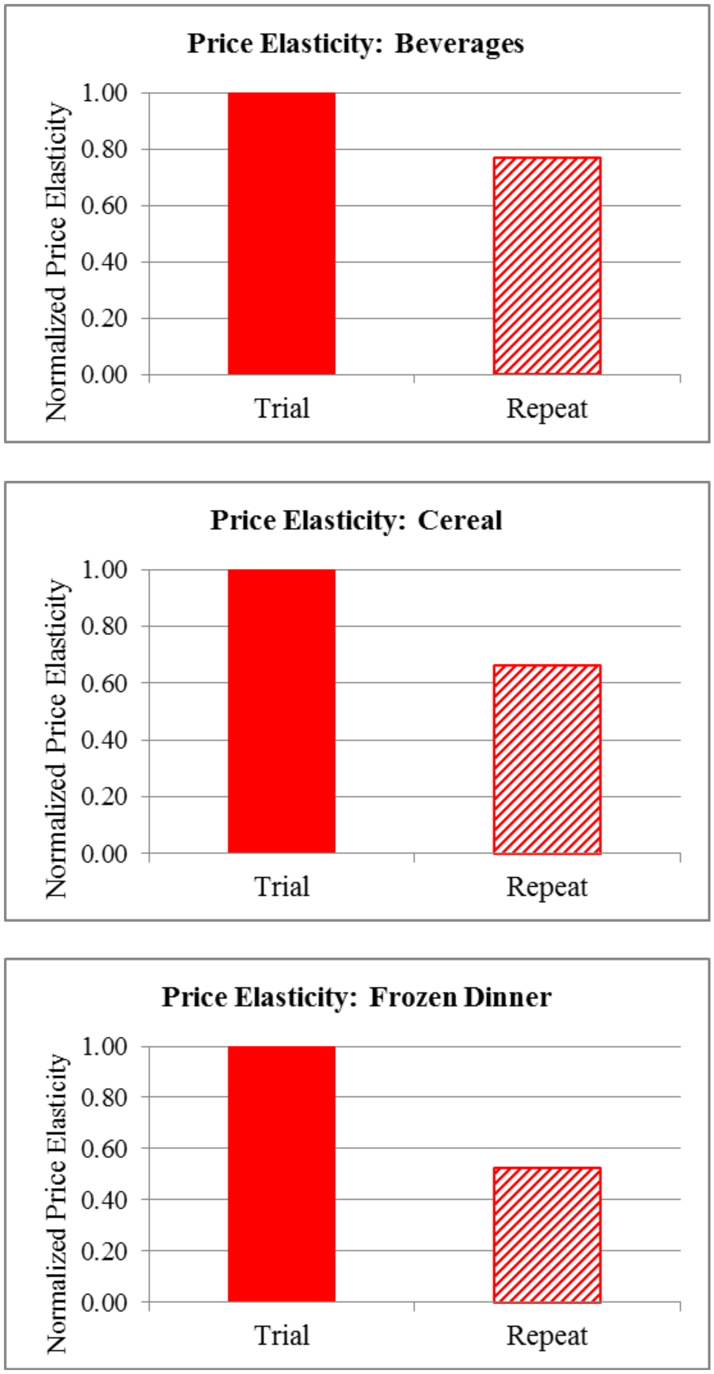

[+]Average simulated elasticities, computed for a 10% and 20% price reduction. Elasticities normalized to Trial elasticity for the corresponding product category.

Figure 2 supports the idea that indeed, trial purchases are more responsive to price reductions than repeat purchases. This result suggests that, at the trial stage, price discounts may help reduce the inherent risk associated with buying an unfamiliar product (i.e., compensate for uncertain preferences for the product). After the consumption experience, consumers learn about their preferences for the new product, which diminishes the relative importance of price for subsequent purchase decisions. This finding is of practical importance for marketing managers, suggesting that promotional tools should be more heavily employed during the early post launch periods, when inducing trial is the main concern; allocating these resources later, when most purchases are repeat purchases, may be less effective. Further, this result indicates that models that do not differentiate between Trial and Repeat purchases would result in biased estimates of price elasticities and hence potentially misguided marketing mix recommendations.

The model estimates also allow us to obtain a competitive map of the products for each category, i.e., find the location of each product in a perceptual trial and repeat space. Some new products are accurately evaluated by consumers even before trial, and the consumption experience does not significantly change consumers' perception of these products (i.e., the trial and repeat perceptual positions are very similar to each other). For other products the consumption experience is highly informative, causing a post-trial shift of the product's location in the perceptual space (i.e., the trial and repeat perceptual positions are quite different). Figures 3, 4, 5 and 6 illustrate both types of situations for several new products in each of the categories analyzed.

Figure 3a. Sample of Products with Similar Trial and Repeat Perceptual Positioning for the Beverages Category.

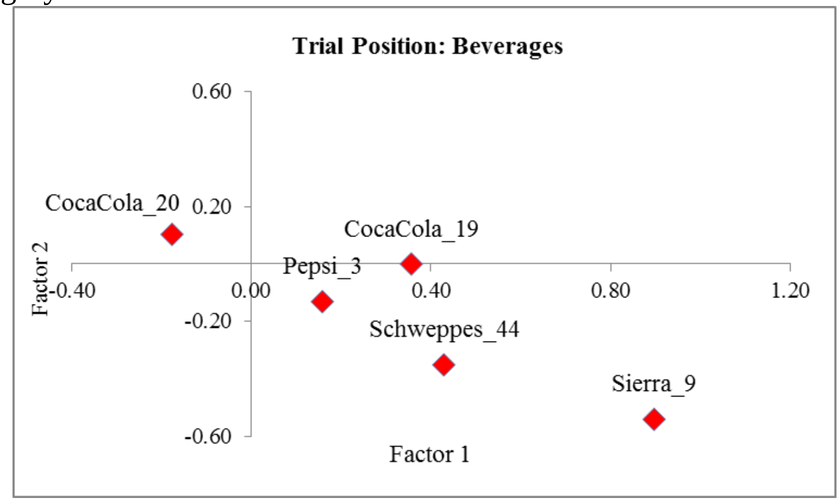

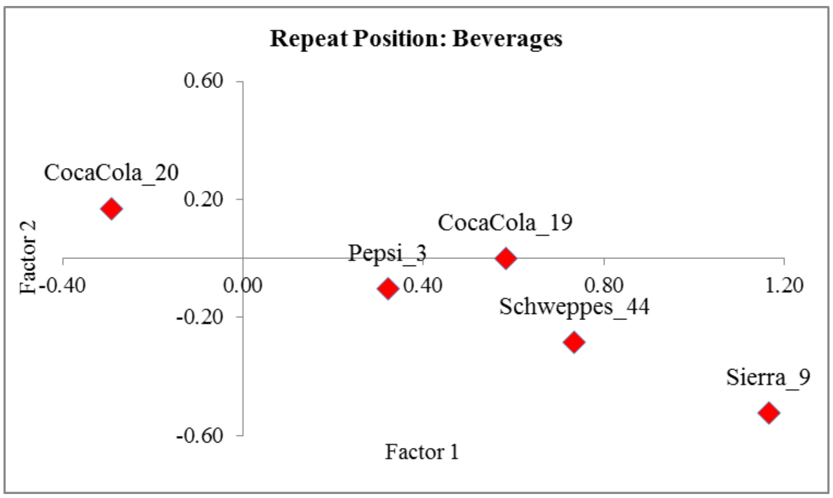

Figure 3b. Sample of Products with Different Trial and Repeat Perceptual Positioning for the Beverages Category.

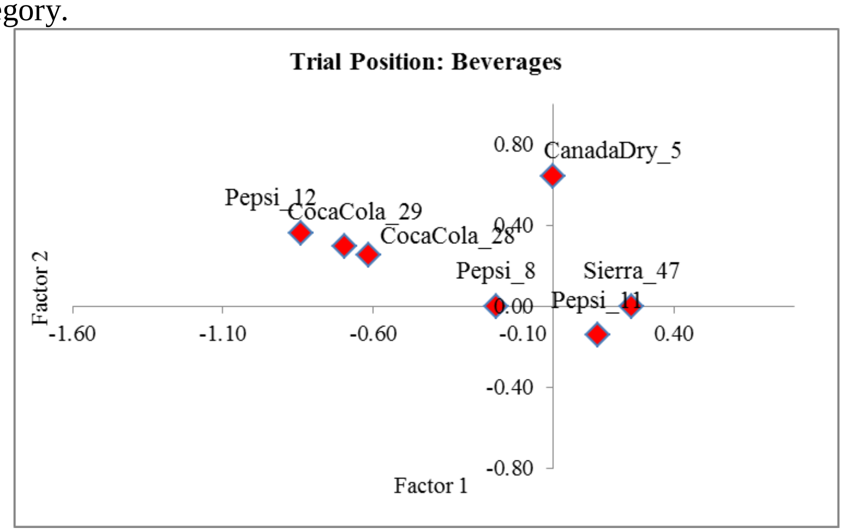

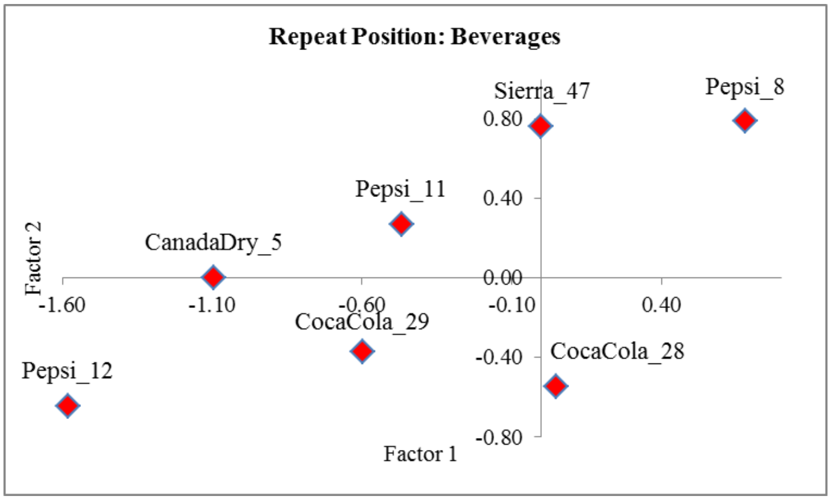

Figure 4a. Sample of Products with Similar Trial and Repeat Perceptual Positioning for the Cereal Category.

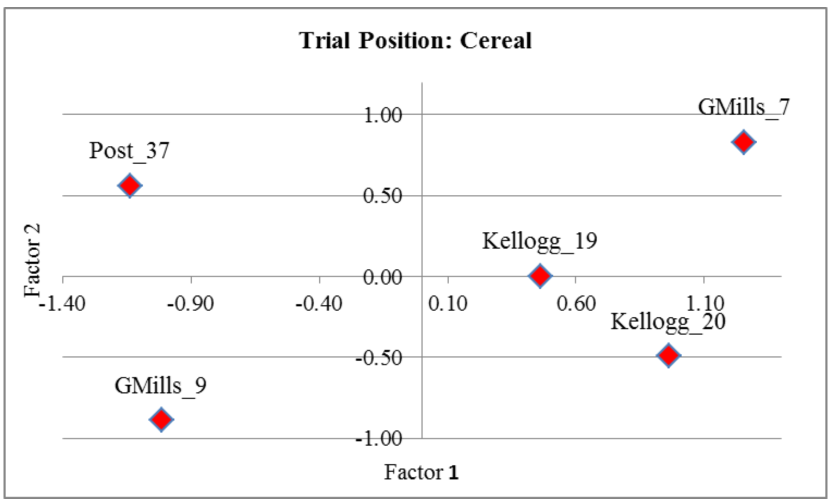

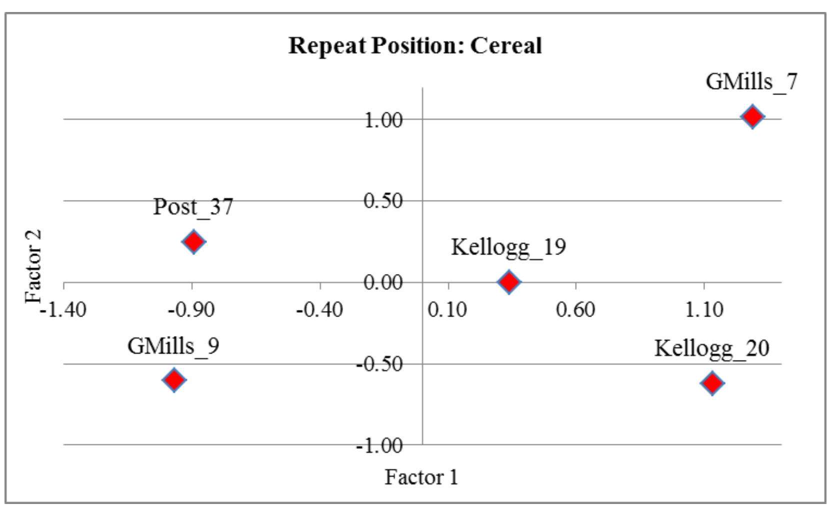

Figure 4b. Sample of Products with Different Trial and Repeat Perceptual Positioning for the Cereal Category.

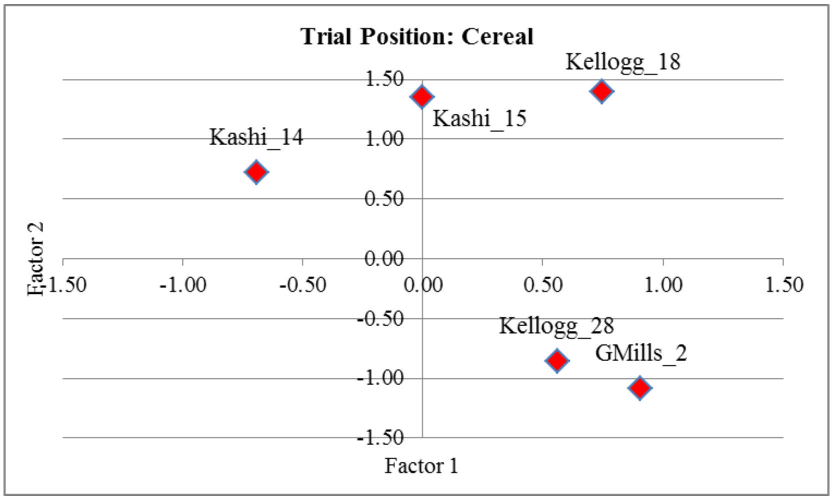

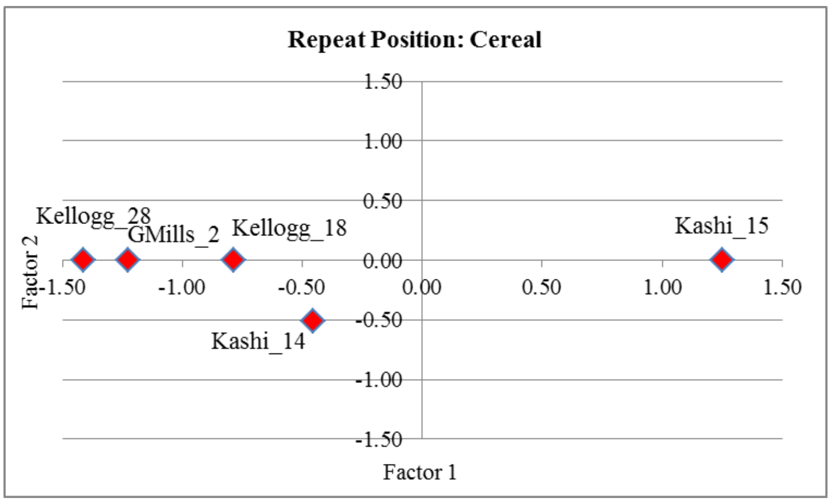

Figure 5a. Sample of Products with Similar Trial and Repeat Perceptual Positioning for the Frozen Dinner Category.

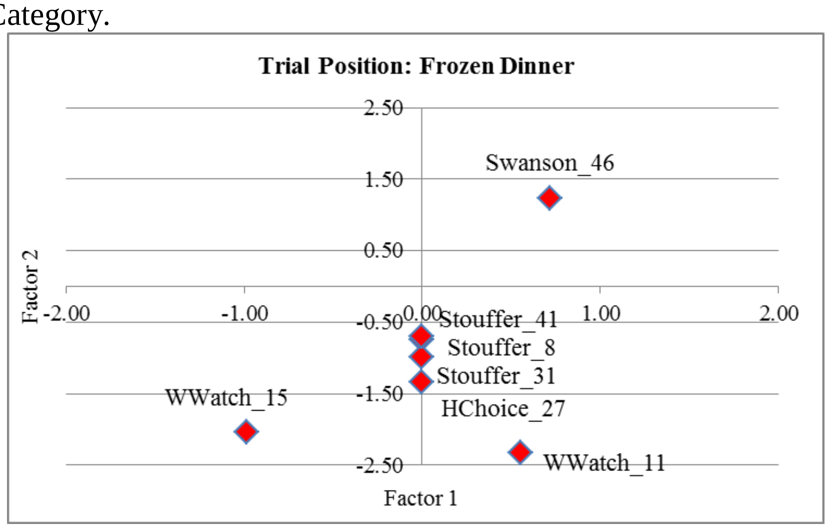

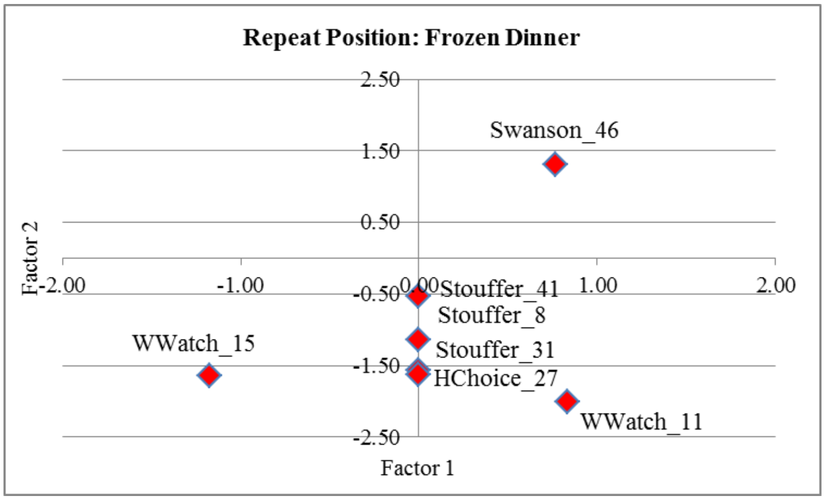

Figure 5b. Sample of Products with Different Trial and Repeat Perceptual Positioning for the Frozen Dinner Category.

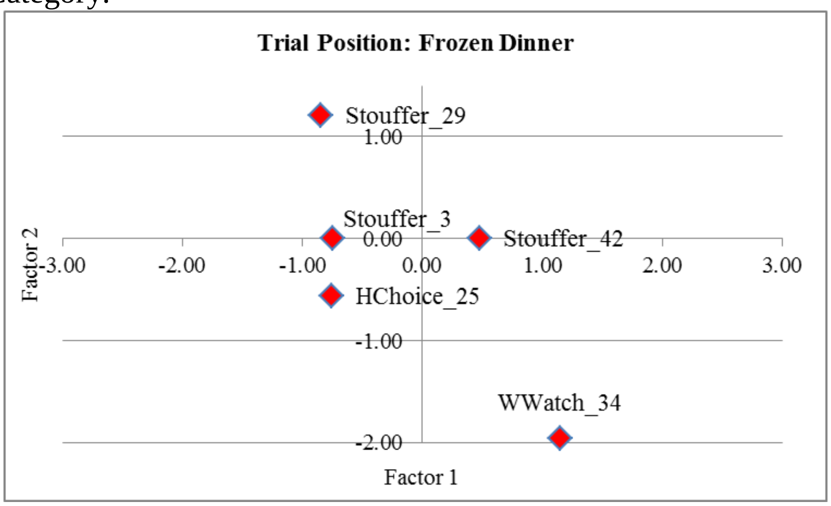

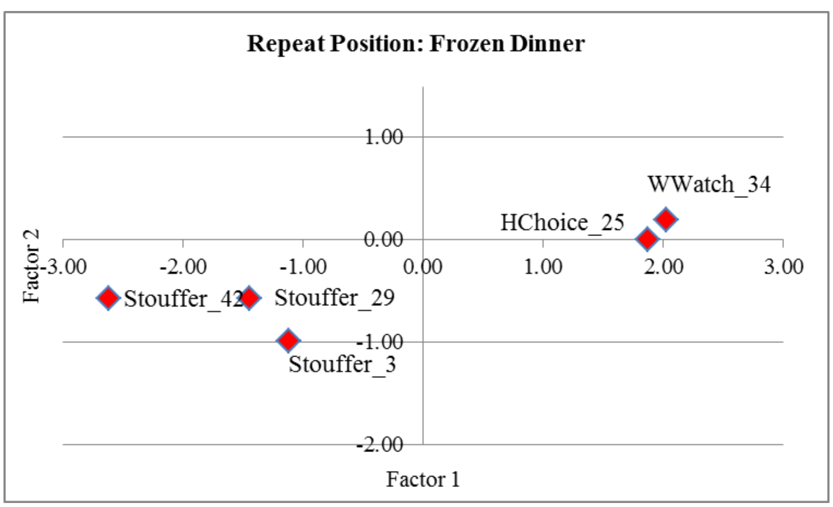

Figure 6a. Sample of Products with Similar Trial and Repeat Perceptual Positioning for the Snack Category.

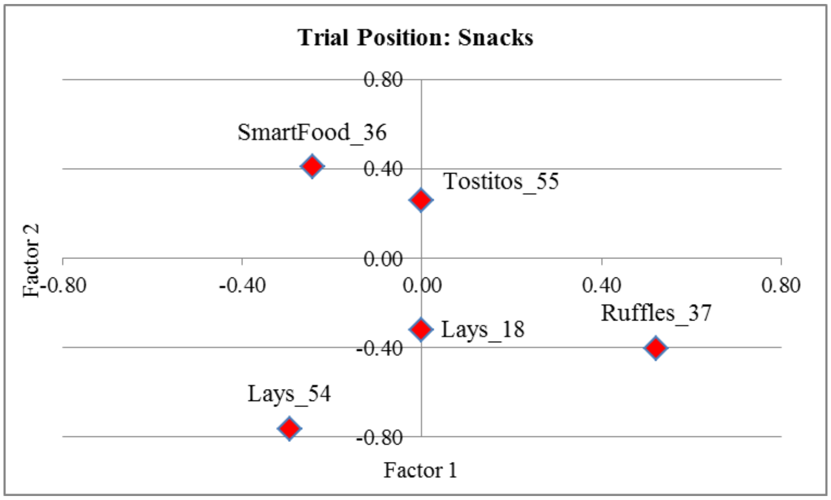

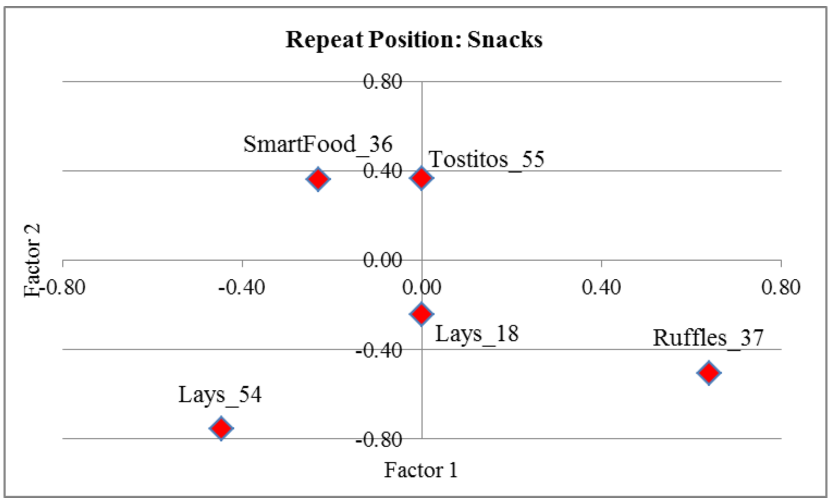

Figure 6b. Sample of Products with Different Trial and Repeat Perceptual Positioning for the Snack Category.

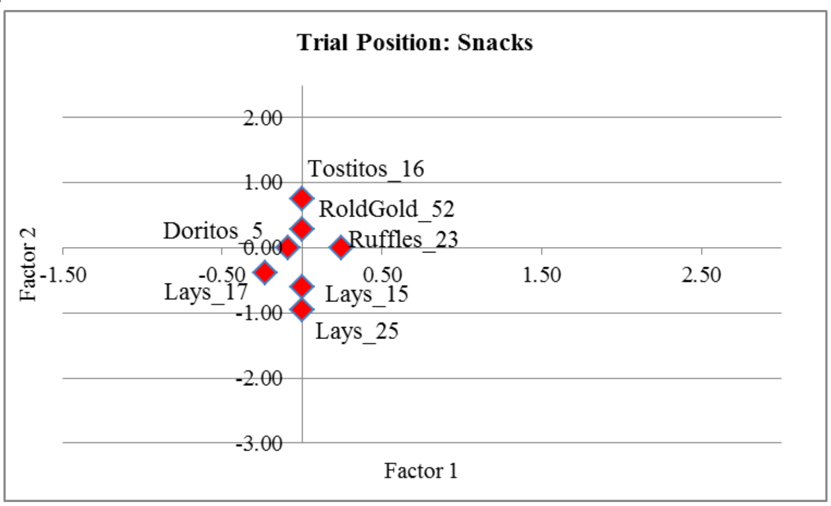

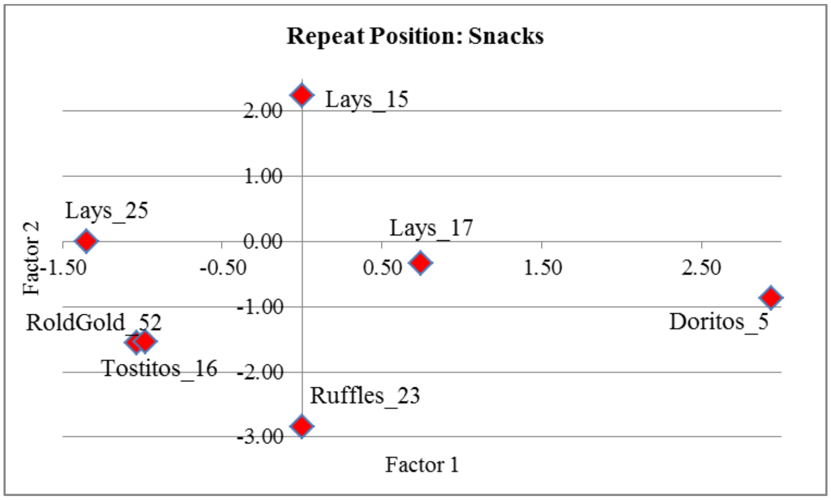

This type of information, which can be obtained very early after the new product launch, is very valuable for marketing managers. Identifying a shift in the product's perceptual position from the trial to the repeat stage provides an opportunity for increasing the product's success by better targeting. A significant difference between a product's trial and repeat locations cautions that the segment of consumers who would be most likely to repeatedly buy the product might be quite different from the segment inclined to actually try it. Managers could take actions[9] aimed at aligning the trial-stage positioning of the product with the post-consumption, repeat stage positioning and increase the product's appeal for the "likely repeat-buyer" segment. In addition, identifying the new product's location in the perceptual space also helps to identify the most relevant set of "competitor products": cross elasticities with respect to marketing actions are likely to be high among products as very similar to each other.

The individual-level model estimates also enable us to investigate potential regularities with respect to the responsiveness of early triers to marketing mix variables. The upper panel in

[9] Such actions might include improved marketing communications through advertising, package redesign, different shelf positioning, etc.

Table 7 shows the correlations between the posterior estimates of the intrinsic "early trier" scores[10] (i.e., z_{i3}^{TP} in equation 4) and posterior estimates of the individual level random effects for price and promotional sensitivity (i.e., the η_i^k terms in equation 3).

Table 7. Correlations of Early Trier/Repeat Scores with Marketing Mix Sensitivities

	Beverages	**Cereal**	**Frozen Dinner**	**Snacks**
Correlations with *Early Trier Score*				
Trial Price Sensitivity	- 0.07**	- 0.15***	- 0.15***	- 0.02
Trial Promotion Sensitivity	0.11***	0.11***	0.19***	0.12***
Repeat Price Sensitivity	- 0.07**	- 0.19***	- 0.11***	- 0.01
Repeat Promotion Sensitivity	0.03	0.11***	0.21***	0.11***
Correlations with *Repeat Score*				
Trial Price Sensitivity	0.05	0.06**	0.01	0.03
Trial Promotion Sensitivity	-0.04	-0.09***	0.03	0.04
Repeat Price Sensitivity	-0.07**	-0.12***	-0.11***	-0.11***
Repeat Promotion Sensitivity	0.01	0.10***	0.10***	0.11***

***p value <0.01; ** p-value <0.05.

While the magnitude of the correlations is small, indicating a relatively weak relation between the propensity to try early and sensitivity to marketing variables, the results are qualitatively similar across categories. For these types of products, early triers tend to have increased responsiveness to marketing variables at the trial stage (i.e., high price and promotional sensitivity). Further, early triers continue to react stronger to price discounts and in-store promotions during the repeat purchase stages. This result is consistent with findings reported in

[10] Another way of calculating an "early trier" score is to include the effect of individual characteristics (i.e. Category Usage and H Index) next to the intrinsic early trier estimates of z_{i3}^{TP}. For the product categories analyzed, the correlations between these two scores range from 0.50 for the Snack category, to 0.73 for the Beverage category (p < 0.01 in all cases) and the corresponding correlations with marketing variables are qualitatively similar.

other studies. For instance, Goldsmith and Flynn (1992) find that early adopters of new fashion products are more likely to pay attention to in-store displays and features than later adopters, while Montgomery (1971) finds a significant positive correlation between the propensity to be an early trier in the dental care category and the general propensity to purchase on price deals in the category. This result supports the idea that early triers may have a higher awareness of the shopping environment, i.e., they may be more likely to notice in-store displays and price fluctuations across shopping trips. In turn, more accurate assessment of price levels may result in higher price sensitivity, as some existing studies suggest (e.g., Thomas and Menon 2004 find that higher accessibility of memory-based price information results in stronger reactions to price increases; Sirvanci 2011 finds that higher price sensitivity is accompanied by more awareness of both price changes and of prices actually paid at the point of purchase). An implication of this result is that, at least for the product categories considered, early triers might not be the most profitable customers, since they will need further promotional incentives in order to repeatedly purchase a new product.

The lower panel in Table 7 shows the correlations between the posterior estimates of the "repeat scores" (i.e., z_{i3}^{RP} in equation 4) and posterior estimates of the individual level random effects for price and promotional sensitivity. All else equal, consumers with higher repeat scores are more likely to continue purchasing the new products tried. Unlike the early triers, these consumers do not exhibit a tendency of increased sensitivity to marketing mix during the trial stage. However, they tend to behave somewhat more price and promotion sensitive at the repeat stage. This pattern of results seems to suggest that consumers with high repeat scores have well defined product preferences (e.g., in terms of brand, flavor, ingredients), which sharply define the boundaries of their consideration sets. At the trial stage, this manifests as a more discriminating trial behavior i.e., only the new products fitting those preferences are tried, regardless of whether they are promoted at the time. At the repeat stage, the new products previously tried become part of a consideration set containing very similar (and relatively equally preferred) products; in turn, this increases the importance of marketing variables as "differentiating attributes".

Next, we employ the procedure described in Section 4 to generate forecasts for the out of sample products in the four product categories. Tables 8, 9, 10 and 11 show the forecasting results for end of year 1 sales. The model's forecasting accuracy for trial purchases is benchmarked against the "Exponential-Gamma model with covariates", which has been shown in the literature (Fader, Hardie and Zeithammer 2003) to produce consistently lower forecasting errors than other competing trial sales models when only short intervals of data are available (i.e., around 12 weeks). The forecasting performance of the model is also benchmarked against the "Erlang 2-Gamma model with covariates", which allows us to obtain a benchmark for both trial and total (trail and repeat) sales. This model has been found in previous research (Gupta 1991) to outperform other commonly used hazard models for interpurchase times, both in terms of in-sample fit and out-of-sample predictive performance[11].

[11] Gupta (1991) does not focus on new product forecasting; hence, his study does not evaluate forecasting performance when only a few weeks of data are available. The study assesses the competing models' forecasting accuracy for a 25-week horizon, based on 40 weeks of available data. However, for the considered horizon, the Erlang 2- Gamma model with covariates achieves a very good performance, with less than 3% forecasting error for total sales.

Table 8.Forecasting Errors (%) – Beverages.

Beverages	Weeks available	Exponential Gamma Trial	Proposed Model Trial	Proposed Model Total	Erlang2 Gamma Trial	Erlang2 Gamma Total
UPC 101						
	4	35.06	3.02	12.82	9.80	73.74
	8	69.39	4.94	3.85	20.76	87.73
	24	34.50	2.58	12.34	23.73	36.51
UPC 102						
	4	169.80	1.71	12.59	36.78	91.62
	8	10.92	6.61	1.28	62.85	121.67
	24	248.92	2.90	11.20	23.80	0.40
UPC 103						
	4	27.37	6.52	9.64	NA[+]	NA[+]
	8	129.98	2.38	11.58	NA[+]	NA[+]
	24	71.67	2.03	8.21	89.33	13.60
UPC 104						
	4	20.24	21.54	24.66	48.86	70.21
	8	20.82	21.82	15.26	55.45	68.07
	24	22.93	2.44	13.73	12.27	37.53
UPC 105						
	4	46.11	10.40	9.53	53.19	100.29
	8	64.05	6.81	13.14	61.23	96.55
	24	243.51	5.16	11.12	4.00	34.65
UPC 106						
	4	81.89	1.60	15.39	NA[+]	NA[+]
	8	10.91	5.30	9.96	NA[+]	NA[+]
	24	61.70	10.91	4.17	46.32	70.72
UPC 107						
	4	27.70	7.90	13.63	30.53	53.67
	8	39.24	3.86	8.23	24.23	55.90
	24	35.80	3.58	15.35	25.54	3.40
UPC 108						
	4	24.23	12.21	13.00	24.90	146.86
	8	34.24	4.81	16.28	24.34	131.82
	24	48.05	11.07	26.10	65.26	12.71
Average						
	4	54.05	8.11	13.90	34.01	89.39
	8	47.44	7.06	9.94	41.51	93.45
	24	95.88	5.08	12.77	36.28	26.19

[+]NA = converged parameter estimates could not be obtained.

Table 9. Forecasting Errors (%) - Cereals

Cereals	**Weeks available**	**Exponential Gamma Trial**	**Proposed Model Trial**	**Proposed Model Total**	**Erlang2 Gamma Trial**	**Erlang2 Gamma Total**
UPC 101						
	4	106.25	33.31	46.39	42.98	158.72
	8	71.44	2.97	16.75	39.80	162.14
	24	62.17	1.72	16.23	30.36	34.43
UPC 102						
	4	92.18	46.81	18.74	44.19	43.55
	8	1.60	22.12	16.73	36.55	94.98
	24	7.10	16.59	10.04	66.96	58.91
UPC 103						
	4	62.55	20.70	26.22	29.45	40.57
	8	75.04	21.04	19.42	25.56	146.80
	24	50.80	14.11	9.70	25.73	103.69
Average						
	4	86.99	33.60	30.45	38.87	80.94
	8	49.36	15.37	17.63	33.97	134.64
	24	40.02	10.80	11.99	41.01	65.67

Table 10. Forecasting Errors (%) – Frozen Dinners

Frozen Dinners	Weeks available	Exponential Gamma Trial	Proposed Model Trial	Proposed Model Total	Erlang2 Gamma Trial	Erlang2 Gamma Total
UPC 101						
	4	39.55	3.31	2.29	11.47	113.96
	8	73.60	0.37	4.71	15.01	118.68
	24	143.66	11.00	7.07	15.75	48.89
UPC 102						
	4	44.43	5.95	7.06	0.67	46.63
	8	2.34	2.56	10.46	15.51	38.64
	24	24.34	9.61	9.68	1.33	82.65
UPC 103						
	4	63.32	6.53	10.89	54.52	0.20
	8	102.64	3.91	14.47	15.33	102.13
	24	30.79	5.42	10.31	36.69	65.15
UPC 104						
	4	9.49	3.64	17.50	71.32	27.94
	8	59.21	3.90	25.07	63.43	3.74
	24	83.04	0.375	12.91	64.00	26.60
UPC 105						
	4	41.64	14.39	21.57	41.98	143.72
	8	85.17	2.41	8.93	14.02	50.18
	24	27.57	14.29	15.03	3.38	110.68
UPC 106						
	4	75.89	4.33	25.30	18.05	38.56
	8	54.40	2.84	9.43	23.13	63.76
	24	56.42	0.84	5.93	44.45	4.40
UPC 107						
	4	67.48	34.47	54.77	76.03	59.47
	8	39.93	2.61	16.26	55.95	23.05
	24	27.79	2.64	1.35	51.61	0.40
UPC 108						
	4	57.03	2.77	5.28	8.19	103.55
	8	76.53	12.86	10.18	2.43	107.94
	24	56.47	0.95	10.17	52.82	15.50
Average						
	4	49.85	9.42	18.08	35.32	66.75
	8	61.72	3.94	12.43	25.60	63.51
	24	56.26	5.64	9.05	33.75	44.28

Table 11. Forecasting Errors (%) - Snacks

Snacks	**Weeks available**	**Exponential Gamma Trial**	**Proposed Model Trial**	**Proposed Model Total**	**Erlang2 Gamma Trial**	**Erlang2 Gamma Total**
UPC 101						
	4	55.55	9.66	3.74	36.96	57.46
	8	20.19	9.30	13.99	10.60	24.60
	24	10.04	4.47	0.87	12.23	27.35
UPC 102						
	4	78.51	32.72	28.19	13.98	69.74
	8	46.57	10.04	6.04	16.13	62.61
	24	11.34	5.20	13.50	28.00	25.67
UPC 103						
	4	55.59	11.35	17.99	34.56	20.11
	8	70.85	4.09	22.51	14.42	36.57
	24	47.01	7.58	4.71	23.84	3.50
UPC 104						
	4	63.16	36.26	0.36	45.91	102.67
	8	34.75	4.87	25.50	3.49	43.22
	24	48.73	2.59	12.47	2.01	39.10
UPC 105						
	4	80.28	20.41	6.86	18.73	6.87
	8	40.69	7.07	4.73	8.38	34.42
	24	39.18	0.41	8.66	1.13	15.58
UPC 106						
	4	79.67	6.43	23.02	58.63	89.20
	8	67.48	3.08	13.94	10.67	10.44
	24	54.25	6.18	11.40	37.27	11.08
UPC 107						
	4	56.35	19.80	10.55	14.21	85.29
	8	7.04	2.90	13.82	16.75	96.90
	24	42.63	2.70	5.60	28.29	18.50
Average						
	4	67.01	19.51	12.95	31.85	61.62
	8	41.08	5.90	14.36	11.49	44.10
	24	36.16	4.16	8.17	18.96	20.11

The forecasts are performed assuming that t = 4, 8 and 24 weeks of post-launch data for the new product are available. The *Prediction Error Trial* column shows the absolute value of the forecasting errors (in percentages) for Trial purchases, cumulated over the interval [t+1 weeks, 52 weeks]. Similarly, the *Prediction Error Total* shows the forecasting errors for the total (trial and repeat) sales made by the end of year 1. More specifically:

$$\text{(8)} \quad PredError\ Trial_t = \left| \frac{\sum_{w=t+1}^{52} Pred\ trial\ sales\ week\ w - \sum_{w=t+1}^{52} Actual\ trial\ sales\ week\ w}{\sum_{w=t+1}^{52} Actual\ trial\ sales\ week\ w} \right| * 100$$

$$PredError\ Total_t = \left| \frac{\sum_{w=t+1}^{52} Pred\ total\ sales\ week\ w - \sum_{w=t+1}^{52} Actual\ total\ sales\ week\ w}{\sum_{w=t+1}^{52} Actual\ total\ sales\ week\ w} \right| * 100$$

For instance, for *UPC 101* in the Beverage category, when 4 weeks of data are available, the model's predictions deviate from the actual number of purchase during weeks 5-52 by 3.02% for trial sales and by 12.82% for total (i.e., trial and repeat) purchases.

Tables 8-11 suggest that the proposed model is quite good at producing rather accurate forecasts. Especially when only short periods of data are available, the forecasts generated have considerably smaller errors than the benchmark models. The poor performance of the benchmark models is due at least partly to parameter instability. For most products, when only 4 or 8 weeks of data are used, both benchmark models encounter severe estimation difficulties (e.g., the maximum likelihood procedure is very slow to converge; different sets of starting values reach different local maxima with quite different parameter values at convergence). This is a problem that can be expected to occur with most traditional approaches to forecasting when just a few weeks of purchases are observed. On the other hand, our proposed model circumvents these difficulties, since its ability to obtain parameter estimates for forecasting does not rely exclusively on the observed data for the focal product. The forecasting accuracy of the models seems to vary across categories, possibly reflecting cross-category differences in the "regularity" of new product purchase behavior. For instance, all three models show higher errors for the Cereal category, while the Erlang2-Gamma model performs particularly unsatisfactory for the 4 and 8 week intervals in the Beverage category.

The end of year 1 prediction errors based on 4 weeks of available data for each category are plotted in Figure 7 (trial sales) and Figure 8 (total sales). For trial sales, the 4-week forecasting errors by category are: Beverages: 8.11%, versus 54.05% (Exponential-Gamma) and 34.01% (Erlang2-Gamma); Cereal: 33.60% versus 89.99% (Exponential-Gamma) and 38.87% (Erlang2-Gamma); Frozen dinners: 9.42% versus 49.85% (Exponential-Gamma) and 35.32% (Erlang2-Gamma); Snacks: 19.51% versus 67.01% (Exponential-Gamma) and 31.85% (Erlang2-Gamma).

For the total sales, the 4-week forecasting errors are: Beverages: 13.90% versus 89.39% (Erlang2-Gamma); Cereal: 30.45% versus 80.94% (Erlang2-Gamma); Frozen dinners: 18.08% versus 66.75% (Erlang2-Gamma); Snacks: 12.95% versus 61.62% (Erlang2-Gamma).

Figure 7. Forecasting Errors for Trial Sales (4 weeks of data available).

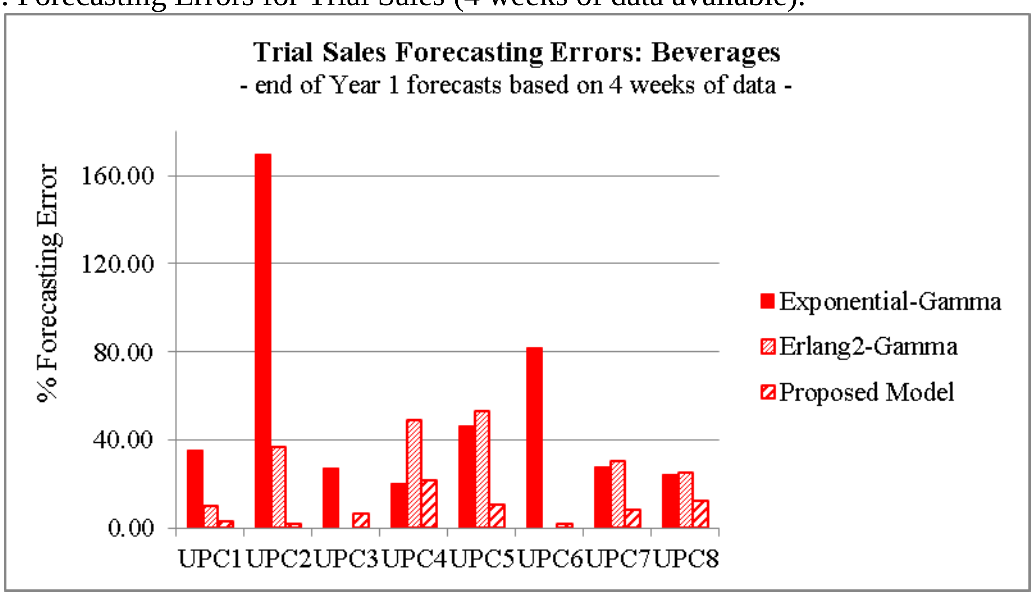

Erlang2-Gamma converged estimates could not be obtained for UPC3 and UPC6.

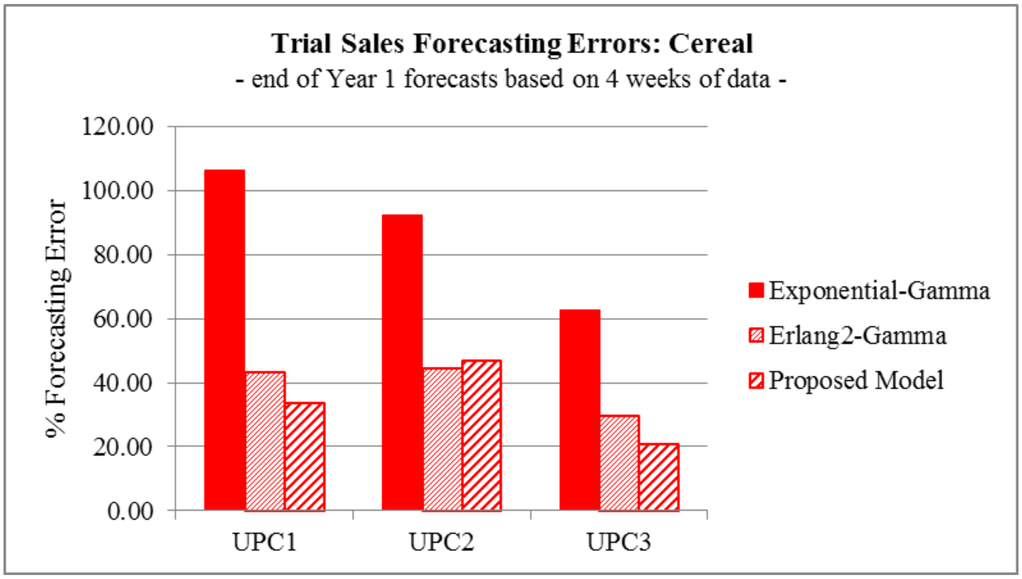
Trial Sales Forecasting Errors: Cereal
- end of Year 1 forecasts based on 4 weeks of data -
% Forecasting Error
120.00
100.00
80.00
60.00
40.00
20.00
0.00
UPC1
UPC2
UPC3
Exponential-Gamma
Erlang2-Gamma
Proposed Model

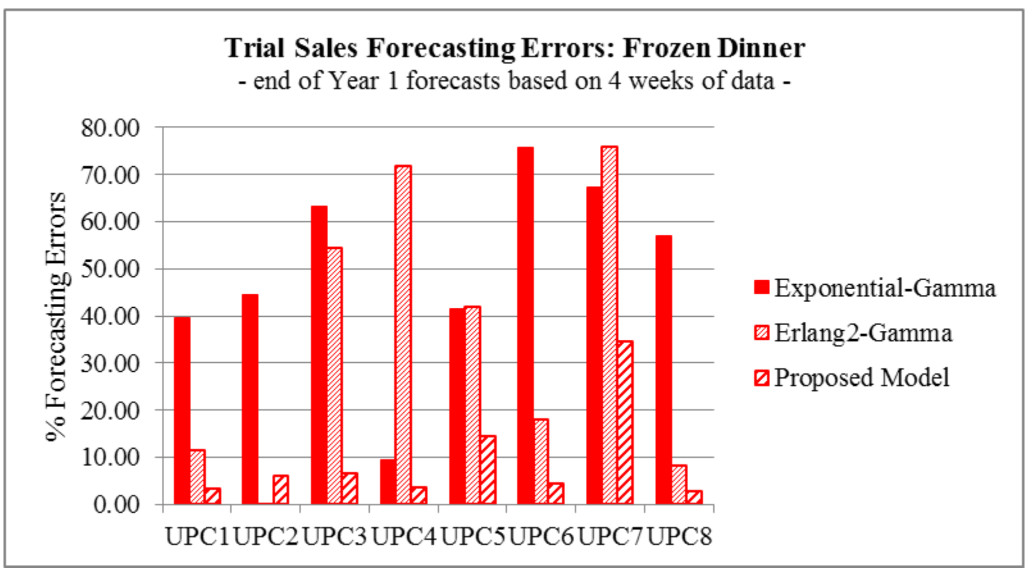
Trial Sales Forecasting Errors: Frozen Dinner
- end of Year 1 forecasts based on 4 weeks of data -
% Forecasting Errors
80.00
70.00
60.00
50.00
40.00
30.00
20.00
10.00
0.00
UPC1
UPC2
UPC3
UPC4
UPC5
UPC6
UPC7
UPC8
Exponential-Gamma
Erlang2-Gamma
Proposed Model

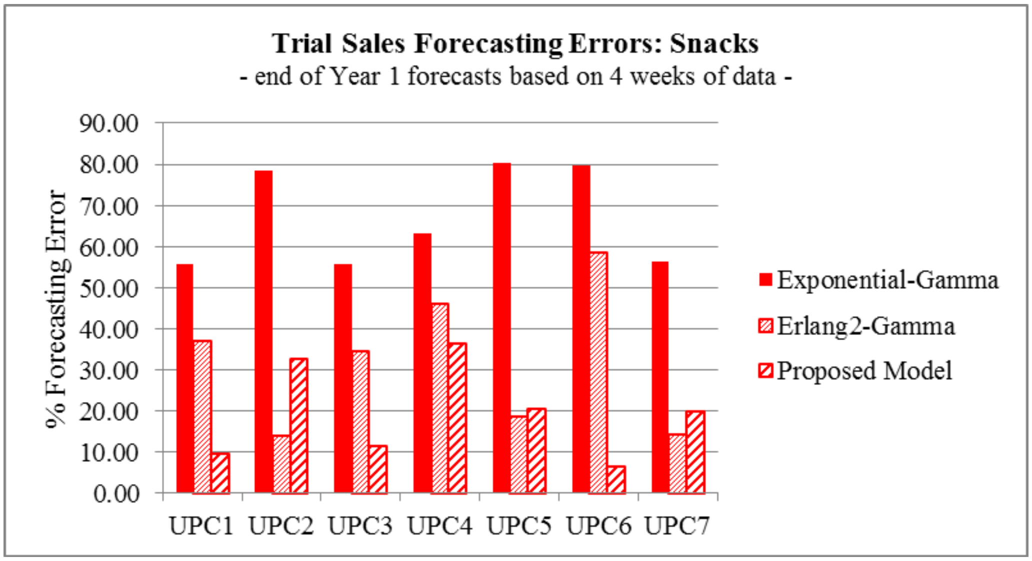

Figure 8. Forecasting Errors for Total Sales (4 weeks of data available).

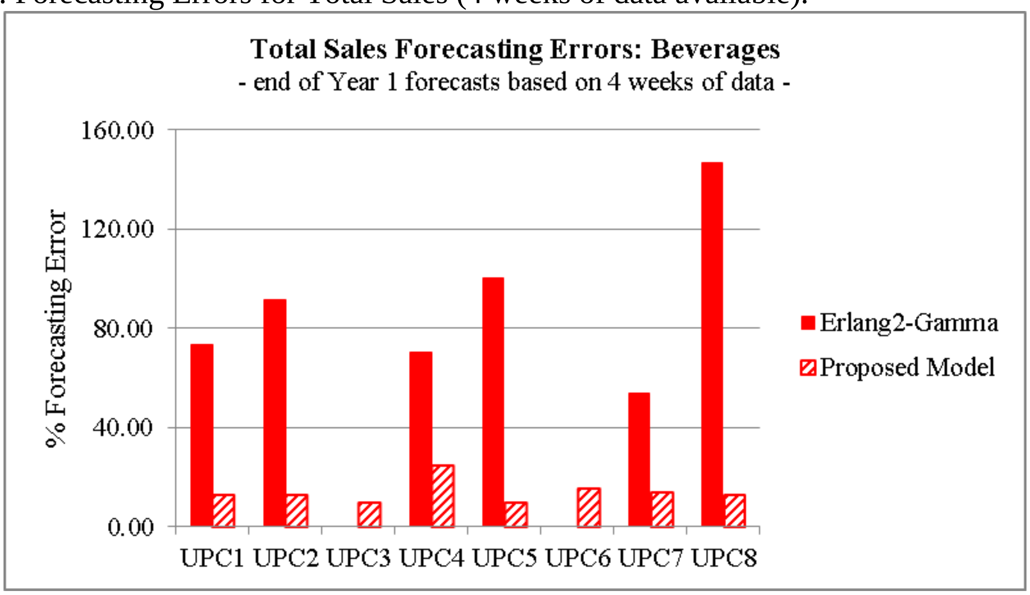

Erlang2-Gamma converged estimates could not be obtained for UPC3 and UPC6.

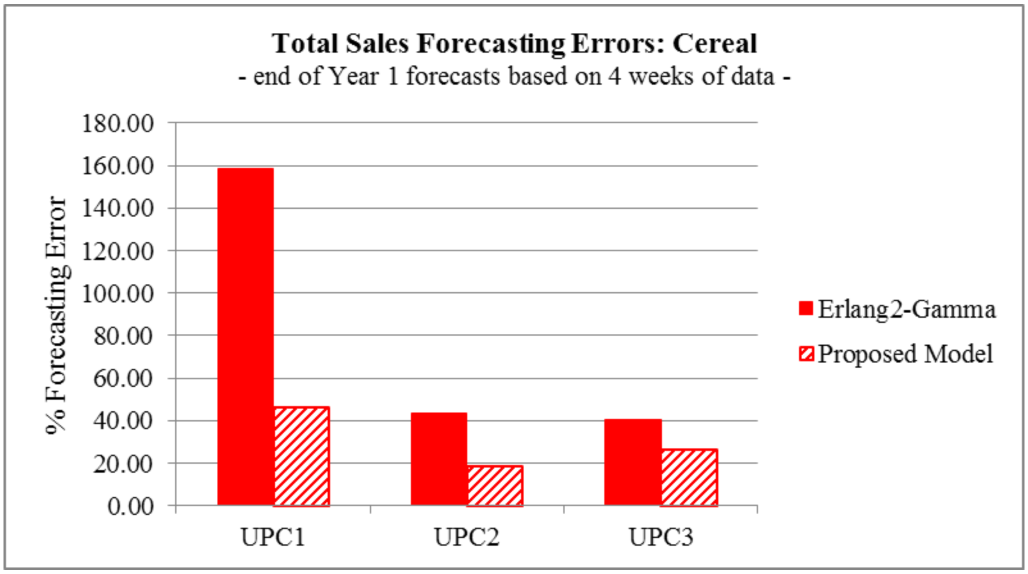
Total Sales Forecasting Errors: Cereal
- end of Year 1 forecasts based on 4 weeks of data -
% Forecasting Error
180.00
160.00
140.00
120.00
100.00
80.00
60.00
40.00
20.00
0.00
UPC1
UPC2
UPC3
Erlang2-Gamma
Proposed Model

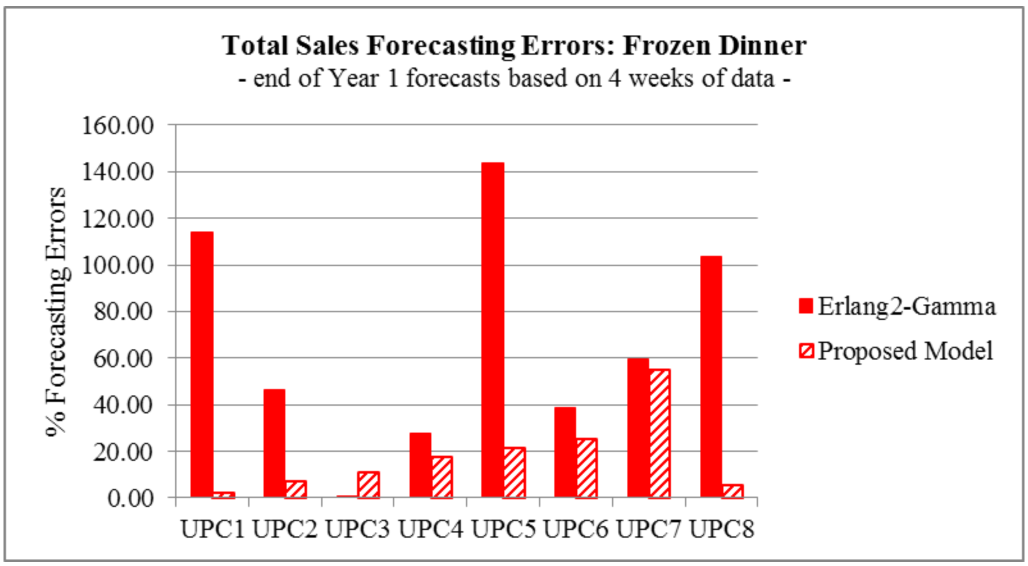
Total Sales Forecasting Errors: Frozen Dinner
- end of Year 1 forecasts based on 4 weeks of data -
% Forecasting Errors
160.00
140.00
120.00
100.00
80.00
60.00
40.00
20.00
0.00
UPC1
UPC2
UPC3
UPC4
UPC5
UPC6
UPC7
UPC8
Erlang2-Gamma
Proposed Model

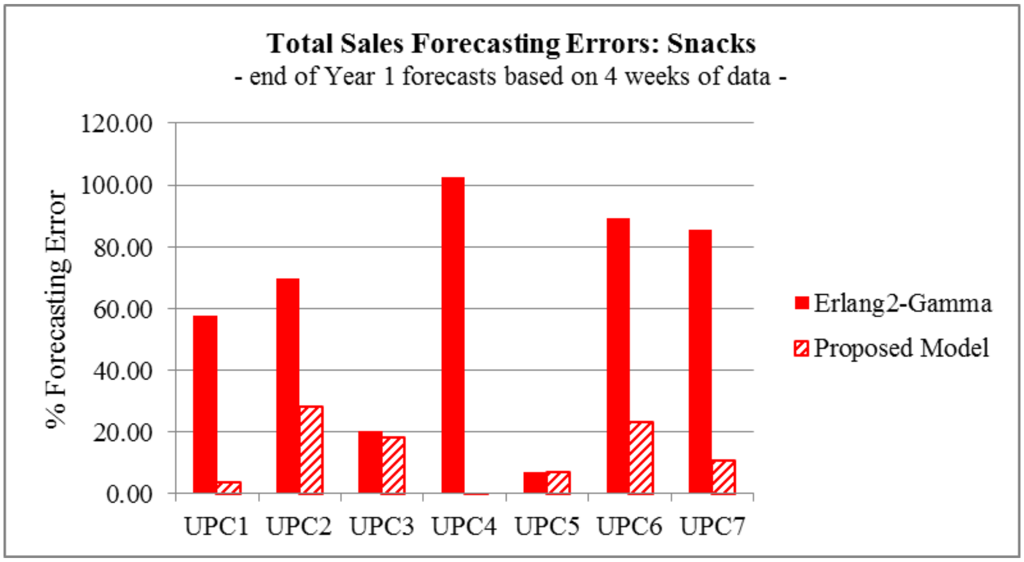

The results above highlight the benefits of our proposed approach. The differences in accuracy relative to the benchmark models tend to diminish when longer post-launch periods are observed (e.g., 24 weeks), but they are quite significant during the early post launch weeks. Due to cross-product/cross-individual information borrowing, the proposed model is able to generate more precise estimates of the parameters describing the sales of the new product, and therefore achieves greater accuracy much earlier. The results above suggest that the proposed model is a very promising approach for the purpose of obtaining early accurate new product sales forecasts.

Besides yielding accurate early new product sales forecasts, the proposed model yields insights that have useful managerial implications. The finding that many of the newly introduced products have different trial and repeat perceptual positions suggests that new product success might be improved by correcting consumers' (mis)perceptions of the product at the trial stage. For the purpose of illustrating this potential application, we perform a counterfactual simulation for one of the new products introduced in the Beverage category. The product is a diet cherry flavor cola beverage which uses a new sugar-derived artificial sweetener (as an alternative to aspartame). As shown in Figure 9, the product occupies quite different perceptual locations in the trial and repeat stage. This suggests that the consumption experience for this product is very informative and consumers' initial (pre-trial) assessment of the product is not very accurate.

Figure 9. Trial and Repeat Perceptual Positions for a New Diet Cherry Flavored Cola.

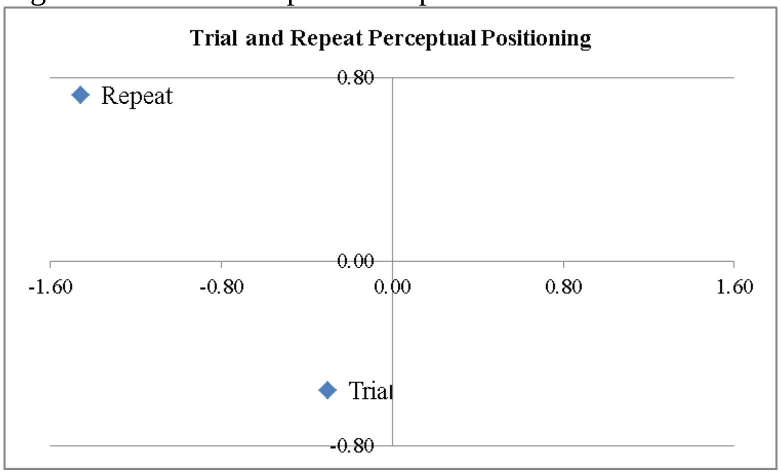

The different positions of the product in trial and repeat imply that the new product is appealing to different consumer segments at these stages. More specifically, the consumers most likely to repeat purchase the product (given trial) are not the consumers most likely to actually try the new beverage, as suggested by Figure 10. The left panel in Figure 10 displays the individual-level intrinsic propensities to repeat purchase the product[12], plotted in descending order. All else equal, the consumers inside the circled area have a higher than average propensity to repeatedly buy the product. The right panel in Figure 10 shows the corresponding individual-level trial propensities. Figure 10 suggests a mismatch between the trial and repeat tendencies of many consumers, which might prevent potentially frequent repeat buyers from even trying the new product.

[12] More specifically, the panel shows the consumer level deviations from the mean propensity to repeat buy the product (all else equal), which is computed as: $\Lambda^R_{F1} z_{i1} + \Lambda^R_{F2} z_{i2}$, where F is the focal product and i denotes the consumer.

Figure 10. Individual Level Preferences (Propensities) for Trial and Repeat.

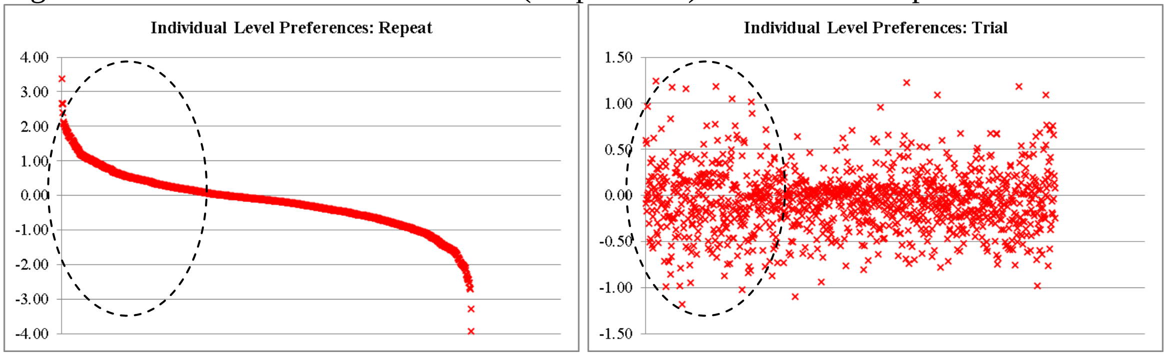

Marketing managers might be able to take actions aimed at shifting the product's location in the trial space closer to the product's location in the repeat space. Such actions could include designing advertising messages that provide more information about what the consumption experience might be like. For instance, in the case of the considered product, the observed trial stage positioning suggests that the new product is expected to be more similar to other cola based diet beverages; on the other hand, the product's repeat location is closer to non-diet products (i.e., sugar and high fructose corn syrup beverages), and not very close to other cola based soft drinks. In this case, advertising could emphasize that the new type of sweetener tastes "like real sugar", and that the dominant flavor is quite different from the typical cola soft drink. Other actions aimed at increasing the trial stage "accuracy" of the product's perception might include in-store activities, such as offering free samples, or changing the product's shelf placing (for instance, place it closer to the non-diet beverages and away from other cola soft drinks). The results of our simulations show that, if the trial stage positioning of the product can be shifted such that at the trial stage consumers perceive the product to be located at the "true" post-consumption (repeat) position, the new product success could be considerably improved. More specifically, the penetration rate would increase from the current 8.87% to 17.74%, and the total (trial and repeat) sales would increase at 198.16% relative to the sales at the current trial positioning. This type of trial repositioning ensures that consumers who would actually like the product (all else equal), would actually try it, which results in a high penetration rate for the segment of potential regular repeat buyers. The scenario in which trial and repeat positions exactly match is quite extreme and used here only for illustrative purposes. However, in practice, actions that shift the trial stage perception of the new product closer to the repeat position would increase the new product's chances of success[13].

Another result of the model that might be useful for generating managerially relevant recommendations is the finding that price elasticities are different at the trial and repeat stages, which suggests that the timing of sales promotions is important for new product success. The model estimates can be used for counterfactual simulations which can help with finding a more profitable schedule for sales promotions. To illustrate this use, we choose a new product in the Beverage category that has a relatively uniform distribution of the sales promotion weeks over the sample period. The product is sold at an average price of $3.87/unit and it is sold at a discounted price about 55% of the time. Approximately half of the promotional weeks occur relatively early after launch i.e., during a period when most purchases are trial purchases. More specifically, 54% of the promotional weeks occur during the first 39 post-launch weeks (Period 1) and 46% occur during weeks 40-79 (Period 2). Moreover, the posterior model estimates for this product indicate that, on average, consumers are more sensitive to the price during trial than during repeat (i.e. the product level estimates are directionally consistent with the category level average). For illustrative purposes we consider two different very simple, heuristic-based scenarios for the sales promotions schedule. Under Scenario 1 most of the promotional

[13] This holds assuming that the products appeal to enough consumers (i.e., is able to generate enough volume) at the repeat stage.

occasions[14] are allocated early post launch, i.e., 75% of the promotional weeks occur during Period 1. Under Scenario 2, the reverse situation is considered, with only 25% of the promotional weeks occurring during Period 1[15].

Table 12.Unit Sales and Revenues under Different Price Promotion Schedules.

	% Promotional Weeks Period 1	**% Promotional Weeks Period 2**	**Total Unit Sales Relative to Baseline**	**Total Revenue Relative to Baseline**
Scenario 1	75%	25%	109.86%	110.95%
Scenario 2	25%	75%	89.87%	89.50%

Table 12 shows the results of the simulations for sales under the two scenarios, relative to the baseline (actual) promotion schedule. With more frequent early period promotions under Scenario 1, the unit sales of the product over the considered 79 week period would increase by 9.86%. Further, the unit sales increase is accompanied by a comparable increase in sales revenue, which suggests that the increase in volume does not necessarily happen only during discount occasions. By contrast, scheduling less frequent promotions during the early post launch period would actually hurt the product's success, with only 89.87% of total unit sales relative to the

[14] For both scenarios the vector of actual (observed in sample) prices is used. The only difference across the Baseline scenario, Scenario 1 and Scenario 2 is the timing of sales (i.e. the week during which a discounted price is used).

[15] It is also possible to derive the "optimal" promotional schedule. However, with no closed form solution, this involves simulations for a complex (high dimensionality) dynamic optimization problem, which is beyond the scope of this study.

baseline. While these scenarios are very simple, heuristic-based counterfactuals are easy to implement in practice and may be used early post-launch to improve the new product's success.

Overall, our results show that the proposed model is able to achieve a good forecasting accuracy and generates insights that could be translated into actionable managerial recommendations.

7. SUMMARY AND DISCUSSION

In this study, we propose a new approach to new product sales forecasting, which we apply to data for the CPG industry. The model aims at improving prediction accuracy by leveraging the information available from multiple past new product introductions within the category, as well as the information available from the observed past behavior of consumers in the target market. The model allows for individual level heterogeneity in purchase behavior, and for differences in responsiveness to marketing mix variables between the trial and repeat stages.

The results indicate that the proposed model outperforms commonly used benchmark models for new product forecasting and it is able to achieve a very good predictive performance with only a few weeks of post-launch data available. The model also allows generating some practical recommendations that could help marketing managers increase the new product's chances of success. Our results show that, for three out of four product categories considered, trial price elasticities are higher than repeat price elasticities, which suggests that promotional budgets should generally be employed more heavily early post launch. In addition, our model yields product-specific estimates of price and promotional elasticities shortly after the new product is introduced; this information could be useful for customizing promotional budget allocation to the focal new product. The model also allows pinpointing the new product's positioning in a perceptual trial and repeat space. This information makes it possible to identify the segment of consumers most likely to repeatedly buy the product, as well as potential "misperceptions" of the product by the target market in the pre-trial stage, which might be correctable by marketing actions. Further, this information allows identifying early post launch the set of most relevant "competitor products" within the category, which could be useful for designing marketing communications (e.g., comparative advertising), or timing promotional activities (given the expected high promotional cross-elasticities with products in this set). The proposed model also allows for investigating potential behavioral differences between early and late triers. For the product categories considered, we find that consumers with an intrinsic propensity to be early triers tend to be more sensitive to marketing variables than later triers, at both trial and repeat stages, which might suggest that early triers are more aware of the in-store informational environment (e.g. price levels, in-store promotional displays or featured products).

We applied the model in the context of the CPG industry. However, the proposed approach is generalizable to any context where individual level data for multiple historical products is available, such as CRM databases. While the absolute forecasting performance of the model will likely depend on the degree of similarity between the focal new product and the historical products, our information borrowing approach can be expected to outperform other existing approaches to forecasting. The model's implications for behavioral questions, such as the relative magnitude of trial and repeat elasticities, or marketing mix responsiveness profile of early triers can be expected to depend on different factors, such as the degree of product newness (e.g. Donnely and Etzel 1973), or various product category characteristics.

REFERENCES

Bass, F. M. (1969) “A New Product Growth Model for Consumer Durables”, *Management Science*, 15 (5), pp. 215-227.

Bass, F.M, T.V. Krishnan and D. C. Jain (1994), “Why the Bass Model Fits without Decision Variables”, *Marketing Science*, 13 (3), pp. 203-223.

Chandrashekaran, M and R. K. Sinha (1995) “Isolating the Determinants of Innovativeness: A Split-Population Tobit (SPOT) Duration Model of Timing and Volume of First and Repeat Purchase”, *Journal of Marketing Research*, 32 (4), pp. 444-456.

Du, R.Y. and W.A. Kamakura (2011), “Measuring Contagion in the Diffusion of Consumer Packaged Goods”, *Journal of Marketing Research*, 48 (1), pp. 28-47.

Donnelly, J. H. and M. J. Etzel (1973), “Degrees of Product Newness and Early Trial”, *Journal of Marketing Research*, 10 (3), pp. 295-300.

Fader, P.S., B.G.S. Hardie and M. Wisniewski (1998), “An Empirical Comparison of New Product Trial Forecasting Models”, *Journal of Forecasting*, 17 (34), pp. 209-229.

Fader, P.S. and B.G.S. Hardie (2003), “Can We Infer Trial and Repeat Numbers from Aggregate Sales Data?”, *Working Paper, University of Pennsylvania*.

Fader, P.S., B.G.S. Hardie and R. Zeithammer (2003), “Forecasting New Product Trial in a Controlled Test Market Environment”, *Journal of Forecasting*, 22 (5), pp. 391-410.

Fader, P.S., B.G.S. Hardie and C.Y. Huang (2004), “A Dynamic Changepoint Model for New Product Sales Forecasting”, *Marketing Science*, 23 (1), pp. 50-65.

Foxall, G. R. (1995), "Cognitive Styles of Consumer Initiators", *Technovation,* 15 (5), pp. 269-288.

Golder, P. N. and G. J. Tellis (1997) “Will It Ever Fly? Modeling the Takeoff of Really New Consumer Durables”, *Marketing Science*, 16 (3), pp. 256-270.

Goldsmith, R. and L. R. Flynn (1992), “Identifying Innovators in Consumer Product Markets”, *European Journal of Marketing*, 26 (12), pp. 42-55.

Gupta, S. (1991), “Stochastic Models of Interpurchase Times with Time-Dependent Covariates”, *Journal of Marketing Research*, 28 (1), pp. 1-15.

Hahn, M., S. Park, L. Krishnamurthi and A. A. Zoltners (1994), “Analysis of New Product Diffusion Using a Four-Segment Trial-Repeat Model”, *Marketing Science,* 13 (3), pp. 224-247.

Iyengar, R., C. Van den Bulte and T. W. Valente (2011), “Opinion Leadership and Social Contagion in New Product Diffusion”, *Marketing Science*, 30 (2), pp. 195-212.

Im, S., B.L. Bayus and C.H. Mason (2003), “An Empirical Study of Innate Consumer Innovativeness, Personal Characteristics, and New-Product Adoption Behavior”, *Journal of the Academy of Marketing Science*, 31(1), pp. 61-73.

Jain, D. C. and N. J. Vilcassim (1991), “Investigating Household Purchase Timing Decisions: A Conditional Hazard Function Approach”, *Marketing Science,* 10 (1), pp. 1-23.

Kamakura, W. A. and S.K. Balasubramanian (1987), “Long-Term Forecasting with Innovation Diffusion Models: The Impact of Replacement Purchases”, *Journal of Forecasting,* 6 (1), pp. 1-19.

Kamakura, W. A. and S.K. Balasubramanian (1988), “Long-Term View of the Diffusion of Durables”, *International Journal of Research in Marketing*, 5 (1), pp. 1-13.

Kamakura, W. A., B. Kossar and M. Wedel (2004), "Identifying Innovators for the Cross-Selling of New Products", *Management Science,* 50 (8), pp. 1120-1133.

Lee, J., P. Boatwright and W.A. Kamakura (2003), "A Bayesian Model for Prelaunch Sales Forecasting of Recorded Music", *Management Science*, 49 (2), pp. 179-196.

Massy, W. F. (1969), "Forecasting the Demand for New Convenience Products", *Journal of Marketing Research*, 6 (4), pp. 405-412.

Moe, W. and P.S. Fader (2001), "Modeling Hedonic Portfolio Products: A Joint Segmentation Analysis of Music Compact Disk Sales", *Journal of Marketing Research,* 38 (3), pp. 376-385.

Montgomery, D. B. (1971), "Consumer Characteristics Associated with Dealing: an Empirical Example", *Journal of Marketing Research,* 8 (1), pp.118-120.

Morgan, F.W. Jr. (1979), "Are Early Triers Heavy Users?", *Journal of Business*, 52 (3), pp. 429-434.

Olson, J. and S. Choi (1985), "A Product Diffusion Model Incorporating Repeat Purchases", *Technological Forecasting and Social Change*, 27 (4), pp. 385-398.

Rao, A. G. and M. Yamada (1988), "Forecasting with a Repeat Purchase Diffusion Model", *Management Science,* 34 (6), pp. 734-752.

Schneider, J. and J. Hall (2011), "Why Most Product Launches Fail", *Harvard Business Review*, 89 (4), pp. 21-23.

Schweidel, D.A. and P. Fader (2009), "Dynamic Changepoints Revisited: An Evolving Process Model of New Product Sales", *International Journal of Research in Marketing*, 26 (2), pp. 119-124.

Seetharaman, P. B. (2004), "The Additive Risk Model for Purchase Timing", *Marketing Science,* 23 (2), pp. 234-242.

Silk, A. J. and G. L. Urban (1978), "Pre-Test-Market Evaluation of New Packaged Goods: A Model and Measurement Methodology", *Journal of Marketing Research*, 15 (3), pp. 171-179.

Sinha, R. K. and M. Chandrashekaran (1992) "A Split Hazard Model for Analyzing the Diffusion of Innovations", *Journal of Marketing Research,* 29 (1), pp.116-127.

Sirvanci, M. B. (2011), "An Empirical Study of Price Thresholds and Price Sensitivity", Journal of Applied Business Research, 9 (2), pp. 43-49.

Steenkamp, Jan-Benedict E.M. and K. Gielens (2003), "Consumer and Market Drivers of the Trial Probability of New Consumer Packaged Goods", *Journal of Consumer Research*, 30 (4), pp. 368-384.

Taylor, J.W. (1977), "A Striking Characteristic of Innovators", *Journal of Marketing Research*, 14 (1), pp.104-107.

Thomas, M. and G. Menon (2004), "Memory-Based Comparison Standards: Decomposing the Effects of Accessibility and Magnitude", *Working Paper, New York University*.

Train, K. (2003), "Discrete Choice Methods with Simulation", *Cambridge University Press.*
Van den Bulte, C. (2000), "New Product Diffusion Acceleration: Measurement and Analysis", *Marketing Science*, 19 (4), pp.366-380.

U.S. Census Bureau (2000), Allcountries.org, Table 889, "New Product Introductions of Consumer Packaged Goods", *http://www.allcountries.org/uscensus/889_new_product_ introductions_of_consumer_packaged.html*.

Wood, S. L. and J. Swait (2002), “Psychological Indicators of Innovation Adoption: Cross-Classification Based on Need for Cognition and Need for Change”, *Journal of Consumer Psychology*, 12 (1), pp. 1-13.

APPENDIX A. OUTLINE OF MODEL ESTIMATION

Let us start by assuming a model with no individual level heterogeneity.

Denote β_j = (M_J-dimensional) vector of product-level coefficients for product *j*; $j = 1,..J;$

Let $\beta_j \sim N\ (b, W^{UPC})$ with priors[16]:

$k(b,W^{UPC}) = k(b)k(W^{UPC})$

k(b) is $N(b_0, S_{0b})$, diffuse

$k(W^{UPC})$ is inverted Wishart $IW(M_J, I)$

The conditional posteriors have the form:

1) $K(\beta_j\ |b, W^{UPC}) \propto \prod_n L(y_{nj}|\beta_j)\ \phi(\beta_j|b, W^{UPC})$

2) $K(b^r|W^{UPC\ r-1}, \beta_j^{r-1} \forall j)$ is given by $N(\frac{\sum_j \beta_j^{r-1}}{J},\ \frac{W^{UPC\ r-1}}{J})$

3) $K\left(W^{UPC}\middle|b, \beta_j\ \forall j\right)$ is inverted Wishart: $IW\left(M_J + J,\ \frac{(M_J I + J\overline{S_J})}{(M_J + J)}\right)$ where $\overline{S_J} = \frac{\sum_j(\beta_j - b)(\beta_j - b)'}{J}$

where r is the current iteration of the algorithm and (with some abuse of notation) $\prod_n L(y_{nj}|\beta_j)$ denotes the likelihood function pertaining to product *j* for all consumers in the sample.

[16] Note that this is conceptually equivalent to specifying mean-zero product effects $\beta_j \sim N\ (0, W^{UPC})$ together with a separate vector of "fixed" parameters *b* with diffuse prior $N(b_0, S_{0b})$, where b_0 and S_{0b} would be treated as given and the parameters *b* would be updated with respect to the likelihood for the entire data (i.e., all products included).

Including individual-level heterogeneity adds several more parameters to be estimated:

β_n = (M_N-dimensional) vector of individual-level coefficients for consumer n; $n = 1,..N;$

α = vector of "fixed" coefficients (includes the latent factors for each product);

with $\beta_n \sim N\ (0, W^N)$ with priors $k(W^N)$ inverted Wishart $IW(M_N, I)$ and where $k(\alpha)$ is $N(\alpha_0, S_{0\alpha})$, diffuse.

The corresponding conditional posteriors are then:

1) $K(\beta_j\ | b, W^{UPC}) \propto \prod_i L(y_{ij}|\beta_j, \beta_n, \alpha, Z_n)\ \phi(\beta_j|b, W^{UPC})$

2) $K(b^r|W^{UPC\ r-1}, \beta_j^{r-1}\forall j)$ is given by $N(\frac{\sum_j \beta_j^{r-1}}{J}, \frac{W^{UPC\ r-1}}{J})$

3) $K(W^{UPC}|b, \beta_j\ \forall j)$ is inverted Wishart: $IW\left(M_J + J, \frac{(M_J I + J\bar{S}_J)}{(M_J + J)}\right)$

4) $K(\beta_n\ |W^N) \propto L(y_n|\beta_j, \beta_n, \alpha, Z_n)\phi(\beta_n|0, W^N)$

5) $K(\alpha|\beta_j, \beta_n, \alpha, Z_n) \propto \prod_n L(y_n|\beta_j, \beta_n, \alpha, Z_n)$

6) $K(W^N|\beta_n\ \forall n)$ is inverted Wishart: $IW\left(M_N + N, \frac{(M_N I + N\bar{S}_N)}{(M_N + N)}\right)$

where Z_n is the 4-dimensional vector of factor weights for individual n, $L(y_{ij}|\beta_j, \beta_n, \alpha, Z_n)$ denotes the individual-level likelihood contribution and $\prod_n L(y_n|\beta_j, \beta_n, \alpha, Z_n)$ denotes the likelihood for the entire sample.