Integrated Factor Models of Variety Seeking Dynamics

by

Adam N. Smith
University College London
a.smith@ucl.ac.uk

Yinxing Li
Tohoku University
yinxing.li.a8@tohoku.ac.jp

Shohei Hasegawa
Hosei University
shasegawa@hosei.ac.jp

Nobuhiko Terui
Tohoku University
terui@tohoku.ac.jp

Greg M. Allenby
Ohio State University
allenby.1@osu.edu

August 26, 2020

Integrated Factor Models of Variety Seeking Dynamics

Abstract

Consumers commonly purchase multiple product varieties both in a given shopping trip and across trips over time. However, existing models of variety seeking usually address one but not both of these dimensions. In this paper, we propose an integrated framework for modeling variety seeking dynamics. We simultaneously account for horizontal variety through a multiple discrete/continuous demand model and temporal variety by imposing dynamic factor structures on utility parameters. These factor models allow high-dimensional sets of utility parameters to be explained by a few common latent factors that can evolve over time. The proposed dynamic models are applied to household-level scanner panel data in the yogurt category where we find distinct shifts in demand patterns over time. We also find that consumers value variety much more than indicated by traditional models. Finally, we show that being able to distinguish between cross-sectional and temporal heterogeneity in preferences is important for measuring the value of an assortment over time.

Keywords: Factor model, multiple discreteness, threshold switching, satiation, heterogeneity, compensating value

1 Introduction

Consumers are commonly observed to purchase a variety of goods in every product category. This variety can arise from multiple goods being purchased at the same time (i.e., horizontal variety) or from multiple goods being purchased over a longer time horizon (i.e., temporal variety). Explaining both dimensions of the demand for variety has been the subject of a long literature in marketing and economics, with reasons including changes in prices and income (Stigler and Becker, 1977), the presence of multiple needs (McAlister and Pessemier, 1982), attribute-based satiation (McAlister, 1982; Trivedi et al., 1994; Kim et al., 2002, 2007; Hasegawa et al., 2012), changes in assortment structure (Kahn and Wansink, 2004; Hong et al., 2016), time of day and mood (Kahn, 1995; Gullo et al., 2019), preference uncertainty and learning (Che et al., 2015), and changes in costs of search (Ershov, 2020). Together, these studies suggest that changes in both internal motivations and external market conditions play important roles in shaping preference and choice.

A practical challenge in modeling the demand for variety is accounting for this multitude of changing internal and external factors. Many variables like an individual's mood or in-store assortment layouts are not commonly observable to researchers. Even if some variables such as past consumption, promotions, or time of day are observable, they are unlikely to constitute an exhaustive list of salient factors in a given purchase occasion. It is also not obvious how these variables should be embedded within a formal model of demand because they are likely to shift different aspects of preference. For example, high temperatures in the summer months may lead to an upward shift the preference for ice cream but not affect the rate of satiation. In contrast, periods of intense repeated consumption of a particular good may not only lead to satiation, but also more long-lasting effects on baseline preference.

In this paper, we propose a model of consumer demand that allows for flexible temporal dynamics in horizontal variety seeking. We model horizontal variety using a direct utility model of multiple discrete/continuous demand (Kim et al., 2002; Bhat, 2005, 2008) in which utility is characterized by a set of baseline utility parameters and satiation parameters. The former shift marginal utility up and down while the latter control the rate at which marginal utility diminishes. In this class of models, horizontal variety arises through a nonlinear utility function which, when maximized subject to a budget constraint, admits both corner and interior solutions.

In order to account for temporal dynamics, we impose dynamic factor structures on each set of utility parameters. That is, we treat baseline utility and satiation parameters as high-dimensional time series objects whose co-movements are driven by a lower-dimensional set of common factors. The latent factors evolve according to a dynamic process that allows for serial correlation over time. A benefit of these factor models is that they allow for the flexible evolution of preferences over time without requiring an exhaustive set of internal preference shifting variables or external market conditions. They also provide a source of dimension reduction, which is critical when models of multiple discrete/continuous demand are applied to large choice sets.

We then explore a variety of specifications for the relationship between the two factor models. We first specify two independent factor models that share no common structure. In reality, however, the latent factors that drive the evolution of baseline preferences are likely related to the factors driving the evolution of satiation. Therefore, we also propose integrating the two factor models using a threshold switching structure in which current-period baseline utility factors are only allowed to change when the prior-period satiation factor exceeds a threshold. Using this threshold switching mechanism, we then allow baseline utility factors to either update through an autoregressive process or move in the direction of satiation preference factors.

In our empirical analysis we use household-level scanner data from the yogurt category. We model the demand for a relatively large product set (35 SKUs) to ensure that we capture the boundaries of variety. Our analysis starts with descriptive evidence that the demand for variety exists, both within a give purchase occasion and over time. We then fit the proposed models to the data and find that the best fitting model integrates the two dynamic factor models and allows baseline utility factors to move with satiation factors. Estimates of model parameters show that there is substantial cross-sectional heterogeneity in the incidence and degree in preference movements.

We demonstrate the value of modeling temporal dynamics in variety in the context of assortment breadth decisions faced by retailers. Decisions to carry wide assortments are often anecdotally justified on the basis of cross-sectional heterogeneity. However, it is not clear how retailers should manage assortments in the presence of demand-side preference dynamics. For example, if tastes for certain brands or product attributes systematically weakens over time, then retailers may want to remove those varieties from the assortment. On the other hand, if as preferences evolve there remains a sufficient amount of cross-sectional heterogeneity at any point in time, then consumers will benefit from having access to wider assortments.

We address these questions by computing compensating values for products and product attributes in the retailer's assortment. Compensating values are utility-based measures of value reflecting the dollar amount required to compensate a household when a product or attribute is removed from the assortment. We first find that the proposed dynamic models generate larger estimates of compensating values relative to a static model. That is, consumers value product variety offered in wider assortments as their preferences evolve over time. We also find that product-level shares of total quarterly value remain fairly stable over time, implying that the retailer should favor wide assortments over narrow assortments with regular modifications. These results point to the importance of accounting for not only cross-sectional preference heterogeneity, but also temporal dynamics in preferences when valuing assortment breadth.

Our work adds to the literature on variety seeking by demonstrating the importance of accounting for both horizontal and temporal variety seeking behavior. Past research has largely explored one dimension at a time, but not together. For example, the models of McAlister (1982), Lattin (1987), Bawa (1990), Trivedi et al. (1994), and Chintagunta (1998) explain changes in consumption patterns over time, but fail to accommodate horizontal variety in a given purchase occasion. In contrast, Kim et al. (2002, 2007) develop models of horizontal variety but assume that baseline utility and satiation parameters are fixed over time. The modeling framework of Hasegawa et al. (2012) is closest to ours, as they also employ a dynamic factor model on the satiation parameters of a discrete/continuous demand model to account for both horizontal and temporal variety.

We make two contributions relative to Hasegawa et al. (2012). The first is terms of the proposed dynamics. In addition to the factor model on the satiation parameters, we add a second dynamic factor model to the baseline utility parameters and propose a mechanism for integrating the two factor models. Because satiation drives variation in purchase *quantities*, the added dynamics in baseline utility provides additional flexibility in modeling dynamics in secondary demand and the consumption of different product attributes over time. The second contribution is in terms of the empirical application. The analysis of Hasegawa et al. (2012) uses data from a laboratory experiment featuring undergraduate students who made weekly purchases from among eight varieties of corn chips (Kim et al., 2007). In contrast, we apply our model to household-level scanner data from the yogurt category with 35 different SKUs, thus providing external validity to many of the results in Hasegawa et al. (2012).

Our work also contributes to the marketing literature that investigates discrete shifts in marketplace behavior. For example, Markov-switching autoregressive processes have been used to model online browsing behavior (Montgomery et al., 2004), customer relationship dynamics (Netzer et al., 2008), and cyclical buying behavior (Park and Gupta, 2011). Evidence of "threshold effects" have also been found in the context advertising dynamics (Vakratsas et al., 2004; Dubé et al., 2005; Terui et al., 2011), choice set formation (Gilbride and Allenby, 2004; Swait and Erdem, 2007), asymmetric price responses (Terui and Dahana, 2006), online group buying (Wu et al., 2015), and smoking habits (Machado and Sinha, 2007). Methodologically, our approach differs from traditional Markov-switching models in that we impose a switching structure on a dynamic *factor model* rather than dynamic *linear model*. Specifically, we allow factor loadings to exhibit structural changes over time to enhance model flexibility and help identify sources of temporal dynamics. The novelty of our switching specification is that it allows us to posit a common set of factors that drive both satiation and baseline utility dynamics. This is also in contrast to the majority of existing regime-switching models embedded within a single factor model (e.g., Liu and Chen, 2016, 2020).

The remainder of the paper is organized as follows. Section 2 introduces the underlying demand model for horizontal variety. Section 3 outlines the two dynamic factor models on baseline utility and satiation parameters that account for temporal variety. Section 4 presents the results from our empirical application. Section 5 demonstrates the importance of dynamics for valuing assortment breadth. Concluding remarks are offered in Section 6.

2 A Demand Model for Horizontal Variety

Our modeling framework begins with a demand model for horizontal variety (Kim et al., 2002; Bhat, 2005, 2008) in which consumers maximize a direct utility function subject to a budget constraint.

$$\max U(\boldsymbol{x}_{ht}) = \sum_{i=1}^{n} \frac{\psi_{iht}}{\gamma_{iht}} \log(\gamma_{iht} x_{iht} + 1)$$
$$\text{s.t. } \boldsymbol{p}'_{ht}\boldsymbol{x}_{ht} = E_{ht} \tag{1}$$

Here $\boldsymbol{x}_{ht} = (x_{1ht}, \ldots, x_{nht})$ is the vector of quantities demanded by consumer h at time t, $\boldsymbol{p}_{ht} = (p_{1ht}, \ldots, p_{nht})$ is the associated vector of prices, and E_{ht} is total expenditure. Utility is additively separable in $\boldsymbol{x}_{ht}$ and the subutility for good i depends on a baseline utility parameter ψ_{iht} and a satiation parameter γ_{iht}. Specifically, ψ_{iht} sets the value of marginal utility when $x_{iht} = 0$ and γ_{iht} governs the rate at which marginal utility decreases.

In this model, horizontal variety arises through the diminishing marginal returns implied by the nonlinear utility specification in (1). That is, consumers are assumed to have preferences that diminish in intensity as consumption quantities increase, leading them to satiate when consuming large quantities. The presence of satiation can in turn admit demand solutions in the interior of the feasible set where multiple product offerings can be chosen simultaneously.

Fitting the model to observed purchase data requires introducing stochastic errors into the utility specification in (1). To ensure strictly positive marginal utility, we add multiplicative errors to the baseline utility parameters: $\psi_{iht} = \exp(\psi^*_{iht} + \varepsilon_{jht})$ where ε_{iht} follows a type I extreme value distribution as in Bhat (2005, 2008). The demand equations are solved for using the method of Lagrange multipliers.

$$\max \ \mathscr{L}(\boldsymbol{x}_{ht}, \lambda) = \sum_{i=1}^{n} \frac{\psi_{iht}}{\gamma_{iht}} \log(\gamma_{iht} x_{iht} + 1) + \lambda(E_{ht} - \boldsymbol{p}'_{ht}\boldsymbol{x}_{ht}) \tag{2}$$

The associated Kuhn-Tucker first-order conditions are given by:

$$\frac{\psi_{iht}}{\gamma_{iht} x_{iht} + 1} - \lambda p_{iht} = 0 \ \text{ if } x_{iht} > 0; \tag{3}$$
$$\frac{\psi_{iht}}{\gamma_{iht} x_{iht} + 1} - \lambda p_{iht} < 0 \ \text{ if } x_{iht} = 0. \tag{4}$$

In order to enforce the budget constraint which projects demand onto a $(n-1)$-dimensional surface, we difference the first-order conditions above with respect to the first good. After taking logs and rearranging, we have:

$$V_{iht} + \varepsilon_{iht} = V_{1ht} + \varepsilon_{1ht} \quad \text{if } x_{iht} > 0;\ i = 2, \ldots, n \tag{5}$$

$$V_{iht} + \varepsilon_{iht} < V_{1ht} + \varepsilon_{1ht} \quad \text{if } x_{iht} = 0;\ i = 2, \ldots, n \tag{6}$$

where $V_{iht} = \psi^*_{iht} - \log(\gamma_{iht} x_{iht} + 1) - \log(p_{iht})$. Assuming, without loss of generality, that the first $m < n$ goods are purchased, the probability of observed demand can then be written as:

$$\begin{aligned} P(\boldsymbol{x}_{ht}) &= P(x_{1ht} > 0, \ldots, x_{mht} > 0, x_{(m+1)ht} = 0, \ldots, x_{nht} = 0) \quad (7) \\ &= |J_{ht}| \int \left(\prod_{i=1}^{m} \phi(V_{1ht} - V_{iht} + \varepsilon_{1ht}) \right) \left(\prod_{i=m+1}^{n} \Phi(V_{1ht} - V_{iht} + \varepsilon_{1ht}) \right) \phi(\varepsilon_{1ht}) d\varepsilon_{1ht} \\ &= \frac{1}{p_{1ht}} \left(\prod_{i=1}^{m} \frac{\gamma_{iht}}{\gamma_{iht} x_{iht} + 1} \right) \left(\sum_{i=1}^{m} p_{iht} \cdot \frac{\gamma_{iht} x_{iht} + 1}{\gamma_{iht}} \right) \left(\frac{\prod_{i=1}^{m} e^{V_{iht}}}{(\sum_{j=1}^{n} e^{V_{jht}})^m} \right) (m-1)! \end{aligned}$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the PDF and PMF of the type I extreme value distribution and $|J_{ht}|$ is the determinant of the Jacobian arising from the change-of-variables mapping the error terms to the observed data. The assumption of extreme value errors gives rise to the closed-form solution shown above (Bhat, 2005, 2008).

3 Factor Models for Temporal Dynamics in Variety

3.1 Satiation Dynamics

Our model for satiation dynamics first projects the log-transformed satiation parameters $\gamma^*_{iht} \equiv \log(\gamma_{iht})$ onto a p-dimensional vector of product characteristics: $\gamma^*_{iht} = \boldsymbol{c}_i \boldsymbol{\alpha}_{ht}$. The coefficient vector $\boldsymbol{\alpha}_{ht}$ explains how the given characteristics drive satiation of product i at a given purchase occasion. We then specify a dynamic factor structure on the attribute coefficients:

$$\boldsymbol{\alpha}_{ht} = \boldsymbol{\alpha}_h + \boldsymbol{a} f_{ht} + \boldsymbol{e}_{ht},\ \boldsymbol{e}_{ht} \sim N(0, \Sigma) \tag{8}$$

where $\boldsymbol{\alpha}_h$ is a consumer-level time-invariant intercept, f_{ht} is a scalar latent factor, $\boldsymbol{a}$ is a time-invariant p-dimensional vector of factor loadings that is common across households, and $\boldsymbol{e}_{ht}$ is a p-dimensional vector of error terms with covariance matrix Σ. This factor structure allows the satiation for each product to follow a single time trend f_{ht} common across goods, while also allowing for product-specific deviations through the intercept and factor loadings. We assume a one-dimensional factor structure following Hasegawa et al. (2012) who do not find material improvements in fit from a more complex representation.

Factor dynamics arise through the evolution of the latent factors in (8), which we assume follow a random walk without drift.

$$f_{ht} = f_{ht-1} + \nu_{ht},\ \nu_{ht} \sim N(0,1) \tag{9}$$

This specification imposes a smoothness prior on changes in the factor over time (Kitagawa and Gersch, 1984; Harvey, 1989; West and Harrison, 1996) and has been successfully employed in a variety of state space modeling applications (e.g., Terui et al., 2010; Hasegawa et al., 2012; Terui and Ban, 2014).

3.2 Baseline Utility Dynamics

We again start by projecting baseline utility parameters onto the same p-dimensional vector of characteristics: $\psi^*_{iht} = \boldsymbol{c}_i \boldsymbol{\beta}_{ht}$. The attribute vector $\boldsymbol{\beta}_{ht}$ is then written as a factor model:

$$\boldsymbol{\beta}_{ht} = \boldsymbol{\beta}_h + \boldsymbol{b}\boldsymbol{g}_{ht} + \boldsymbol{u}_{ht},\ \boldsymbol{u}_{ht} \sim N(0,\Omega) \tag{10}$$

where $\boldsymbol{\beta}_h$ is consumer-level time-invariant intercept, $\boldsymbol{g}_{ht}$ is a k-dimensional vector of latent factors, $\boldsymbol{b}$ is a time-invariant $p \times k$ matrix of factor loadings, and $\boldsymbol{u}_{ht}$ is a p-dimensional vector of stochastic errors with covariance matrix Ω. Each row of the factor loadings matrix represents a product, so each row vector defines coordinates for that offering in a k-dimensional space. We use a two-dimensional factor structure ($k = 2$) for ease of interpretation, but this can easily be extended.

We consider four possible specifications of the latent factors $\boldsymbol{g}_{ht}$ which are outlined below.

M1 **(Static)**: The baseline utility factor does not change over time.

$$\boldsymbol{g}_{ht} = \boldsymbol{g}_{h} \tag{11}$$

M2 **(Dynamic) Random walk without drift**: The baseline utility factors follow a random walk, mirroring the dynamics in (9).

$$\boldsymbol{g}_{ht} = \boldsymbol{g}_{ht-1} + \boldsymbol{\eta}_{ht}, \ \boldsymbol{\eta}_{ht} \sim N(0, I_k) \tag{12}$$

M3 **(Dynamic) Threshold switching**: The baseline utility factor follows a random walk but only changes when the corresponding satiation factor exceeds a threshold value, therefore integrating the satiation and preference factor models.

$$\boldsymbol{g}_{ht} = \begin{cases} \boldsymbol{g}_{ht-1} + \boldsymbol{\eta}_{ht}, \ \boldsymbol{\eta}_{ht} \sim N(0, I_k) & \text{if } f_{ht-1} \geq r_h \\ \boldsymbol{g}_{ht-1} & \text{otherwise} \end{cases} \tag{13}$$

M4 **(Dynamic) Threshold switching with common factors**: The satiation and preference factor models are again integrated by a threshold switching structure, however this specification allows both sets of utility parameters to be driven by a common set of factors.

$$\boldsymbol{g}_{ht} = \begin{cases} \boldsymbol{\delta}_{h} f_{ht-1} + \boldsymbol{\eta}_{ht}, \ \boldsymbol{\eta}_{ht} \sim N(0, I_k) & \text{if } f_{ht-1} \geq r_h \\ \boldsymbol{g}_{ht-1} & \text{otherwise} \end{cases} \tag{14}$$

The two focal differences between the dynamic specifications above are whether baseline utility factor dynamics depend on satiation factor dynamics, and what the mean function of the dynamic updating equation looks like. A priori, one may believe that the latent drivers of baseline utility are not entirely separate from the drivers of satiation, which motivates an integrated factor model, as in M3 and M4. Further, one may want to allow a common set of latent drivers for dynamics in both sets of parameters, as in M4. Both integrated models

(M3 and M4) also include a threshold switching structure that allows baseline utility to shift when satiation exceeds some threshold. This type of threshold switching mechanism is in part motivated by the many models of discrete behavioral changes which have been found to provide an accurate description of temporal dynamics in consumer behavior (e.g., Vilcassim and Jain, 1991; Montgomery et al., 2004; Terui and Dahana, 2006; Netzer et al., 2008; Park and Gupta, 2011).

3.3 Heterogeneity and Factor Identification

In the proposed models, preference heterogeneity arises both across households and within a household over time. Temporal heterogeneity is modeled by the dynamic factor structures in (8) and (10). We model cross-sectional heterogeneity by allowing the attribute intercepts $\boldsymbol{\alpha}_h$ and $\boldsymbol{\beta}_h$ to follow a normal distribution in the population.

$$\boldsymbol{\alpha}_h \sim N(\bar{\boldsymbol{\alpha}}, V_\alpha) \tag{15}$$

$$\boldsymbol{\beta}_h \sim N(\bar{\boldsymbol{\beta}}, V_\beta) \tag{16}$$

We next turn to identification of the proposed dynamic factor models. The standard factor model is not identified due to the rotational invariance of the factors and factor loadings. This indeterminacy has been studied for both static factor models (Anderson and Rubin, 1956) and dynamic factor models (Bai and Wang, 2015). In both cases, identification can be achieved by imposing constraints on factor loadings. For example, a common strategy is to assume that the factor loading matrix is a lower-triangular matrix with strictly positive elements along the diagonal (Geweke and Zhou, 1996; Aguilar and West, 2000). However, these restrictions are usually applied to factor models where the response variable is observable. In our case, the response variables for both factor models correspond to latent parameters in the consumer's utility function. We therefore impose a slightly stronger restriction on factor loadings.

$$\boldsymbol{a} = \begin{pmatrix} 1 \\ a_2 \\ a_3 \\ \vdots \\ a_p \end{pmatrix} \qquad \boldsymbol{b} = \begin{pmatrix} 1 & 0 \\ b_{21} & 1 \\ b_{31} & b_{32} \\ \vdots & \vdots \\ b_{(p-1)1} & b_{(p-1)2} \end{pmatrix} \tag{17}$$

In addition to the factor loadings, each factor model contains a set of time-invariant intercept parameters. The baseline utility intercepts $\boldsymbol{\beta}_h$ are identified by the relative variation in purchase shares for each product. However, the differencing applied to the first-order conditions (required to enforce the budget constraint) means that not all product-specific intercepts can be identified. We therefore set $\psi^*_{nht} = 0$ for the last product n and for all $h = 1, \ldots, H$ and $t = 1, \ldots, T_h$. Moreover, the satiation intercepts $\boldsymbol{\alpha}_h$ are identified by contemporaneous variation in purchase quantities over time. For identification, we fix $\alpha_{h1} = 0$ for each $h = 1, \ldots, H$.

The proposed models contain two sets of covariance matrices: (Σ, Ω) are the covariance matrices on the error terms of each factor model and (V_α, V_β) parameterize the distributions of cross-sectional preference heterogeneity. The factor model covariance matrices can in principle be identified from both the autocorrelation function defining dynamics as well as the contemporaneous cross-product variation in utility parameters (Bai and Wang, 2015). However, it is often infeasible to obtain precise estimates of covariance terms for a large number of cross-sectional units (i.e., products). The contemporaneous cross-product variation in utility parameters is also confounded with cross-sectional heterogenetiy. We therefore restrict $\Sigma = I_p$ and $\Omega = I_p$ and model the heterogeneity covariance matrices as $V_\alpha = \text{diag}(v^2_{\alpha 1}, \ldots, v^2_{\alpha p})$ and $V_\beta = \text{diag}(v^2_{\beta 1}, \ldots, v^2_{\beta p})$.

The proposed models also contain switching coefficients $\boldsymbol{\delta}_h$ and latent threshold r_h. The switching coefficients in M4 are identified from the differences between the current-period baseline utility factors and lagged satiation factors. In both threshold switching models (M3 and M4), we set $r_h = 0$.

4 Empirical Application

4.1 Data Description

Our empirical analysis uses two years of data from the IRI household panel data set (Bronnenberg et al., 2008). We focus on the yogurt category, which is known for having many brands, product lines, and flavors. In our sample we include four national manufacturers (Colombo, Dannon, Stonifield Farm, and Yoplait) and a private label. Two manufacturers carry multiple product lines: Yoplait (Original vs. Light) and Dannon (Fruit on the Bottom and Light & Fit). For each of the seven total product lines, we choose the top five SKUs within each line based on market share. This yields a total of 35 SKUs (see product descriptions and summary statistics in Appendix A). Our final sample consists of 387 households who make at least five purchases from these 35 SKUs in a single retail store in Pittsfield, Massachusetts.

4.2 Descriptive Evidence of Variety Seeking Dynamics

We observe characteristics in the data which point to the need to accommodate both horizontal and temporal variety seeking. Evidence of horizontal variety seeking is provided in Table 1. We find that 60.4% of all trips involve the purchase of more than one SKU, with an average of about two unique SKUs purchased per trip. Moreover, multiple product lines are chosen in 7.5% of all trips with an average of one product line purchased per trip. We also find that 44.4% of all households have had a trip for which they purchased multiple product lines and 96.1% of households have had a trip for which they purchased multiple SKUs.

Table 1: Horizontal Variety Seeking

	SKUs	Product Lines
Average number of unique per trip	2.02	1.08
Trips with ≥ 1 chosen	60.4%	7.59%
Households with at least one trip where ≥ 1 chosen	96.1%	44.4%

We also find evidence of changes in purchase patterns over time. Table 2 reports three measures of purchase dynamics. First, we calculate the number of unique items purchased in the sample period divided by the number of trips. This is referred to as "share of unique" and will be close to one if many different SKUs or product lines are purchased over time. Second, for each household, we find the SKU or product line that is purchased most often and report the share of weeks for which it was purchased. This is referred to as "maximum share" and is indicative of state-dependence. Third, we split each household's sample in half and compute the number SKUs or product lines purchased in the second half of the sample period that were not purchased in the first half. Across all three statistics, we find that consumers exhibit strong temporal variety seeking behavior, especially at the SKU-level.

Table 2: Temporal Variety Seeking

	SKUs	Product Lines
Share of unique	0.741	0.357
Maximum share	0.587	0.694
Unique in second half	0.444	0.291

The various dynamics discussed above can also be seen in Figure 1, which plots the purchase patterns for three households in our sample. The top panel plots the consumption of different product lines over time, while the bottom panel plots the consumption of different flavors. Each point denotes the purchase incidence of a given product line or flavor and the size denotes the total number of SKUs consumed. We find evidence of horizontal and temporal variety, but in ways that differ drastically across households. For example, the first household (left panel) only purchases one product line (Yoplait) and mostly one flavor (strawberry) across all trips. The second household (middle panel) purchases only Yoplait Light & Fit for a while and then switches to Dannon Light & Fit. However this household regularly purchases multiple flavors within and across trips. The third household (right panel) purchases many different product lines over time and many flavors within a line for any given trip. Together, these patterns highlight the importance of allowing for horizontal variety, temporal variety, and cross-sectional preference heterogeneity.

(a) Product Line

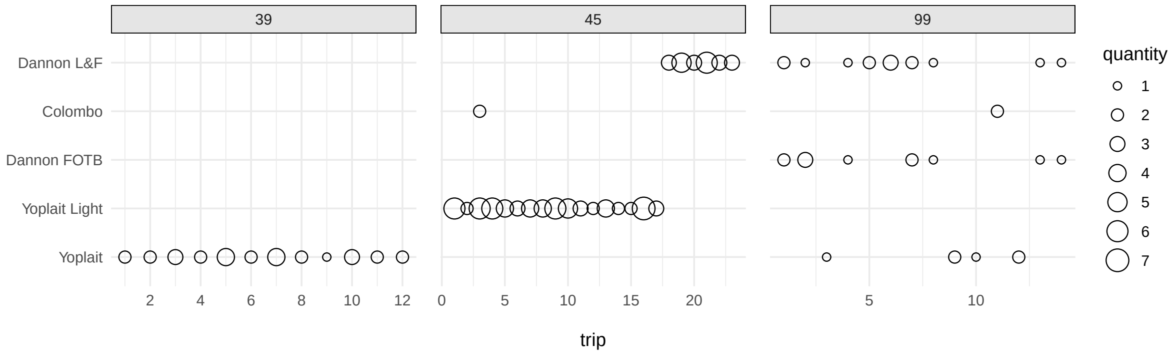

(b) Flavor

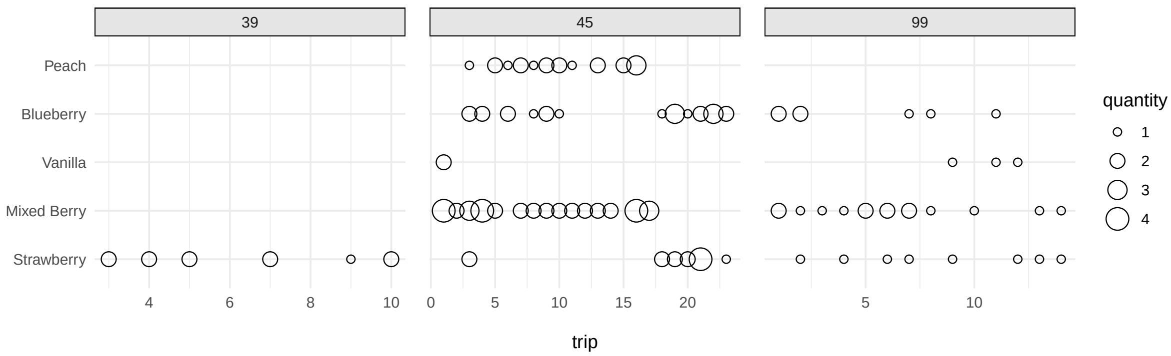

Figure 1: Purchase patterns for a sample of three households.

4.3 Estimation Results

We fit a total of five models to the data: the four candidate models outlined in Section 3.2 (M1-M4) and a purely static model (M0) in which both baseline utility and satiation parameters are constant over time. The complete set of models are summarized in Table 3 below. All models are specified as Bayesian hierarchical models where the hierarchical priors are defined by the factor models in (8) and (10) and models of cross-sectional heterogeneity in (15) and (16). We place normal diffuse priors on the elements of the factor loadings matrices $\boldsymbol{a}$ and $\boldsymbol{b}$ as well as the switching coefficients, and inverse gamma priors on the diagonal elements of each covariance matrix. Markov chain Monte Carlo methods are used to sample from each posterior distribution. Detailed descriptions of the sampling algorithms and choices of hyperparameters are provided in Appendix B.

Table 3: Summary of Models

	Satiation		Baseline Utility		
Model	Dynamics	Specification	Dynamics	Specification	Switching
M0	No	$f_{ht} = f_h$	No	$\boldsymbol{g}_{ht} = \boldsymbol{g}_h$	–
M1	Yes	$f_{ht} = f_{ht-1} + \nu_{ht}$	No	$\boldsymbol{g}_{ht} = \boldsymbol{g}_h$	–
M2	Yes	$f_{ht} = f_{ht-1} + \nu_{ht}$	Yes	$\boldsymbol{g}_{ht} = \boldsymbol{g}_{ht-1} + \boldsymbol{\eta}_{ht}$	–
M3	Yes	$f_{ht} = f_{ht-1} + \nu_{ht}$	Yes	$\boldsymbol{g}_{ht} = \boldsymbol{g}_{ht-1} + \boldsymbol{\eta}_{ht}$	if $f_{ht-1} \geq r_h$
M4	Yes	$f_{ht} = f_{ht-1} + \nu_{ht}$	Yes	$\boldsymbol{g}_{ht} = \boldsymbol{\delta}_h f_{ht-1} + \boldsymbol{\eta}_{ht}$	if $f_{ht-1} \geq r_h$

Model Fit Statistics

Table 4 reports in-sample and out-of-sample fit statistics. In-sample fit is evaluated using the log marginal density (LMD) (calculated using the Newton-Raftery algorithm of Newton and Raftery, 1994), the deviance information criterion (DIC), and the root-mean-squared error (RMSE). Out-of-sample fit is calculated using the predictive RMSE. All model fit statistics point to the need to accommodate dynamics in both satiation and baseline utility parameters. The model with most empirical support is the integrated model (M4) allowing a common set of factors to drive both satiation and baseline utility dynamics. Going forward, we focus our attention on the parameters of this model.

Table 4: Model Fit Statistics

	In-Sample			Predictive
Model	LMD	DIC	RMSE	RMSE
M0	-30,853	180,336	0.469	0.598
M1	-29,715	175,457	0.457	0.598
M2	-24,546	150,634	0.412	0.621
M3	-25,487	156,414	0.405	0.602
M4	-22,774	140,127	0.404	0.593

Estimates of Utility Parameters

We start by interpreting the parameters in the baseline utility and satiation attribute vector equations in (8) and (10), which reveal information about how the given attributes contribute to average tastes and satiation levels, respectively. The estimates of the population-level averages of the baseline utility intercepts $\bar{\beta}_j$, variances of the baseline utility intercepts $v^2_{\beta j}$, and factor loadings $\boldsymbol{b}_j = (b_{j1}, b_{j2})$ are reported in Table 5. We find that the estimates of $\bar{\beta}_j$ are roughly in line with the ranking of observed product-line-level and flavor-level market shares (see Appendix A), which provides some face validity to our results. For example, the most preferred brands are Yoplait and the private label, while the most preferred flavors are strawberry and blueberry. The estimates of the variances $v_{\beta j}^2$ suggest that significant cross-sectional heterogeneity exists, especially for the brand attributes. Turning to the two-dimensional factor loadings, we find that while the dimensions are directionally consistent with each other, there is a large difference in scale with the first dimension being 3-15 times larger than the second dimension. The first dimension especially penalizes the Dannon FOTB and Dannon L&F product lines, while favoring Colombo and Yoplait Light. In terms of flavor attributes, the first dimension assigns the largest value to fruity flavors such as peach and strawberry, whereas the second dimension assigns the largest value to vanilla.

Table 5: Baseline Utility Parameters

	$\bar{\beta}_j$		$v_{\beta j}^2$	b_{j1}		b_{j2}	
	Mean	SD	Mean	Mean	SD	Mean	SD
Yoplait	-0.495	(0.305)	52.129	1.000	-	0.000	-
Yoplait Light	-5.304	(0.155)	55.416	4.373	(0.142)	1.000	-
Dannon FOTB	-6.648	(0.101)	32.012	-6.103	(0.183)	-0.685	(0.028)
Colombo	-5.904	(0.180)	94.793	6.862	(0.243)	0.661	(0.029)
Stonyfield	-6.327	(0.176)	53.700	-3.462	(0.139)	-0.648	(0.027)
Dannon L&F	-5.572	(0.107)	33.983	-8.037	(0.138)	-0.533	(0.031)
PL	-0.600	(0.116)	1.360	-0.719	(0.087)	-0.224	(0.010)
Strawberry	-0.046	(0.141)	2.131	0.749	(0.064)	0.175	(0.011)
Mixed Berry	-0.148	(0.125)	2.571	0.741	(0.064)	0.193	(0.010)
Vanilla	-0.833	(0.174)	7.298	0.706	(0.074)	0.211	(0.015)
Blueberry	0.387	(0.165)	3.018	0.533	(0.076)	0.156	(0.011)
Peach	-0.223	(0.132)	2.539	0.763	(0.071)	0.161	(0.010)

The estimates of the population-level averages of the satiation intercepts, variances of the satiation intercepts, and satiation factor loadings are reported in Table 6. The estimates of the satiation intercepts $\bar{\alpha}_j$ suggest that, among all product lines, average levels of satiation are highest for varieties of Yoplait and lowest for Dannon Light & Fit and the private label. Among all flavors of yogurt, blueberry has the highest reported level of satiation while vanilla has the lowest (consistent with previous studies of satiation in the yogurt category, such as Kim et al., 2002). Cross-sectional heterogeneity exists as evidenced by the variances $v_{\alpha j}^2$, however the magnitude of this heterogeneity is smaller than that of the baseline utility model. Finally, the sign and magnitude of the factor loadings are consistent with the intercept estimates suggesting that the latent factor is capturing additional satiation effects.

Table 6: Satiation Parameters

	$\bar{\alpha}_j$		$v^2_{\alpha j}$	a_j	
	Mean	SD	Mean	Mean	SD
Yoplait	1.753	(0.094)	1.000	1.000	-
Yoplait Light	0.766	(0.149)	0.109	0.334	(0.029)
Dannon FOTB	0.656	(0.096)	0.149	0.358	(0.026)
Colombo	0.584	(0.091)	0.132	0.371	(0.033)
Stonyfield	0.953	(0.104)	0.188	0.283	(0.053)
Dannon L&F	0.323	(0.087)	0.187	0.273	(0.041)
PL	0.345	(0.138)	0.319	0.224	(0.043)
Strawberry	-1.133	(0.089)	0.095	-0.279	(0.037)
Mixed Berry	-1.480	(0.077)	0.083	-0.329	(0.034)
Vanilla	-1.589	(0.098)	0.189	-0.314	(0.030)
Blueberry	-0.906	(0.073)	0.071	-0.287	(0.031)
Peach	-0.967	(0.096)	0.083	-0.264	(0.038)

Household-Level Switching Propensities

The best fitting dynamic factor model contains a threshold switching equation that allows the baseline utility factor model to move with (or away from) the satiation factor model when satiation exceeds some threshold. The direction of these changes is summarized by the coefficients on the satiation factors, which are reported in Table 7. The negative coefficient on the first dimension and positive coefficient on the second dimension suggest that preferences tend to rotate "counter-clockwise" around the latent two-dimensional space defined by the baseline utility factor loadings reported in Table 5.

Table 7: Switching Coefficients

	Mean	SD
$\bar{\delta}_1$	-0.019	(0.055)
$\bar{\delta}_2$	0.501	(0.235)

We also explore the heterogeneity in each household's propensity to undergo preference "switches" during the sample period. For household h, we compute the proportion of times the switching equation in (14) is turned on over the course of the R post-burn-in iterations for each trip $t = 2, \dots, T_h$

$$k_{ht} = \frac{1}{R} \sum_{r=1}^{R} I(f_{ht-1}^{(r)} \geq r_h). \tag{18}$$

For each household, we then compute the mean and the range across trips. Large mean values indicate that the household has a high average switching propensity, while large range values suggest that the household undergoes periods of both being likely to switch and not very likely to switch.

Figure 2 plots household-level average switching propensities against their dispersion. The marginal distribution of each statistic is also shown in the margin of each axis. We find that a large share of households fall in the lower-right quadrant which corresponds to a high average switching propensity and a low range. These households are likely to undergo switches across all observed trips. In contrast, households with a high average switching propensity and high range are likely to have purchase patterns exhibiting some form of state-dependence for some time, but then eventually switch to consume products in different regions of the attribute space.

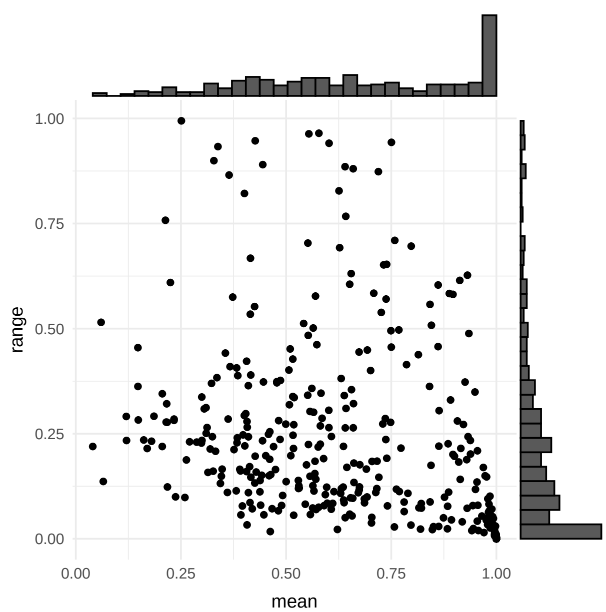

Figure 2: The mean and range of household-level switching propensities.

Temporal Dynamics in Utility Parameters

We further explore the temporal dynamics in preference by plotting the attribute coefficients from both factor models ($\boldsymbol{\beta}_{ht}$ and $\boldsymbol{\alpha}_{ht}$). Figure 3 separately plots the evolution of product line coefficients (top panel) and the flavor coefficients (bottom panel) for the same three households shown in Figure 1. The first household's coefficient vectors (left panel) exhibit relatively little variation over time, which is reflective of the steady purchase patterns shown in Figure 1. The lack of dynamics in utility parameters for this first household can also be explained by its relatively small switching coefficient: mean(k_{39}) = 0.23 and range(k_{39}) = 0.10.

In contrast, the second and third households (middle and right panels) exhibit more notable dynamic trends. For example, the second household's baseline utility for Yoplait Light starts high initially but then falls over time, while the baseline utility for Dannon Light & Fit starts low but increases over time. These trends are clearly driven by the raw purchase patterns in Figure 1 showing this household purchasing Yoplait Light for the first 17 trips and then switching to Dannon Light & Fit for all trips thereafter. Moreover the estimates of the second household's switching coefficients (mean(k_{45}) = 0.25, range(k_{45}) = 0.99) capture the fact that preferences appear steady for some time but eventually make drastic moves around the attribute space. Similarly, the final household has a high switching propensity (mean(k_{99}) = 0.77, range(k_{99}) = 0.71) and thus exhibits dynamics in utility parameters reflecting the frequent shifts in purchase patterns shown in Figure 1.

A final point is to notice the different dimensions of heterogeneity shown in Figure 3. In addition to the within-household heterogeneity across trips, there we also observe evidence of heterogeneity across households. That is, some households tend to have higher levels of satiation or baseline utility, and the ranking of these parameters is also likely to change across households. Together, these features highlight the importance of accounting for both temporal and cross-sectional heterogeneity in models of demand.

(a) Product Line

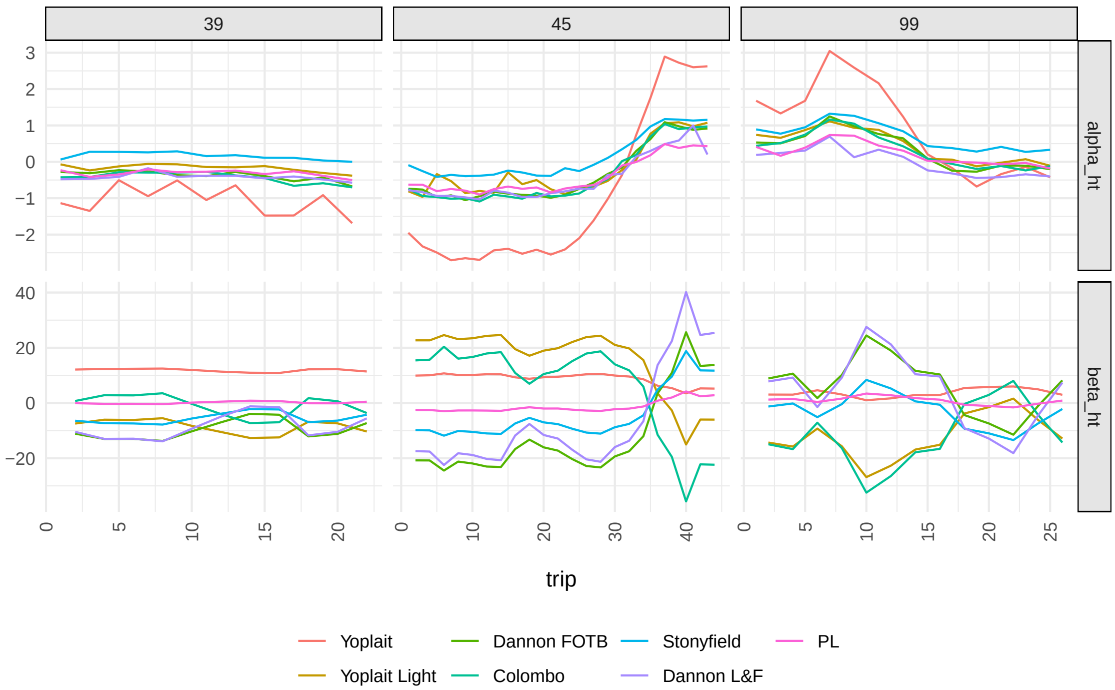

(b) Flavor

Figure 3: Temporal dynamics in product line and flavor attribute coefficients are shown for a sample of three households.

Factor Loadings and Future Attribute Consumption

Here we explore the ways in which the proposed model's dynamic factor structure can be used to capture trends in future preference and understand future consumption of different product attributes. The idea is that if the baseline utility factor loadings $\boldsymbol{b}_j = (b_{j1}, b_{j2})$ sufficiently explain average tastes in the population, then knowledge of the trajectory of a consumer's preferences at time t (as predicted by our model) should be informative of which attributes are consumed at time t. Formally, we use the factor model in (10) to measure the "match" between attribute j and a household's preferences in the hold-out sample (i.e., $t = T_h + 1$). The match is formally defined as the cosine of the angle between the vector of attribute coefficients $\tilde{\boldsymbol{g}}_{ht} = (\beta_{hj}, \boldsymbol{g}_{ht})$ and factor loadings $\tilde{\boldsymbol{b}}_j = (1, \boldsymbol{b}_j)$.

$$\text{match}_{hjt}(\tilde{\boldsymbol{b}}_j, \tilde{\boldsymbol{g}}_{ht}) = \frac{\tilde{\boldsymbol{b}}_j \tilde{\boldsymbol{g}}_{ht}}{||\tilde{\boldsymbol{b}}_j||\ ||\tilde{\boldsymbol{g}}_{ht}||} \tag{19}$$

When a household's preference vector points in the same direction as attribute j's factor loading, then the cosine of the angle will be close to one. In this case, the attribute constitutes a high match with that household's preferences. In contrast, when a household's preference vector points in the exact opposite direction as the factor loading, then the cosine of the angle will be close to negative one and will represent a low match. If a product contains attributes which have a high match with a household's predicted preferences, then a priori we would expect that attribute j should have a relatively high chance of being consumed.

Figure 4 plots the average match value for each attribute conditional on whether that attribute was consumed at time $t = T_h + 1$. We find that across all attributes, the model's factor structure is highly predictive of future consumption. That is, the attributes that are consumed in the hold-out trip tend to have a positive and relatively large match value with the model's prediction of preferences at time $t = T_h + 1$. Similarly, the product lines that are not consumed in the hold-out trip tend to have a small (close to zero) or large and negative magnitude.

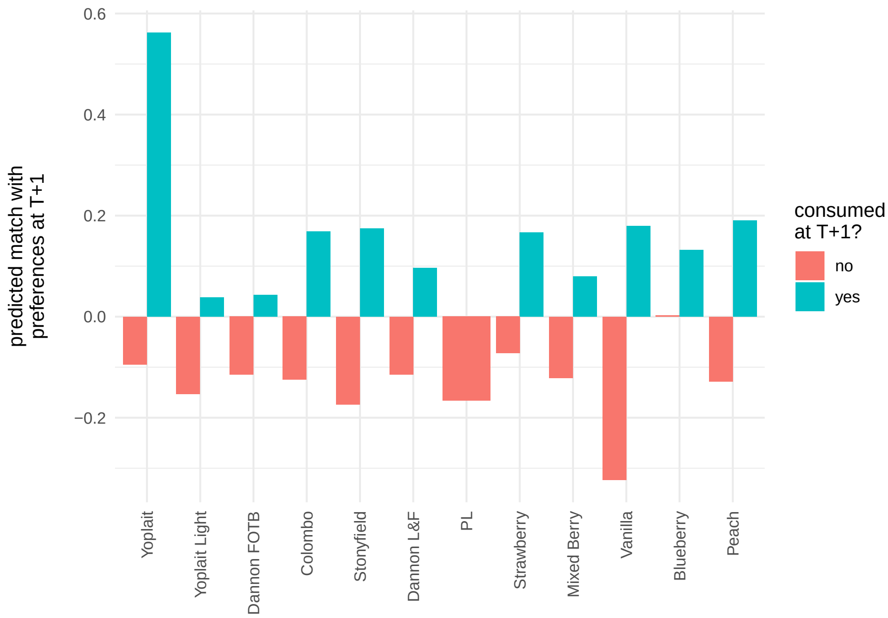

Figure 4: Average "match" values for each product attribute are shown conditional on whether the attribute was consumed in the hold-out sample.

5 Implications for Assortment Breadth

In this section we explore the implications of variety seeking dynamics on the required breadth of products assortments carried by the retailer. The vast majority of grocery retailers carry a wide variety of SKUs within every product category. Explanations for assortment breadth usually rely on cross-sectional heterogeneity: consumers have very different preferences, so retailers can benefit from offering wide and differentiated assortments to increase the likelihood that consumers finds a product that matches their tastes (e.g., Baumol and Ide, 1956; Hoch et al., 1999; Brynjolfsson et al., 2003; Waldfogel, 2003). However, assortment planning also depends on temporal heterogeneity and the evolution of consumer preferences. Being able to measure the incidence and extent of changes in preferences can help retailers determine how and when to modify existing assortments.

The proposed modeling framework allows us to assess the evolution of preferences and its impact on the demand for variety. We take a utility-based approach and compare the compensating values implied by the integrated dynamic factor model (M4) to those from a traditional static model (M0). Compensating values represent the amount of money needed to compensate an individual for changes in the price, quality, and/or breadth of products offered in an assortment. We specifically compute compensating values associated with the removal of a product from the given assortment. To calculate the compensating value of product i, we first need to measure the maximum attainable utility under both the complete and restricted assortments, denoted $V_{ht}(\boldsymbol{p}_{ht}, E_{ht})$ and $V_{ht}^{(i)}(\boldsymbol{p}_{ht}, E_{ht})$, respectively.

$$V_{ht}(\boldsymbol{p}_{ht}, E_{ht}) = \max_{\boldsymbol{x}_{ht}} U(\boldsymbol{x}_{ht}|\boldsymbol{\psi}_{ht}, \boldsymbol{\gamma}_{ht}) \quad \text{s.t.} \quad \boldsymbol{p}_{ht}'\boldsymbol{x}_{ht} = E_{ht} \tag{20}$$

$$V_{ht}^{(i)}(\boldsymbol{p}_{ht}, E_{ht}) = \max_{\boldsymbol{x}_{ht}} U(\boldsymbol{x}_{ht}|\boldsymbol{\psi}_{ht}, \boldsymbol{\gamma}_{ht}) \quad \text{s.t.} \quad \boldsymbol{p}_{ht}'\boldsymbol{x}_{ht} = E_{ht} \text{ and } x_{iht} = 0 \tag{21}$$

The compensating value of product i at time t is then defined as the value $CV_{ht}^{(i)}$ that must be added to the budget E_{ht} so that maximum attainable utility is left unchanged under the restricted assortment.

$$V_{ht}(\boldsymbol{p}_{ht}, E_{ht}) = V_{ht}^{(i)}(\boldsymbol{p}_{ht}, E_{ht} + CV_{ht}^{(i)}) \tag{22}$$

Table 8 reports estimates of attribute-level compensating values that are averaged across all households and purchase occasions. We also report the compensating values as a percent of total expenditure: $PCV_h^{(i)} = \sum_{t=1}^{T_h} CV_{ht}^{(i)} / \sum_{t=1}^{T_h} E_{ht}$. Both static and dynamic models tend to agree on which product lines and flavors are valued most, on average. For example, under both models, Yoplait generates the highest compensating value among product lines and vanilla generates the highest compensating value among flavors. Here the high value of Yoplait can be attributed to its relatively high baseline utility reported in Table 5 while the high value of vanilla can be attributed to its relatively low levels of satiation reported in Table 6. Across, models we find differences in the magnitude in of the compensating values. Specifically, the dynamic model generates estimates that are, on average, two and a half times larger than those of the static model. This is evidence of the appreciable contribution that temporal baseline preference dynamics have on driving value in wide assortments over time.

We further explore the cross-sectional heterogeneity of compensating values across products in Figure 5, which plots the distribution of average log compensating values across households

Table 8: Attribute-Level Compensating Values

	Static (M0)		Dynamic (M4)	
	CV	PCV	CV	PCV
Yoplait	3.06	1.76	4.44	3.54
Yoplait Light	1.17	0.55	2.26	1.49
Dannon FOTB	0.80	0.39	2.71	1.59
Colombo	1.17	0.68	3.68	2.41
Stonyfield	0.63	0.34	2.90	1.92
Dannon L&F	1.43	1.14	3.84	2.53
PL	1.02	0.70	3.05	1.72
Strawberry	0.36	0.19	1.73	0.98
Mixed Berry	1.01	0.47	2.24	1.09
Vanilla	1.06	0.43	2.31	1.76
Blueberry	0.72	0.38	1.66	1.00
Peach	0.36	0.23	1.34	0.73
Average	1.07	0.61	2.68	1.73

for each product line and for both models. The distributions of estimates from the dynamic model tend to have broader shoulders and a longer upper tail, in line with the attribute-based estimates reported above. Across all product lines, we find that there is a clear rightward shift in the distributions of compensating values under the proposed dynamic model.

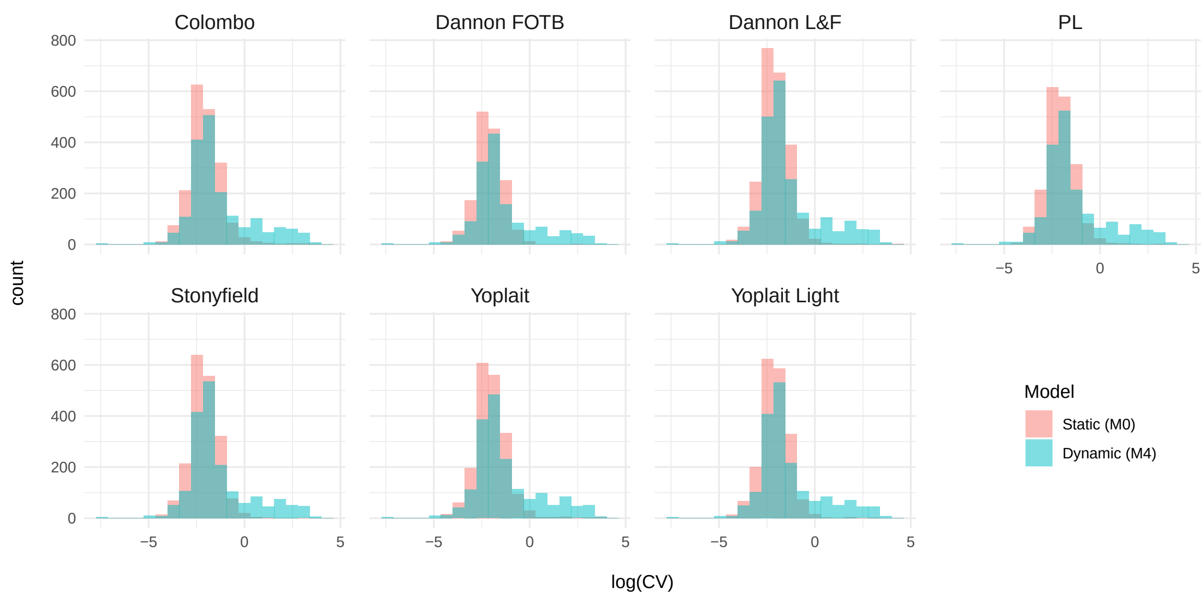

Figure 5: Distribution of compensating values across households for each product.

Our empirical evidence thus far suggests that consumers tend to value products more than what would be suggested under a model of static preferences, even after accounting for cross-sectional heterogeneity. However, what this implies for retailers also depends on the temporal dimension of heterogeneity. Consider the following illustrative example. Suppose there are two products (A and B) and two time periods ($t = 1$, $t = 2$) with per-period compensating values given by the table below.

	Scenario 1			Scenario 2		
	$t = 1$	$t = 2$	Mean	$t = 1$	$t = 2$	Mean
A	2	0	1	1	1	1
B	0	2	1	1	1	1

Although both scenarios generate the same average value of each product, the key difference is in the share that each product contributes to total value in each period. In the first scenario, product A generates 100% of the value in period one and 0% of the value in period two (with product B yielding the reverse). In the second scenario, each product generates 50% of the value in every period. Each scenario carries different implications for the assortment decisions made by retailers. When there is systematic variation in per-period shares (as in the first scenario), then retailers may be better off carrying a narrow assortment that is modified over time (e.g., only offering A and period one and B in period two). If each product contributes the same amount to total value in each period, then retailers would be better off carrying a wide assortment.

The underlying dynamics in our model allow us to also explore this temporal dimension of heterogeneity. Figure 6 plots the evolution of shares of total quarterly value associated with each product. Each line corresponds to a product and the area below each line corresponds to that product's cumulative share of quarterly value (within the product line). We find that these shares exhibit some possibly seasonal variation but do not show systematic trends over time.

These results, taken together with the results on the cross-sectional differences in compensating values, suggest that narrowing the present assortments would be very costly to consumers. The empirical support for the proposed model provides evidence that temporal dynamics exist and these dynamics lead to higher compensating values that persist over time. This persistence is meaningful for retailers, as it supports the practice of carrying wide and stable assortments rather than narrow assortments that can be modified over time.

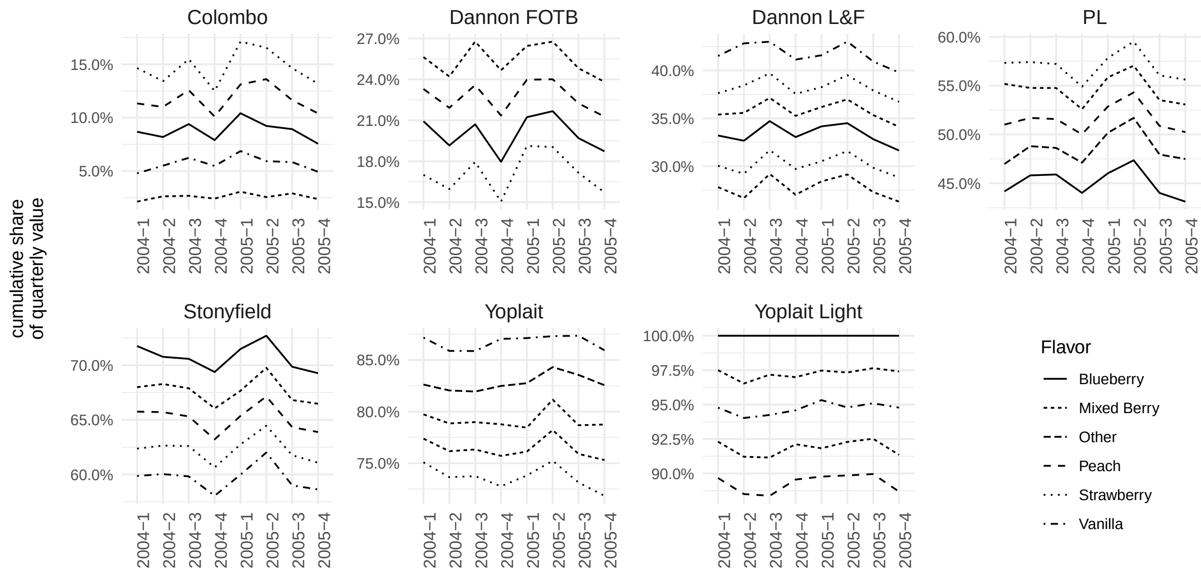

Figure 6: The evolution of shares that each product contributes to total quarterly value.

6 Conclusion

In this paper we study the demand for horizontal and temporal variety. The backbone of the proposed modeling framework is a model of multiple discrete/continuous demand which accounts for the simultaneous purchase of multiple products in a given purchase occasion. Utility is characterized by two sets of parameters, baseline utility and satiation, which each have a role in shaping preference and choice. We then add dynamic factor models onto both sets of utility parameters to allow for the evolution of preferences and changes in the demand for variety over time. We propose and test a variety of specifications for preference dynamics, ranging from a purely static model with no temporal dynamics to a model that integrates dynamic factor models of baseline preference and satiation.

Our empirical application uses data from the yogurt category, where model-free evidence suggests that households not only tend to buy multiple varieties within a given trip, but also different varieties over time. We take the proposed set of models to the data and find empirical support for an integrated model of baseline preference and satiation dynamics. In particular, the best fitting model uses a threshold switching mechanism to allow a common set of factors to drive both satiation and preference dynamics. We find substantial heterogeneity in the incidence and magnitude of these threshold switches, which also generates differences in the implied dynamics of the underlying utility parameters. We then demonstrate the importance of accounting for both horizontal and temporal dimensions of variety seeking for the purposes of valuing assortment breadth. Specifically, the proposed dynamic models allow us to disentangle cross-sectional and temporal heterogeneity, which is crucial for assortment planning decisions.

Future research is needed to better understand the context of purchase and consumption. Our results indicate that the unit of analysis is not the respondent, but instead the respondent at a specific purchase occasion who is influenced by their past decisions and other factors. Additional work is needed to identify and integrate these factors and past events into models of consumer decision making that allow marketers to anticipate shifts in preference. We believe that the temporal study of satiation, and other triggers of preference change, is a fruitful area for future research.

References

Aguilar, O. and West, M. (2000). Bayesian dynamic factor models and portfolio allocation. *Journal of Business & Economic Statistics*, 18(3):338–357.

Anderson, T. W. and Rubin, H. (1956). Statistical Inference in Factor Analysis. In *Third Berkeley Symposium on Mathematical Statistics and Probability*, volume 5, pages 111–150. University of California Press.

Bai, J. and Wang, P. (2015). Identification and bayesian estimation of dynamic factor models. *Journal of Business & Economic Statistics*, 33(2):221–240.

Baumol, W. J. and Ide, E. A. (1956). Variety in retailing. *Management Science*, 3(1):93–101.

Bawa, K. (1990). Modeling inertia and variety seeking tendencies in brand choice behavior. *Marketing Science*, 9(3):263–278.

Bhat, C. R. (2005). A multiple discrete-continuous extreme value model: Formulation and application to discretionary time-use decisions. *Transportation Research Part B*, 39(8):679–707.

Bhat, C. R. (2008). The multiple discrete-continuous extreme value (MDCEV) model: Role of utility function parameters, identification considerations, and model extensions. *Transportation Research Part B*, 42(3):274–303.

Bronnenberg, B. J., Kruger, M. W., and Mela, C. F. (2008). Database paper – the IRI marketing data set. *Marketing Science*, 27(4):745–748.

Brynjolfsson, E., Hu, Y. J., and Smith, M. D. (2003). Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science*, 49(11):1580–1596.

Carter, C. K. and Kohn, R. (1994). On gibbs sampling for state space models. *Biometrika*, 81(3):541–553.

Che, H., Erdem, T., and Öncü, T. S. (2015). Consumer learning and evolution of consumer brand preferences. *Quantitative Marketing and Economics*, 13(3):173–202.

Chintagunta, P. K. (1998). Inertia and variety seeking in a model of brand-purchase timing. *Marketing Science*, 17(3):253–270.

Dubé, J.-P., Hitsch, G. J., and Manchanda, P. (2005). An empirical model of advertising dynamics. *Quantitative Marketing and Economics*, 3:107–144.

Ershov, D. (2020). Consumer product discovery costs, entry, quality and congestion in online markets. *Working Paper*.

Geweke, J. and Zhou, G. (1996). Measuring the pricing error of the arbitrage pricing theory. *Review of Financial Studies*, 9(2):557–587.

Gilbride, T. J. and Allenby, G. M. (2004). A choice model with conjunctive, disjunctive, and compensatory screening rules. *Marketing Science*, 23(3):391–406.

Gullo, K., Berger, J., Etkin, J., and Bollinger, B. (2019). Does time of day affect variety-seeking? *Journal of Consumer Research*, 46(1):20–35.

Harvey, A. (1989). *Forecasting Structural Time Series Models and the Kalman Filter*. Cambridge University Press, London.

Hasegawa, S., Terui, N., and Allenby, G. M. (2012). Dynamic brand satiation. *Journal of Marketing Research*, 49(6):842–853.

Hoch, S. J., Bradlow, E. T., and Wansink, B. (1999). The variety of an assortment. *Marketing Science*, 18(4):527–546.

Hong, S., Misra, K., and Vilcassim, N. J. (2016). The perils of category management: The effect of product assortment on multicategory purchase incidence. *Journal of Marketing*, 80(5):34–52.

Kahn, B. E. (1995). Consumer variety seeking among goods and services. *Journal of Retailing and Consumer Services*, 2(3):139–148.

Kahn, B. E. and Wansink, B. (2004). The influence of assortment structure on perceived variety and consumption quantities. *Journal of Consumer Research*, 30(4):519–533.

Kim, J., Allenby, G. M., and Rossi, P. E. (2002). Modeling consumer demand for variety. *Marketing Science*, 21(3):229–250.

Kim, J., Allenby, G. M., and Rossi, P. E. (2007). Product attribute and models of multiple discreteness. *Journal of Econometrics*, 138(1):208–230.

Kitagawa, G. and Gersch, W. (1984). A smoothness priors-state space modeling of time series with trends and seasonalities. *Journal of the American Statistical Association*, 79(386):378–389.

Lattin, J. M. (1987). A model of balanced choice behavior. *Marketing Science*, 6(1):48–65.

Liu, X. and Chen, R. (2016). Regime-switching factor models for high-dimensional time series. *Statistica Sinica*, 26(4):1427–1451.

Liu, X. and Chen, R. (2020). Threshold factor models for high-dimensional time series. *Journal of Econometrics*, 216(1):53–70.

Machado, F. S. and Sinha, R. K. (2007). Smoking cessation: A model of planned vs. actual behavior for time-inconsistent consumers. *Marketing Science*, 26(6):834–850.

McAlister, L. (1982). A dynamic attribute satiation model of variety-seeking behavior. *Journal of Consumer Research*, 9(3):141–150.

McAlister, L. and Pessemier, E. (1982). Variety seeking behavior: An interdisciplinary review. *Journal of Consumer Research*, 9(3):311–322.

Montgomery, A. L., Li, S., Srinivasan, K., and Liechty, J. C. (2004). Modeling online browsing and path analysis using clickstream data. *Marketing Science*, 23(4):579–595.

Netzer, O., Lattin, J. M., and Srinivasan, V. (2008). A hidden markov model of customer relationship dynamics. *Marketing Science*, 27(2):185–204.

Newton, M. A. and Raftery, A. E. (1994). Approximate bayesian inference with the weighted likelihood bootstrap. *Journal of the Royal Statistical Society, Series B*, 56(1):3–48.

Park, S. and Gupta, S. (2011). A regime-switching model of cyclical category buying. *Marketing Science*, 30(3):469–480.

Stigler, G. J. and Becker, G. S. (1977). De gustibus non est disputandum. *The American Economic Review*, 67(2):76–90.

Swait, J. and Erdem, T. (2007). Brand effects on choice and choice set formation under uncertainty. *Marketing Science*, 26(5):679–697.

Terui, N. and Ban, M. (2014). Multivariate structural time series models with hierarchical structure for over-dispersed discrete outcome. *Journal of Forecasting*, 33(5):376–390.

Terui, N., Ban, M., and Allenby, G. (2011). The effect of media advertising on brand consideration and choice. *Marketing Science*, 30(1):384–391.

Terui, N., Ban, M., and Maki, T. (2010). Finding market structure by sales count dynamics – multivariate structural time series models with hierarchical structure for count data. *Annals of the Institute of Statistical Mathematics*, 62(91):92–107.

Terui, N. and Dahana, W. (2006). Estimating heterogeneous price thresholds. *Marketing Science*, 25(4):384–391.

Trivedi, M., Bass, F. M., and Rao, R. C. (1994). A model of stochastic variety-seeking. *Marketing Science*, 13(3):274–297.

Vakratsas, D., Feinberg, F. M., Bass, F. M., and Kalyanaram, G. (2004). The shape of advertising response functions revisited: A model of dynamic probabilistic thresholds. *Marketing Science*, 23(1):109–119.

Vilcassim, N. J. and Jain, D. C. (1991). Modeling purchase-timing and brand-switching behavior incorporating explanatory variables and unobserved heterogeneity. *Journal of Marketing Research*, 28(1):29–41.

Waldfogel, J. (2003). Preference externalities: an empirical study of who benefits whom in differentiated-product markets. *RAND Journal of Economics*, 34(3):557–568.

West, M. and Harrison, J. (1996). *Bayesian Forecasting and Dynamic Models*. Springer: New York, 2nd edition.

Wu, J., Shi, M., and Hu, M. (2015). Threshold effects in online group buying. *Management Science*, 61(9):2025–2040.

APPENDIX

A Product Descriptions

	Product	Incidence	Quantity	Corner Solutions		Interior Solutions	
1	Yoplait Original Strawberry	383	680	90	(0.23)	293	(0.77)
2	Yoplait Original Mountain Berry	392	819	80	(0.20)	312	(0.80)
3	Yoplait Original Mixed Berry	432	886	111	(0.26)	321	(0.74)
4	Yoplait Original Key Lime Pie	383	824	145	(0.38)	238	(0.62)
5	Yoplait Original French Vanilla	330	744	116	(0.35)	214	(0.65)
6	Yoplait Light Blackberry	327	604	26	(0.08)	301	(0.92)
7	Yoplait Light Vanilla	260	714	64	(0.25)	196	(0.75)
8	Yoplait Light Red Raspberry	285	491	25	(0.09)	260	(0.91)
9	Yoplait Light Blueberry	318	656	31	(0.10)	287	(0.90)
10	Yoplait Light Peach	269	482	20	(0.07)	249	(0.93)
11	Dannon FOTB Strawberry	147	232	12	(0.08)	135	(0.92)
12	Dannon FOTB Blueberry	228	457	50	(0.22)	178	(0.78)
13	Dannon FOTB Peach	121	239	26	(0.21)	95	(0.79)
14	Dannon FOTB Mixed Berry	170	295	29	(0.17)	141	(0.83)
15	Colombo Classic Blackberry	351	619	76	(0.22)	275	(0.78)
16	Stonyfield Farm Vanilla	122	247	29	(0.24)	93	(0.76)
17	Stonyfield Farm Strawberry	156	218	13	(0.08)	143	(0.92)
18	Stonyfield Farm Peach	188	302	18	(0.10)	170	(0.90)
19	Stonyfield Farm Raspberry	197	339	21	(0.11)	176	(0.89)
20	Stonyfield Farm Blueberry	201	395	32	(0.16)	169	(0.84)
21	Dannon FOTB Raspberry	145	247	21	(0.14)	124	(0.86)
22	Dannon L&F Strawberry Banana	302	522	46	(0.15)	256	(0.85)
23	Dannon L&F Blueberry	462	790	102	(0.22)	360	(0.78)
24	Dannon L&F Raspberry	414	700	81	(0.20)	333	(0.80)
25	Dannon L&F Strawberry	319	537	59	(0.18)	260	(0.82)
26	Dannon L&F Vanilla	402	1014	177	(0.44)	225	(0.56)
27	Colombo Classic Vanilla	358	974	134	(0.37)	224	(0.63)
28	Colombo Classic Blueberry	503	1223	138	(0.27)	365	(0.73)
29	Colombo Classic Peach	387	763	64	(0.17)	323	(0.83)
30	Colombo Classic Strawberry	453	866	77	(0.17)	376	(0.83)
31	PL Blueberry	453	975	55	(0.12)	398	(0.88)
32	PL Cherry	462	982	87	(0.19)	375	(0.81)
33	PL Peach	423	697	41	(0.10)	382	(0.90)
34	PL Raspberry	386	674	39	(0.10)	347	(0.90)
35	PL Strawberry	484	1004	60	(0.12)	424	(0.88)
	TOTAL	11213	22211	2195	(0.20)	9018	(0.80)

Notes: FOTB = Fruit on the Bottom, L&F = Light & Fit, PL = Private Label

B MCMC Algorithm

The following outlines the steps of a Metropolis-within-Gibbs sampler used to generate draws from the posterior distribution of model M4. The algorithms for the other three models are nested versions of the algorithm below.

For each iteration:

For household $h = 1, \ldots, H$ (do in parallel):

For trip $t = 1, \ldots, T_h$:

1. Update attribute coefficients $\boldsymbol{\alpha}_{ht}$ and $\boldsymbol{\beta}_{ht}$ (Metropolis-Hastings step)
   Let $P(\boldsymbol{x}_{ht})$ denote the likelihood function for consumer $h = 1, \ldots, H$ at purchase time $t = 1, \ldots, T_h$ given in (7). Then the target posteriors are:

$$p(\boldsymbol{\alpha}_{ht}|\text{else}) \propto |\Sigma|^{-1/2} \exp\left\{-\frac{1}{2}(\boldsymbol{\alpha}_{ht} - \boldsymbol{\alpha}_h - \boldsymbol{a}f_{ht})'\Sigma^{-1}(\boldsymbol{\alpha}_{ht} - \boldsymbol{\alpha}_h - \boldsymbol{a}f_{ht})\right\} P(\boldsymbol{x}_{ht}) \tag{B.1}$$

$$p(\boldsymbol{\beta}_{ht}|\text{else}) \propto |\Omega|^{-1/2} \exp\left\{-\frac{1}{2}(\boldsymbol{\beta}_{ht} - \boldsymbol{\beta}_h - \boldsymbol{b}\boldsymbol{g}_{ht})'\Omega^{-1}(\boldsymbol{\beta}_{ht} - \boldsymbol{\beta}_h - \boldsymbol{b}\boldsymbol{g}_{ht})\right\} P(\boldsymbol{x}_{ht}) \tag{B.2}$$

   For each set of parameters, we use random-walk proposals of the form:

$$\boldsymbol{\beta}_{ht}^{(r)} \sim N(\boldsymbol{\beta}_{ht}^{(r-1)}, k_\beta I_p). \tag{B.3}$$

$$\boldsymbol{\alpha}_{ht}^{(r)} \sim N(\boldsymbol{\alpha}_{ht}^{(r-1)}, k_\alpha I_p). \tag{B.4}$$

   The two step sizes above (k_α, k_β) are adaptively tuned during the burn-in iterations to achieve acceptance rates of 25%.

2. Update intercepts $\boldsymbol{\alpha}_h$ and $\boldsymbol{\beta}_h$ (Gibbs step)

$$\boldsymbol{\alpha}_h|\text{else} \sim N\Big(V_\alpha^{-1}\bar{\boldsymbol{\alpha}} + I_p(\boldsymbol{\alpha}_{ht} - \boldsymbol{f}'_h\boldsymbol{a}), I_p\Big) \tag{B.5}$$

$$\boldsymbol{\beta}_h|\text{else} \sim N\Big(V_\beta^{-1}\bar{\boldsymbol{\beta}} + I_p(\boldsymbol{\beta}_{ht} - \boldsymbol{g}'_h\boldsymbol{b}), I_p\Big) \tag{B.6}$$

3. Update latent factors $\boldsymbol{f}_h$ and $\boldsymbol{g}_h$ (Forward-Filtering Backward-Sampling)
   We reformulate the measurement equation as

$$\begin{pmatrix} \boldsymbol{\alpha}_{ht} - \boldsymbol{\alpha}_h \\ \boldsymbol{\beta}_{ht} - \boldsymbol{\beta}_h \end{pmatrix} = \begin{pmatrix} \boldsymbol{a} & 0 \\ 0 & \boldsymbol{b} \end{pmatrix} \begin{pmatrix} f_{ht} \\ \boldsymbol{g}_{ht} \end{pmatrix} + \begin{pmatrix} \boldsymbol{e}_{ht} \\ \boldsymbol{u}_{ht} \end{pmatrix}, \quad \begin{pmatrix} \boldsymbol{e}_{ht} \\ \boldsymbol{u}_{ht} \end{pmatrix} \sim N\left(0, \begin{pmatrix} \Sigma & 0 \\ 0 & \Omega \end{pmatrix}\right) \tag{B.7}$$

as well as the system equation as

$$\begin{pmatrix} f_{ht} \\ \boldsymbol{g}_{ht} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ k_{ht}\boldsymbol{\delta}_h & (1-k_{ht})I_k \end{pmatrix} \begin{pmatrix} f_{ht-1} \\ \boldsymbol{g}_{ht-1} \end{pmatrix} + \begin{pmatrix} \nu_{ht} \\ \boldsymbol{\eta}_{ht} \end{pmatrix}, \quad \begin{pmatrix} \nu_{ht} \\ \boldsymbol{\eta}_{ht} \end{pmatrix} \sim N\left(0, \begin{pmatrix} 1 & 0 \\ 0 & k_{ht}I_k \end{pmatrix}\right) \tag{B.8}$$

where $k_{ht} = 1(f_{ht-1} \geq r_h)$. We use the method of Carter and Kohn (1994) to allow for time-varying coefficients in the state space model described above.

4. Update switching coefficients $\boldsymbol{\delta}_h$ (Gibbs step)

$$\delta_{h\ell}|\text{else} \sim N\Big((\tilde{\boldsymbol{f}}'_h\tilde{\boldsymbol{f}}_h + 1)^{-1}(\tilde{\boldsymbol{f}}'_h\tilde{\boldsymbol{g}}_{hd} + \bar{\delta}_\ell), (\tilde{\boldsymbol{f}}'_h\tilde{\boldsymbol{f}}_h + 1)^{-1}\Big) \qquad \ell = 1, \ldots, k \tag{B.9}$$

Here $\tilde{\boldsymbol{f}}_h$ and $\tilde{\boldsymbol{g}}_{h\ell}$ are both factor vectors collected in the case of $f_{ht-1} \geq r_h$. If $f_{ht-1} < r_h$ (i.e., $k_{ht} = 0$ for all t), then the posterior is $\delta_{h\ell} \sim N(\bar{\delta}_\ell, 1)$ by homogeneity. Hyperparameters are set to $\bar{\delta}_\ell = 0$ for all dimensions $\ell = 1, \ldots, k$.

---

5. Update factor loadings $\boldsymbol{a}$ and $\boldsymbol{b}$ (Gibbs step)

Satiation: Define $\boldsymbol{f}_h = (f_{h1}, \ldots, f_{hT_h})'$ to be a T_h-dimensional vector and let $\boldsymbol{f} = (\boldsymbol{f}'_1, \ldots, \boldsymbol{f}'_H)'$ be a $(\sum_h T_h)$-dimensional vector. Similarly, define $\boldsymbol{\alpha}^*_{hj} = (\alpha_{hj1} - \alpha_{hj}, \ldots, \alpha_{hjT_h} - \alpha_{hj})'$ and $\boldsymbol{\alpha}^*_j = (\boldsymbol{\alpha}^{*'}_{1j}, \ldots, \boldsymbol{\alpha}^{*'}_{Hj})'$. The posterior for a_j can then be derived from a normal regression equation with response variable $\boldsymbol{\alpha}^*_j$ and explanatory variable $\boldsymbol{f}$.

$$a_j|\text{else} \sim N\Big((A_0^{-1} + \sigma_j^{-2}\boldsymbol{f}'\boldsymbol{f})^{-1}(A_0^{-1}a_0 + \sigma_j^{-2}\boldsymbol{f}'\boldsymbol{\alpha}^*_j), (A_0^{-1} + \sigma_j^{-2}\boldsymbol{f}'\boldsymbol{f})^{-1}\Big) \qquad j = 1, \ldots, p \tag{B.10}$$

Hyperparameters are set to $a_0 = 0$ and $A_0 = 100$.

Baseline utility: As above, we define $\boldsymbol{g}_h = (\boldsymbol{g}_{h1}, \ldots, \boldsymbol{g}_{hT_h})'$ to be a $T_h \times k$ matrix, $\boldsymbol{g} = (\boldsymbol{g}'_1, \ldots, \boldsymbol{g}'_H)'$ a $(\sum_h T_h) \times k$ matrix, $\boldsymbol{\beta}^*_{hj} = (\beta_{hj1} - \beta_{hj}, \ldots, \beta_{hjT_h} - \beta_{hj})'$ a T_h-dimensional vector, and $\boldsymbol{\beta}^*_j = (\boldsymbol{\beta}^{*'}_{1j}, \ldots, \boldsymbol{\beta}^{*'}_{Hj})'$ a $(\sum_h T_h)$-dimensional vector. The posterior of $\boldsymbol{b}_j$ can then be written using normal conjugacy results.

$$\boldsymbol{b}_j|\text{else} \sim N\Big((B_0^{-1} + \omega_j^{-2}\boldsymbol{g}'\boldsymbol{g})^{-1}(B_0^{-1}\boldsymbol{b}_0 + \omega_j^{-2}\boldsymbol{g}'\boldsymbol{\beta}^*_j), (B_0^{-1} + \omega_j^{-2}\boldsymbol{g}'\boldsymbol{g})^{-1}\Big) \qquad j = 1, \ldots, p \tag{B.11}$$

Hyperparameters are set to $\boldsymbol{b}_0 = 0$ and $B_0 = 100I_k$.

6. Update population-level attribute means $(\bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{\beta}})$ and variances (V_α, V_β) (Gibbs step)

Satiation: Define $\boldsymbol{\alpha}_j = (\alpha_{j1}, \ldots, \alpha_{jH})$.

$$\bar{\alpha}_j|\text{else} \sim N\Big((\mathbf{1}'_H\mathbf{1}_H + v_{\alpha 0}^{-2})^{-1}(\mathbf{1}'_H\boldsymbol{\alpha}_j + v_{\alpha 0}^{-2}\bar{\alpha}_0), (\mathbf{1}'_H\mathbf{1}_H + v_{\alpha 0}^{-2})^{-1}\Big) \qquad j = 1, \ldots, p \tag{B.12}$$

Hyperparameters are set to $\bar{\alpha}_0 = 0$ and $v^2_{\alpha 0} = 100$.

$$v^2_{\alpha j}|\text{else} \sim IG\left(a_{\alpha 0} + \frac{H}{2}, b_{\alpha 0} + \frac{1}{2}\sum_{h=1}^{H}(\alpha_{jh} - \bar{\alpha}_j)^2\right) \qquad j = 1, \ldots, p \tag{B.13}$$

Hyperparameters are set to $a_{\alpha 0} = b_{\alpha 0} = 0.5$.

Baseline utility: Define $\boldsymbol{\beta}_j = (\beta_{j1}, \ldots, \beta_{jH})$.

$$\bar{\beta}_j|\text{else} \sim N\Big((\mathbf{1}'_H\mathbf{1}_H + v^{-2}_{\beta 0})^{-1}(\mathbf{1}'_H\boldsymbol{\beta}_j + v^{-2}_{\beta 0}\bar{\beta}_0), (\mathbf{1}'_H\mathbf{1}_H + v^{-2}_{\beta 0})^{-1}\Big) \qquad j = 1, \ldots, p \tag{B.14}$$

Hyperparameters are set to $\bar{\beta}_0 = 0$ and $v^2_{\beta 0} = 100$.

$$v^2_{\beta j}|\text{else} \sim IG\left(a_{\beta 0} + \frac{H}{2}, b_{\beta 0} + \frac{1}{2}\sum_{h=1}^{H}(\beta_{jh} - \bar{\beta}_j)^2\right) \qquad j = 1, \ldots, p \tag{B.15}$$

Hyperparameters are set to $a_{\beta 0} = b_{\beta 0} = 0.5$.

7. Update population-level switching parameters $\bar{\boldsymbol{\delta}}$ (Gibbs step)

$$\bar{\delta}_\ell|\text{else} \sim N\left((H + v_0^{-2})^{-1}\Big(\sum_h \delta_{h\ell} + v_0^{-2}\bar{\delta}_0\Big), (H + v_0^{-2})^{-1}\right) \qquad \ell = 1, \ldots, k \tag{B.16}$$

Hyperparameters are set to $\bar{\delta}_0 = 0$ and $v_0^2 = 100$.