Diversification in Multi-Product Choices: Bias or Rational Utility Maximization?

Edward J. Fox[1*], Hristina Pulgar[2] and John H. Semple[3]

[1*]Department of Marketing, Cox School of Business, 6212 Bishop Blvd., Dallas, Texas, US.

[2]formerly JCPenney Center for Retail Excellence, Department of Marketing, Cox School of Business, 6212 Bishop Blvd., Dallas, Texas, US.

[3]Department of Information Technology and Operations Management Department of Marketing, Cox School of Business, 6212 Bishop Blvd., Dallas, Texas, US.

*Corresponding author. E-mail: efox@smu.edu;
Contributing authors: hristinadishkova@gmail.com; jsemple@smu.edu

Abstract

This paper tests a theory of rational multi-product choice (*RMPC*) using empirical evidence from a large-scale choice experiment, two smaller longitudinal choice experiments, and multi-market panel data. Multi-product choice involves the selection of a set of substitutable products, where "set" incorporates both the product alternatives included and the number of units of each. Derived from a basic random utility framework, *RMPC* theory predicts that multi-product choices will reflect a tradeoff between the intrinsic utility of products in a set and the consumption flexibility afforded by the structure of that set. Data from the large-scale choice experiment show that modeling this tradeoff offers superior predictions compared to other models for 3-, 4-, and 5-product choice sets in two different categories. Interestingly, we also find that the modal variety is exactly three product alternatives for all choice set sizes and categories studied. Data from the longitudinal choice experiments provide support for *RMPC* theory's prediction that consumption choice probabilities will be proportional to the inventories of available product alternatives. While consumption choices are found to be consistent with *RMPC* theory, they are not consistent with variety seeking. Finally, *RMPC* theory predicts that the variety in a consumer's multi-product choice set will be higher for lower consumption rates, and lower for higher consumption rates. Evidence from panel data of yogurt purchases supports this proposition. Taken in its entirety, the empirical evidence presented in this paper confirms that multi-product choices, in particular the product diversification previously attributed to bias, are actually consistent with rational utility maximization.

1 Introduction

> "Sometimes you feel like a nut, sometimes you don't" - lyric from 1970s/1980s advertising jingle and tagline for Peter Paul's Mounds and Almond Joy candy bars

Marketing practitioners have long recognized that consumers' tastes reflect uncertainty over time. In the case of hedonic products, consumers' long run preferences may be stable over time, yet exhibit considerable variation from one consumption occasion to the next. This element of randomness was dramatized in an iconic ad campaign for Almond Joy and Mounds candy bars. As the jingle suggests, ccnsumers are aware that their future preferences are uncertain...so they are better off having a choice.

When shopping for products for future consumption, it is rational for consumers to consider this preference uncertainty. For example, a consumer shopping for soup might buy multiple cans of vegetable soup (their favorite) but also include other alternatives that they prefer on occasion; for example, chicken soup for when they are not feeling well or tomato soup to pair with a grilled cheese sandwich. Consumers therefore often buy multiple products (in different quantities) in a category, which they store at home to be consumed later.

The primary finding in the literature addressing the selection of multiple products in a category is that consumers consistently include "too much" variety when choosing a set for future consumption (Simonson, 1990; Read and Loewenstein, 1995). This finding is known as *diversification bias* and is based on the empirical regularity that consumers include more variety (i.e., more different product alternatives) when choosing a set of products for future consumption—known as *simultaneous choice*—compared to products chosen one-at-a-time on each consumption occasion—known as *sequential choice.* Strangely, the experimental evidence of diversification bias imposed an extraneous requirement on simultaneous choice. Participants were required to precommit to the exact order in which all products in their set would be consumed. In actuality, consumers can choose any can of soup they have at home...or none of them (the same is true of any product category). This unnatural precommitment requirement suppresses the impact of future preference uncertainty and reduces simultaneous choice to a forecasting exercise. Further, Reed and Lowenstein found that 44% of their simultaneous choice participants wanted to change the precommitted order during the consumption sequence, evidence that the requirement prevented them from accommodating their true preferences once uncertainty is resolved, as they would naturally. Because imposing such an arbitrary, utility-reducing requirement is inconsistent with actual behavior, neither our utility-maximizing models nor our empirical tests impose any restrictions on consumption sequences. To make this distinction clear, we refer to a set selected without any consumption precommitment restrictions as *multi-product-choice.*

Diversification bias refers to the product variety in simultaneous choice compared to the variety across an equal number of sequential choices. Each sequential choice involves the selection of a product— any product—from the full assortment, so it is not restricted by the simultaneous choice set nor previous consumption choices from that set.

In contrast, multi-product choice involves the construction of a set of products from which subsequent consumption choices can be made. This set is then reduced by one unit after each consumption choice, leaving fewer products for subsequent consumption choices. Because previous consumption choices impact subsequent choices, forward-looking dynamic models are appropriate for multi-product choice. The theory we test relies on such dynamic models (due to Walsh 1995 and Fox, Norman, and Semple 2018), which predict that including product variety in a multi-product choice set can be a rational response to future preference uncertainty. For expositional simplicity, we will abbreviate rational multi-product choice as "*RMPC*." In this paper, we assess whether the observed sets selected by consumers are consistent with this dynamic model of expected utility theory and thus provide an alternative explanation to diversification bias.

To develop a basic intuition for *RMPC* theory, assume a universe of only three products, A, B, and C, and assume a consumer who is indifferent between these three products; i.e., on any given consumption occasion, they prefer each with probability $\frac{1}{3}$. If this consumer is asked to select a set of three products, which set should they select? *RMCP* theory shows that the set maximizing their expected utility is the set A, B, C i.e., one of each[1]. This result is intuitive, and can be proven based on the analysis in §3.1. Let us now suppose that the same consumer chooses a product from the full assortment on three sequential consumption occasions (sequential choice). There are 27 equally probable choice sequences, of which 21 have a variety of two or fewer product alternatives (21/27 =78%). In this scenario, a utility-maximizing consumer who selects three products sequentially will exhibit less variety than their utility-maximizing multi-product set 78% of the time. This example demonstrates that the variety included in sequential choices cannot be used to reliably evaluate the rationality of multi-product choices.

RMPC theory is based on a simple random utility specification. It was introduced by Walsh (1995), who analyzed the case of a two-product set chosen for future consumption. His analysis was extended by Fox, Norman, and Semple (2018) to the case of a consumer choosing a set of n products, then consuming them one-at-a-time until the set is exhausted. Fox, Norman, and Semple's *base RMPC* model (which they called the canonical model) assumes that [i] a product is chosen on each of n consecutive consumption occasions, [ii] the future utility of each remaining product in the set is uncertain due to a stochastic error component, and [iii] the consumer's goal at each consumption occasion is to maximize the current utility of the product selected plus the expected future utility from consuming the remaining set. The *general RMPC* model (which Fox, Norman, and Semple similarly termed the generalized model) relaxes the assumption that a product is selected from the set on every consumption occasion by introducing an outside option. This general model is applicable when different categories of products compete for the consumer's attention, and so a product need not be chosen from the set on every consumption occasion. It is particularly relevant when consumers have different usage rates. As noted previously, the sequence in which products are chosen in the consumption stage is not fixed, so consumers are free to choose whichever available product maximizes their current utility plus expected future utility.

Together, the base and general *RMPC* models imply three testable propositions. These propositions are modified here so as to be understandable without notation.

1. Consumers' multi-product choices will reflect a tradeoff between the intrinsic utility of products in the set and the consumption flexibility afforded by the structure of that set
2. When consuming a multi-product set, the probability of selecting a product alternative for consumption will be proportional to its current inventory
3. For a given set size, the multi-product choices of consumers with lower consumption rates will include more variety than the multi-product choices of consumers with higher consumption rates

To the best of our knowledge, the consumption flexibility identified in the first proposition has not been studied previously. Yet modeling the tradeoff between consumption flexibility and intrinsic utility is key to determining the value a consumer would expect from a set of products, and therefore to predicting multi-product choices. Modeling that tradeoff could assist in the design and pricing of multi-product packs—both fixed "variety packs" and customizable options such as "build your own 6-pack." In addition, modeling the first proposition's tradeoff offers a precise method for determining *the value of variety* in multi-product choice. These are arguably our paper's most important contributions.

The second proposition relates to the consumption choices made from a given set. It enables us to determine whether consumption choices are consistent with forward-looking utility maximization. If consumers are not forward looking, we would expect them to follow a myopic consumption

[1] See Fox, Norman, and Semple (2018) for proof of optimality in n-product choice sets.

policy and select whichever product available in their set has the highest utility. In contrast, the second proposition's dynamic stochastic consumption policy, which is implied by both *RMPC*'s base and general models (see proof in Appendix A), depends instead on the current inventory of each product alternative in the set. This consumption policy accommodates future preference uncertainty by probabilistically matching products in the set to the most opportune occasions on which to consume them. *RMPC* theory predicts that consumption choices will appear, at least probabilistically, to be proportional to a product's current inventory.

The third proposition relates a product category's consumption rate to the product variety included in multi-product choice sets. This proposition comes from the general *RMPC* model which features an outside option. Generally speaking, more variety in the multi-product choice set is needed to compete with a more attractive outside option (implied by a lower consumption rate). We formalize this intuition in §3.2. To our knowledge, no other theory addresses the relationship between usage rate and choice set variety nor has that relationship previously been investigated.

Together, these three propositions represent the predictions of *RMPC* theory, a rational theory of multi-product choice derived from the canonical random utility model. Though individual propositions may have alternative explanations, empirical support for all three propositions together provides strong evidence of rationality in multi-product choices. Indeed, a large-scale multi-product choice experiment, two longitudinal consumption choice experiments, and multi-market panel data provide that empirical support. The evidence suggests that consumers are not biased but instead act rationally in the face of future preference uncertainty. Consumers may not be able to forecast when they will want a specific product, but they know...sometimes you feel like a nut, sometimes you don't.

2 Literature Review

Three streams of literature have addressed consumers' choices of multiple substitutable products. Research in consumer psychology has compared the product variety in a simultaneous choice set chosen for future consumption with the variety of products in sequential choices made on successive consumption occasions. Recall that the primary finding in this literature is that simultaneous choices generally include more variety (i.e., more product alternatives) than the same number of sequential choices (e.g., Simonson, 1990; Read and Loewenstein, 1995). Researchers argued that this diversification bias (cf. Read and Loewenstein 1995) was due to poor forecasting (Kahnemann and Snell, 1992 p.304); more specifically, overestimating the probability of satiating on one's favorite products over time (Simonson, 1990; Read and Loewenstein, 1995; Kahn and Ratner, 2005).

Simonson (1990) was the first to document diversification bias, finding that people systematically choose a greater variety of product alternatives in simultaneous choices for future consumption than when choosing sequentially at the time of consumption. In a series of experiments, two groups of participants were asked to choose three snacks. In one group, participants were allowed to select an item on each consumption occasion from the full universe of possible snacks—sequential choice. In the second group, participants were asked to first select a set of three snacks and commit to the precise order in which they would be consumed—simultaneous choice. Those in the simultaneous choice group selected greater variety than those in the sequential choice group.

Read and Lowenstein (1995) confirmed and extended Simonson's findings across several experiments, one of which was designed to test whether the greater variety in simultaneous choice is due to preference uncertainty. They attempted to eliminate future preference uncertainty by allowing participants in one of the simultaneous choice cells to pre-taste each of the six well-known snacks. They found no difference in the amount of variety selected across simultaneous choice cells, concluding that preference uncertainty could not be the cause of diversification bias. Read and Loewenstein also found that participants assigned the simultaneous choice obtained less utility than those assigned the sequential choice task. In addition, Read and Loewenstein investigated whether diversification is actually a bias, or is consistent with rational utility maximization. They identified two sources of bias: (i) the tendency to overestimate the time between consumption occasions which causes people to overestimate satiation, and (ii) mental bracketing induced by simultaneous choice (combining multiple choices into one), which causes people to mistakenly choose a portfolio of products. Other studies of diversification bias also used the portfolio metaphor, comparing simultaneous choice to the selection of a stock portfolio to hedge against future uncertainty (Simonson 1990, Kahn and Ratner 2005).

Diversification bias implies that, as the shopper purchases for more future consumption occasions, the variety of product alternatives selected will also increase. Simonson and Winer (1992) tested this implication using scanner panel data for the yogurt category, finding a positive relationship between the number of products and the variety of flavors purchased.

Though Read and Loewenstein (1995) did not find diversification to be be a rational hedge against uncertainty, Salisbury and Feinberg (2008) used simulation studies to show that the extent of diversification in multiproduct choice should depend on the level of future preference uncertainty, along with the relative attractiveness of product alternatives and uncertainty about their attractiveness.

Recent econometric studies have considered the purchase of multiple products in the same category, which may vary by brand, flavor, etc. These studies modified existing discrete choice models to accommodate multiple products, with the objective of investigating price and promotion response in a multi-product context. Dube (2004) assumed that shoppers' purchases would be consumed over an unknown number of future consumption occasions. To accomodate diverse multi-product purchases, he assumed that the consumption utility for each product is concave and monotonically increasing in quantity. The resulting model was applied to carbonated soft drink purchase data. Kim, Allenby and Rossi (2002) developed a demand model based on an additive utility structure with a mixed distribution of continuous density and probability mass points. They then used this econometric approach to address retailers' assortment/pricing tradeoffs. Richards, Gomez and Pofahl (2012) applied a slightly different model that accommodated diverse multi-product purchases with a satiation parameter, implying that consumers prefer variety when buying for future consumption. They applied this model to fresh produce, specifically different varieties of apples. Lee and Allenby (2014) derived a model that incorporated differences in package size, in addition to brand and flavor variety. To investigate issues with the estimation of this model, they applied it to simulated data and to yogurt purchase data. They found that ignoring the complexity of diverse multi-product purchases leads to biased parameter estimates and improper attribution of many zero purchase quantities. What these econometric studies have in common is a decreasing marginal utility specification so that every unit of a product alternative that is purchased decreases the utility of buying another unit on the same trip. Kim, Allenby, and Rossi (2002, p. 231) argued that "it is important for the utility function to have diminishing marginal returns to capture satiation as we model demand situations where more than one unit is purchased and consumed." Satiation is the antecedent of variety seeking, so the specifying decreasing marginal return specification implies that multi-product purchases are explained by variety seeking.

Walsh (1995) proposed and analyzed the first dynamic model of shopping and consumption. In that model, the consumer chooses n total products for future consumption from two product alternatives. Walsh showed that it may not be optimal for consumers to select exclusively their favorite alternative. They may be better off also choosing a smaller quantity of the less preferred alternative, even in the absence of variety-seeking. He also showed that adding a unit to a set added more than the unit's expected utility to the value of that set. Further, Walsh found that it is optimal to consume strategically, with consumption choice probabilities depending on the consumer's inventories of the two product alternatives.

Fox, Norman, and Semple (2018) extended Walsh's dynamic model of shopping and consumption to more than two product alternatives. Assuming that preference uncertainty is Gumbel-distributed, they derived a closed-form value function for two cases—a base case without an outside option and a more general case with an outside option, thereby allowing for different consumption rates. They determined that the value forward-looking consumers enjoy from a multi-product set is a function of both the expected utilities of products in that set and the consumption flexbility that set's structure provides. Further, they found that the implied optimal consumption policy, also available in closed form, depends on the current inventory of available product alternatives.

What the relevant literatures for this study have in common is that they address the contemporaneous choice of multiple product alternatives for future consumption. Of course, there are vast and varied literatures addressing the choice of a single product (vs. multiple products); some of those focus on choices over time (e.g., variety seeking, habit persistence). Our empirical analysis will therefore include a test for variety seeking in consumption choices.

3 *RMPC* Theory

RMPC theory is based on the analyses of Walsh (1995) and Fox, Norman, and Semple (2018). We borrow notation from Fox, Norman, and Semple in this section. More specifically, we will lay out two related models: [i] the *base model*, which defines consumption occasions based on actual consumption of an item from the chosen set, and [ii] the more *general model*, which includes an outside option to allow for variation in usage rates. We introduce our notation as follows. In the shopping phase, n units are chosen from the store's full assortment of M product alternatives in the category. The n products chosen in a set can be represented by the integer vector (k_1, k_2, ..., k_M) with $k_i \geq 0$ and $\sum_{i=1}^{M} k_i = n$. Importantly, multiple units of a product alternative may be selected. Consumer and category subscripts are omitted for expositional clarity.

The theory is based on a simple random utility formulation. The utility of any product alternative i ($i = 1$, 2, ..., M) is assumed to be the sum of a consumer-specific deterministic component (U_i) and a random component (ϵ_{it}) for alternative i on the t^{th} consumption occasion. The deterministic component reflects the consumer's long-run preference for that product and so is time invariant. We assume without loss of generality that alternatives have been ordered and subscripted so that $U_1 \geq U_2 \cdots \geq U_M$. The random component captures the consumer's preference uncertainty; it changes with each consumption occasion and is revealed immediately before consumption. The ϵ_{it} are assumed to be independent across product alternatives and time. After a suitable translation of the ϵ_{it} and U_i, one can assume each ϵ_{it} has zero mean and each U_i represents the *expected utility* of alternative i.

3.1 Base Model

For the base model, we assume that a product from the choice set is selected on each consumption occasion. To illustrate the dynamics, suppose that there are $M = 4$ alternatives and that a consumer has selected the set of $n = 3$ products $(2, 1, 0, 0)$; that is, two units of their favorite alternative and one unit of their second-favorite. On the first consumption occasion, the consumer can select a unit of alternative 1 (favorite) or a unit of alternative 2 (second-favorite). The current utility associated with each choice is $U_1 + \epsilon_{11}$ and $U_2 + \epsilon_{21}$, respectively (recall the current period errors are observed immediately before consumption). A myopic consumer would select alternative 1 if $U_1 + \epsilon_{11} > U_2 + \epsilon_{21}$ and select alternative 2 if $U_2 + \epsilon_{21} > U_1 + \epsilon_{11}$ (ties can be broken arbitrarily). However, a forward-looking consumer would consider both the current utility and the *expected future utility*. Letting $V(q)$ represent the value of expected total future utility for a vector of quantities q, the strategic consumer chooses the alternative that maximizes $U_1 + \epsilon_{11} + V(1,1,0,0)$ vs. $U_2+\epsilon_{21}+V(2,0,0,0)$. Note that the future values are different because they incorporate different reductions in future inventory. The same "current utility" plus "expected future utility" comparison is done at each subsequent consumption occasion. In general, the hard part is determining a manageable expression for the expected future utility, or "value," function V.

This framework necessarily abstracts shopping and consumption behavior. For example, the total number of products, n, is assumed to be exogenous.[2] Clearly, factors such as trip type (major versus fill in, cf. Kollat and Willet, 1967) and incentives for multiple purchases could affect n. Also, the consumer's deterministic component of utility could be a function of store-specific factors such as price, or time-varying factors such as satiation. Introducing these complexities would not only greatly complicate the modeling effort, but would also obscure the basic insights on which our propositions are based.

If the random errors are assumed to follow a standard (zero-mean) Gumbel distribution—as one does in the classical logit choice framework—then we can obtain some compact structural results for V. Assuming that the consumer follows an optimal consumption strategy (described shortly), then the value $V(k_1, k_2, ..., k_M)$ that consumer obtains from an arbitrary n-pack $(k_1, k_2, ..., k_M)$ is given by the formula

$$V(k_1, k_2, ..., k_M) = \left[\sum_{i=1}^{M} k_i U_i\right] + \left[ln(n!) - \sum_{i=1}^{M} ln(k_i!)\right] \tag{3.1}$$

A previously unpublished induction proof is in Appendix A. The value function consists of two distinct components, shown in square brackets in (3.1). The first is a linear function of quantities (k_i) and expected utilities (U_i). This component is increased by choosing alternatives that have higher expected utility; i.e., the consumer's favorites. If the consumer had to precommit to a consumption sequence or if the random component was vanishingly small, the linear component would represent the consumer's expected utility for the set. We will refer to this component as the *intrinsic utility* of products in the set. However, if the consumer is free to choose any alternative from the set at the time of consumption (in the presence of future preference uncertainty), then the logarithmic second component in (3.1) must also be present. We will refer to this component as *consumption flexibility*—the incremental value of making consumption choices from a set using all available information at the time of consumption. Observe that this logarithmic component does not depend on the U_i but on the distributional properties of the quantities $(k_1, k_2, ..., k_M)$. In contrast to the linear component in (3.1), the logarithmic component favors variety. The logarithmic component is maximized by selecting a single unit of n different alternatives, which represents the most possible variety.[3]The two components in (3.1) capture the tradeoff between opposing objectives: one favoring intrinsic utility, the other favoring consumption flexibility.[4] Analysis of the general

[2]The assumption that set size is exogenous is common in diversification bias research.

[3]Adding a unit of alternative j to a set with n units increases the expected utility of that set by an amount $U_j + ln(n+1) + ln(k_j) - ln(k_j+1)$. Thus, the incremental benefit of adding a unit exceeds the expected utility of the choice made from those alternatives (McFadden, 1978; Ben-Akiva and Lerman, 1979). This property was demonstrated by Walsh (1995) for the two-alternative case.

[4]Even for non-Gumbel errors, the value function can be shown to consist of a linear component $\sum_{i=1}^{M} k_i U_i$ plus a nonlinear component $f_n(k_1, k_2, ..., k_M)$ that does not depend on the U_i and favors variety. Like the logarithmic com-

model in §3.2 will show that the weighting of these objectives may not be equal. The consumer's utility-maximizing set must balance these two objectives.

Proposition 1. *Consumers' multi-product choices (i.e., set selection) will reflect a tradeoff between the intrinsic utility of products in the set and the consumption flexibility afforded by the structure of that set.*

In order for (3.1) to capture the expected future value of a given set, that set would have to be consumed according to a policy that maximizes its value. Suppose that, on the t^{th} consumption occasion, the set has q_{it} units of alternative i remaining. Then the policy that maximizes the set's future value is to consume the alternative that maximizes $ln\,(q_{it}) + \epsilon_{it}$ (see Appendix A). Observe that this consumption policy incorporates each product's current inventory, (q_{it}) and its random component, ϵ_{it}, but not its deterministic component, U_i. This policy implies that the probability of choosing a product alternative for consumption will be proportional to its current inventory. By favoring consumption of product alternatives with higher inventory levels, this policy generally preserves options through the consumption sequence.

Proposition 2. *The probability that a product alternative is selected for consumption will be proportional to its current inventory.*

More generally, Proposition 2 implies that the probability of selecting a product alternative from one's set is increasing in its inventory. This general property of consumption probability can be shown to apply regardless of the error distribution (see Alptekinoglu and Semple 2018). Of course, alternative theories of consumption from a multi-product set might lead to a similar (or equivalent) proposition. However, this consumption policy and the value function in (3.1) are implied by *RMPC* theory and so are derived from a canonical random utility specification.

3.2 General Model

The general *RMPC* model incorporates an outside option into the framework developed in the previous subsection, which complicates the value function. Suppose that the outside option is "product 0," with expected utility U_0 and error term ϵ_{0t}. If U_0 is large relative to the expected utilities of product alternatives in the set, then the outside option is attractive and the consumption rate for products in the set will be lower. Conversely, if U_0 is small relative to the expected utilities of product alternatives in the set, then the outside option is unattractive and the consumption rate (of products in the set) will be higher. Note that the consumption horizon T is no longer defined to be n periods, because the outside option is inexhaustible and may be consumed in any or all periods.

As before, let the chosen set be represented by the vector of integer quantities $(k_1, k_2, ..., k_M)$ where $k_i \geq 0$, $\sum_{i=1}^{M} k_i = n$. Now define a new set S_T of all possible consumption possibilities $(x_0, x_1, \cdots, x_M)$ over a horizon of T periods by

$$S_T(k_1, k_2, \ldots, k_M) \quad = \quad \left\{ \; (x_0, x_1, \cdots, x_M) : \sum_{i=0}^{M} x_i = T;\, 0 \leq x_0 \leq T;\, 0 \leq x_i \leq k_i \; \right\}.$$

Observe that consumption occasions can be decomposed into two subsets: (i) the number of times that products from the chosen set are consumed, represented by the $(x_1, \cdots, x_M)$ and necessarily

ponent, the function $f_n(k_1, k_2, ..., k_M)$ is minimized by any set having n units of a single alternative and maximized by any set having n distinct alternatives. Unfortunately, the nonlinear component cannot typically be expressed in closed-form for general error distributions (see Alptekinoglu and Semple, 2018).

satisfying $x_i \leq k_i$ and (ii) the number of times that the outside option is consumed, x_0, which must satisfy $x_0 = T - \sum_{i=1}^{M} x_i$. Then

$$V_T(k_1, k_2, \ldots, k_M) \quad = \quad ln \left[\sum_{(x_0, x_1, \ldots, x_M) \in S_T(k_1, k_2, \ldots, k_M)} \frac{T!}{x_0! x_1! \cdots x_M!} e^{\sum_{j=0}^{M} x_j U_j} \right] \qquad (3.2)$$

Note that this value function nests the base model's value function (3.1) when the utility of the outside option $U_0 = -\infty$.

While the general model's value function is somewhat unwieldy, Fox, Norman, and Semple (2018) determined that the set which optimizes (3.2) has at least as much variety as the set which optimizes (3.1). A simple thought experiment provides the intuition. Imagine that a consumer has r units in inventory on the final consumption occasion of their horizon. If $r > 1$, then some unit(s) will go to waste because only one unit may be consumed on that final occasion. This happens when utility has been maximized on any prior consumption occasion(s) by choosing the outside option rather than consuming a unit from the set. In the case of $r > 1$ units in inventory and only one consumption occasion remaining, the optimal set would have r different product alternatives. Maximizing the consumer's options is necessary to maximize their expected utility on that final consumption occasion. This "end of horizon" effect, as it is known in dynamic programming, encourages consumers with low consumption rates to include more variety in their chosen sets than consumers with high consumption rates.

Proposition 3. *For a given set size* n*, the multi-product choices of consumers with lower consumption rates will include more variety than the multi-product choices of consumers with higher consumption rates.*

4 Large-Scale Multi-Product Choice Experiment

In this section, we present evidence from a large-scale experiment designed to determine whether Proposition 1's tradeoff between intrinsic utility and consumption flexibility explains actual multi-product choices.

4.1 Multi-Product Choice Experimental Design

Unlike previous studies of multi-product choice, our between-subjects design varies choice set size. Specifically, participants chose either a 3-, 4-, and 5-product set. While *RMPC* theory does not make materially different predictions depending on choice set size, varying set size is nevertheless important to assess the theory's predictive accuracy. Multi-product choice data was collected for two snack categories—candy and salty-snacks. The category was not manipulated; consumers selected their preferred snack product category. The experiment proceeded as follows.

Data was collected by administering an online questionnaire to participants from a representative panel managed by Dynata, which purports to be "[*t*]he world's largest first-party data platform for insights, activation & measurement" (https://www.dynata.com/). Randomly selected Dynata panelists were invited to participate in the study. Those who opted into the study answered two demographic questions, then were presented with the vending machine scenario: "Imagine that you purchase a single-serving snack from a vending machine twice a week for a full year (more than 100 times). What type of snack would you purchase most often?" The available snack type choices were "candy," "cookies," "healthy snack (granola, trail mix, etc.)," "salty snack," and "none." Only participants who selected "candy" or "salty snack" were allowed to proceed with the questionnaire.

Those participants were presented with an assortment of 12 products from their preferred snack category, assortments of products stocked in local vending machines and pretested in a pilot study. The same product assortment was offered to every participant who selected "candy" or "salty snack," respectively. Presentation order was randomized.

We elicited participants' long-run choice probabilities for products in the assortment by asking them to "[*p*]lease assign a choice percentage to each product, so that they add up to 100 (it's OK to assign a choice percentage of 0 to a product)." These long-run choice probabilities were subsequently used to compute participants' utilities for products in the assortment. Next, participants were then asked to identify:

[i] their favorite product in the assortment: "Of the candies available in the vending machine, which would you say is your favorite?"

[ii] their second favorite product: "Of the candies available in the vending machine, which would you say is your second favorite? (it's OK to pick a product that you might not actually choose during the year)"

[iii] their third favorite product: "Of the candies available in the vending machine, which would you say is your third favorite? (again, it's OK to pick a product that you might not actually choose during the year)"[5]

Based on pretesting, we determined that eliciting more than three ordered favorites was cognitively taxing and yielded unreliable data. These ordinal preferences were used to construct participant-specific multi-product choice sets in a way that makes the empirical analysis tractable.

We then asked a series of questions about category usage rate, attitudes, and perceptions. Neither these responses nor the demographic responses collected previously were analyzed in this paper; however, they are available as additional explanatory variables if required. Next, participants made their multi-product choices. The choice task was substantially different for 3-, 4-, and 5-product sets. Using participant's self-reported favorites, multi-product choices were constrained by requiring that $k_f \geq k_{f+1}$, where k represents the quantity in the choice set and the subscript indicates favorite ordering. Applying this constraint results in three possible 3-product choice sets, five possible 4-product choice sets, and seven possible 5-product choice sets. Table 1 shows the possible choice sets, including choice set notation in parentheses as well as the exact language from the questionnaire. Observe that, for 4- and 5-product choice sets, participants could include a product alternative that was not identified as one of their three favorites. Observe also that, for 4- and 5-product choice sets, different sets may include the same *variety*; defined here as the number of different alternatives available in the set. For example, (3,1,0,0) and (2,2,0,0) both have two product alternatives and therefore the same amount of variety.

4.2 Analysis of Multi-Product Choice Experiment

A sample of 5,140 qualifying completed questionnaires was collected.[6] Questionnaires were then screened to ensure that the product alterative to which the participant assigned the highest long-run choice probability was included among their three favorites. This screen was designed to

[5] Note that participants were prevented from duplicating favorite selections.

[6] Questionnaires were qualified if:

[i] the participant's preferred snack category was either "candy" or "salty snack."

[ii] the participant responded correctly to an attention check question within the questionnaire.

[iii] the participant was not a 'speeder;' i.e., did not complete the questionnaire in less than 1/3 of the median completion time.

Table 1: Multi-Product Choice Sets

3-Product Set

"Now, choose one of the sets of 3 candies/salty snacks below for the next 3 occasions you eat a single-serving candy/salty snack."

(3,0,0)	"A set including 3 [*1st favorite*]"
(2,1,0)	"A set including 2 [*1st favorite*] and 1 [*2nd favorite*]"
(1,1,1)	"A set including 1 [*1st favorite*], 1 [*2nd favorite*], and 1 [*3rd favorite*]"

4-Product Set

"Now, choose one of the sets of 4 candies/salty snacks below for the next 4 occasions you eat a single-serving candy/salty snack."

(4,0,0,0)	"A set including 4 [*1st favorite*]"
(3,1,0,0)	"A set including 3 [*1st favorite*] and 1 [*2nd favorite*]"
(2,2,0,0)	"A set including 2 [*1st favorite*] and 2 [*2nd favorite*]"
(2,1,1,0)	"A set including 2 [*1st favorite*], 1 [*2nd favorite*], and 1 [*3rd favorite*]"
(1,1,1,1)	"A set including 1 [*1st favorite*], 1 [*2nd favorite*], 1 [*3rd favorite*], and 1 other (different) candy/salty snack from the vending machine"

5-Product Set

"Now, choose one of the sets of 5 candies/salty snacks below for the next 4 occasions you eat a single-serving candy."

(5,0,0,0,0)	"A set including 5 [*1st favorite*]"
(4,1,0,0,0)	"A set including 4 [*1st favorite*] and 1 [*2nd favorite*]"
(3,2,0,0,0)	"A set including 3 [*1st favorite*] and 2 [*2nd favorite*]"
(3,1,1,0,0)	"A set including 3 [*1st favorite*], 1 [*2nd favorite*], and 1 [*3rd favorite*]"
(2,2,1,0,0)	"A set including 2 [*1st favorite*], 2 [*2nd favorite*], and 1 [*3rd favorite*]"
(2,1,1,1,0)	"A set including 2 [*1st favorite*], 1 [*2nd favorite*], 1 [*3rd favorite*], and 1 other (different) candy/salty snack from the vending machine
(1,1,1,1,1)	"A set including 1 [*1st favorite*], 1 [*2nd favorite*], 1 [*3rd favorite*], and 2 other (different) candies/salty snacks from the vending machine"

ensure test/retest reliability of participants' long-run preferences. After applying this screen, we analyzed the remaining 4,191 questionnaires.[7]

Table 2: Choice Set Variety

	3-Product Set		4-Product Set		5-Product Set	
	Candy	Salty Snack	Candy	Salty Snack	Candy	Salty Snack
n	413	959	402	972	413	1032
variety						
1	22.8%	19.3%	15.9%	15.3%	16.1%	14.0%
2	28.6%	28.1%	27.0%	28.4%	23.4%	23.0%
3	47.9%	52.5%	39.8%	38.7%	40.5%	45.2%
4			17.4%	17.7%	11.0%	9.2%
5					9.0%	8.6%

Table 2 summarizes the actual variety of the multi-product sets, organized by set size and category. We find that, although maximum variety increases with the choice set size, actual mean variety increases much less, median variety increases for only one category×set size combination, while modal variety does not increase at all. Diversification bias implies that variety will increase with choice set size (Simonson and Winer, 1992). The data do not support this.

Recall Proposition 1's tradeoff between the intrinsic utility of products in the choice set and the consumption flexbility afforded by the choice set's structure. Intrinsic utility is the sum of expected utilities of products in the choice set. We determined each participant's expected utilities $U_j = ln\,(p_j/p_1)$, where their favorites are ordered by the indicator variable j and p_j is the the long-run choice probability of their j^{th} favorite. Consumption flexibility is captured by the expression $ln\,(n!) - \sum_{i=1}^{M} ln\,(k_i!)$ from (3.1).

To assess how well Proposition 1's tradeoff explains multi-product choices, we estimated four multinomial logit [*MNL*] choice models. *MNL* choice models were estimated separately for candy and salty snack categories' 3-, 4-, and 5-product choice sets.[8] For each category×set size combination, we estimate models A through D:

[i] A is a two-parameter model with separate intrinsic utility and consumption flexibility coefficients to allow differential weighting

[ii] B is a one-parameter nested model in which the intrinsic utility and consumption flexibility coefficients are restricted to be equal

[iii] C is a one-parameter nested model in which the consumption flexibility coefficient is restricted to be zero.

[iv] D is a multi-parameter choice set intercepts model.[9]

[7] We checked the robustness of our results to this reliability screen. Specifically, we replicated our choice set analyses both after applying a relaxed screening criterion (eliminating only the few questionnaires for which all three favorites were assigned zero long-run choice probabilities) and after applying two more stringent screens. Regardless of the screening criteria applied, our results were substantially the same.

[8] Recall that the choice set configurations are different for 3-, 4-, and 5-product sets. This prevented us from pooling over set sizes. Further, we estimated separate models for candy and salty snacks to allow for systematic differences between categories.

[9] Regardless of set size, the single-variety set—either (3,0,0), (4,0,0,0), or (5,0,0,0,0)—is the baseline for choice set intercepts model estimation.

There are no clear alternatives in the literature to the tradeoff model A, so we estimate nested models and less parsimonious choice set intercepts models for comparison.[10] Model B's parameter restriction reflects the *RMPC* base case, in which a product is consumed on each consumption occasion. Comparing model C to the tradeoff model A permits us to assess the incremental predictive contribution of consumption flexibility. The choice set intercepts model D incorporates both consumption flexibility and intrinsic utility via set structure, so we would expect it to offer comparable predictive accuracy to the more parsimonious tradeoff model A. Model fits are assessed in sample using *AIC* and *BIC*. Model fits are assessed out of sample using a ten-fold validation. For this validation, we randomly partitioned each dataset into ten equally-sized subsets. Each subset then served as a validation sample, while the other nine were used for estimation. Hit rates were averaged over the ten validation samples.[11]

Table 3 shows fit statistics and parameter estimates for the *MNL* choice models. Models A through D are arranged in vertical panels; categories and set sizes are arranged horizontally. Across the six category×set size combinations, model A explains multi-product choice data better and offers superior predictive accuracy compared to the other three models. In fact, model A dominates the two nested models, B and C, and is clearly superior to the choice set intercepts model D.[12] Thus, the *MNL* choice model fits provide support for Proposition 1 and evidence that consumers' multi-product choices are consistent with a tradeoff between intrinsic utility and consumption flexibility. Note that, across categories and choice set sizes, all model A parameter estimates in the top panel of Table 3 are positive and significant at the 0.001 level. Comparing those parameter estimates, we find that the consumption flexibility parameter estimate is higher than the intrinsic utility parameter in all six category×set size combinations. Testing the difference between the two parameters (incorporating standard errors of the parameter estimates), we find that the difference is statistically significant for five of the six category×set size combinations. The higher weighting of consumption flexibility vis-a-vis intrinsic utility in multi-product choices is consistent with *RMPC*'s more general model, which allows for an outside option. The implication is that consumers do not eat candy or a salty snack on every consumption occasion, but rather consume those snacks less frequently.

We now consider the predictive contribution of consumption flexibility (resulting from the structure of the choice set) after accounting for intrinsic utility by comparing the hit rates of model A and model C. As noted above, the hit rate of model A is higher than model C for every category×set size combination. Across the combinations, the mean increase in hit rate from adding consumption flexibility (i.e., model A compared to model C) is 14.4%—a material improvement in the explanation of multi-product choice. To put that improvement in context, consider that the mean hit rate of model C (without consumption flexibility) is only 25.0%. More importantly, the hit rate of model C *above what would be expected by random chance* is only 2.4%.[13] Finally, we apply model A's parameter estimates to the data for each category×set size combination to see if we recover the aggregate choice set variety patterns. Table 4 shows the choice set varieties predicted by model A with the corresponding actual choice set varieties, as reported in Table 2, in parentheses.

Observe that model A's predictions recover the aggregate patterns of Table 2's actual choice

[10] Note that the econometric models of multi-product choice are estimated using time series purchase data.

[11] In sample hit rates were also computed. They are very similar to hit rates computed in the ten-fold validation and result in the same conclusions.

[12] In five of the six category×set size combinations, model A offers a higher hit rate and lower *AIC* and *BIC* than all other models. For 4-product candy choices, however, the 5-parameter choice set intercepts model D has a lower hit rate (38.8%<39.3%) and slightly lower *AIC* (1201.2<1202.1) compared to model A, but a higher *BIC* (1217.2>1210.1).

[13] The hit rate that would be expected by chance is simply the inverse of the number of alternative sets shown in Table 1.

Table 3: Multi-Product Choice Set Model Estimates

	Choice Set Size = 3		Choice Set Size = 4		Choice Set Size = 5	
	Candy	Salty Snacks	Candy	Salty Snacks	Candy	Salty Snacks
obs =	413	959	402	972	413	1032
Model A – Intrinsic Utility + Consumption Flexibility						
Hit Rate ‡	51.8%	53.7%	38.8%	39.2%	25.9%	26.6%
Log Likelihood	-406.8	-927.5	-599.0	-1448.0	-769.6	-1896.4
AIC	817.5	1859.0	1202.1	2900.0	1543.1	3796.9
BIC	825.6	1868.8	1210.1	2909.7	1551.1	3806.7
Coefficients:						
Intrinsic Utility	0.225 ***	0.201 ***	0.225 ***	0.208 ***	0.149 ***	0.152 ***
Flexibility	0.818 ***	0.878 ***	0.625 ***	0.580 ***	0.210 ***	0.230 ***
Model B – Intrinsic Utility + Consumption Flexibility with Restricted Coefficients						
Hit Rate ‡	48.4%	47.0%	31.8%	34.7%	23.2%	26.6%
Log Likelihood	-436.1	-1018.1	-622.5	-1499.2	-770.8	-1901.5
AIC	874.3	2038.2	1247.0	3000.4	1543.6	3804.9
BIC	878.3	2043.0	1251.0	3005.2	1547.6	3809.9
Coefficient:						
Intrinsic Utility+Flexibility	0.156 ***	0.142 ***	0.140 ***	0.141 ***	0.136 ***	0.137 ***
Model C – Intrinsic Utility Only						
Hit Rate ‡	36.6%	33.0%	21.1%	20.9%	20.8%	17.4%
Log Likelihood	-447.5	-1039.1	-634.4	-1528.7	-778.5	-1923.6
AIC	897.0	2080.2	1270.9	3059.5	1559.1	3849.2
BIC	901.0	2085.0	1274.9	3064.3	1563.1	3854.1
Coefficient:						
Intrinsic Utility	0.059 *	0.012	0.055 **	0.058 ***	0.086 ***	0.084 ***
Model D – Choice Set Intercepts						
Hit Rate ‡	47.9%	51.8%	39.3%	38.1%	15.7%	24.3%
Log Likelihood	-429.5	-958.0	-596.6	-1446.3	-779.5	-1890.7
AIC	863.1	1920.0	1201.2	2900.6	1571.0	3793.4
BIC	871.1	1929.7	1217.2	2920.1	1595.2	3823.1
Coefficients:						
Intercept (2,1,0)	0.228	0.374 ***				
Intercept (1,1,1)	0.745 ***	0.999 ***				
Intercept (3,1,0,0)			-0.211	-0.247 *		
Intercept (2,2,0,0)			-0.118	0.079		
Intercept (2,1,1,0)			0.919 ***	0.930 ***		
Intercept (1,1,1,1)			0.091	0.146		
Intercept (4,1,0,0,0)					-0.429 *	-0.434 **
Intercept (3,2,0,0,0)					-0.219	-0.007
Intercept (3,1,1,0,0)					0.229	0.570 ***
Intercept (2,2,1,0,0)					0.229	0.372 ***
Intercept (2,1,1,1,0)					-0.383 *	-0.423 **
Intercept (1,1,1,1,1)					-0.579 **	-0.490 ***

‡ Using 10-fold validation
* Significant at $\alpha = 0.05$
** Significant at $\alpha = 0.01$
*** Significant at $\alpha = 0.001$

Table 4: Predicted vs. Actual* Choice Set Variety

	3-Product Set		4-Product Set		5-Product Set	
	Candy	Salty Snack	Candy	Salty Snack	Candy	Salty Snack
1	13.6% (22.8%)	5.6% (19.3%)	9.0% (15.9%)	9.1% (15.3%)	15.3% (16.1%)	13.7% (14.0%)
2	18.2% (28.6%)	17.8% (28.1%)	19.2% (27.0%)	20.2% (28.4%)	22.8% (23.4%)	20.5% (23.0%)
3	68.3% (47.9%)	76.5% (52.5%)	49.1% (39.8%)	44.0% (38.7%)	48.7% (40.5%)	51.4% (45.2%)
4			22.7% (17.4%)	26.7% (17.7%)	3.4% (11.0%)	4.2% (9.2%)
5					9.9% (9.0%)	10.3% (8.6%)

* Actual variety percentages are reported in parentheses; corresponding predicted percentages are not

set variety. Specifically, we find that, although maximum variety increases with the choice set size, mean predicted variety increases much less, and median and modal predicted variety of three product alternatives does not increase at all.

5 Longitudinal Experiments

Proposition 2 specifies a rational consumption policy for multi-product choice sets. This consumption policy—that available product alternatives will be chosen for consumption in proportion to their current inventory—is implied by *RMPC*. As products are consumed one-by-one, the effect of this policy is to maintain flexibility by probabilistically balancing the inventory of product alternatives in the set. To test Proposition 2, we conducted two longitudinal experiments [i] to determine whether actual consumption patterns are consistent with Proposition 2, [ii] to test the robustness of Proposition 2 to different choice set concentrations, and also [iii] to test variety seeking as an alternative explanation for actual consumption choices.

5.1 Experimental Design

The first experiment was conducted to determine whether participants' consumption choices are consistent with Proposition 2. The second experiment was conducted to test the robustness of Proposition 2 across a variety of inventory levels and to test for variety seeking in consumption choices. Following Simonson (1990), these experiments involved students consuming snack products once or twice per week over a series of consumption occasions. The experimental design was approved for human participants by our University's Institutional Review Board.

The first phase of these longitudinal experiments was exactly the same preference elicitation process used in the large-scale multi-product choice experiment and detailed in §4.1. Specifically, participants [i] provided demographic information, [ii] selected a preferred snack category, [iii] assigned long-run choice percentages to all products in that category's assortment and [iv] identified their three (ordered) favorite products from that assortment, then [v] provided information about their usage rate, attitudes, and perceptions about products in the category. The second (i.e., consumption) phase of the longitudinal experiments required participants to sequentially consume a set of five snacks. Each student participant was assigned a box with five snacks to consume one-at-a time at consecutive class meetings over multiple weeks. Each box contained a combination of the participant's favorite and second-favorite product alternatives.[14] The snack boxes were wheeled

[14]We selected the participant's two most preferred snacks to ensure, to the extent possible, that they had a choice between products that they like. In a pilot study, some participants had received sets with a combination of their favorite and second-favorite product alternatives; others had received a combination of their favorite and third-favorite alternatives. That design manipulated the relative preference between product alternatives. We found in the pilot study that relative preference for alternatives in the set did not have a significant effect on consumption choices, so

into the classroom in a cart before each class meeting. At the beginning of class, participants chose one snack from their box for personal consumption that day. They were instructed not to trade snacks or to select a snack for someone else to consume. At the end of class, researchers removed the cart and recorded which product alternative each participant had chosen to consume. Each consumption choice reduced the participant's inventory of either their favorite or second-favorite alternative. This procedure was repeated until all snacks were consumed.

5.2 Longitudinal Experiment 1

Participants' set of five snacks included either: [i] four units of their favorite product alternative and one unit of their second-favorite, or [ii] one unit of their favorite and four units of their second favorite. Sixty-nine graduate students completed the first phase questionnaire; 67 undertook the second phase consumption task.

Figure 1: Experiment 1 - Percent Favorite Chosen vs. Favorite Inventory

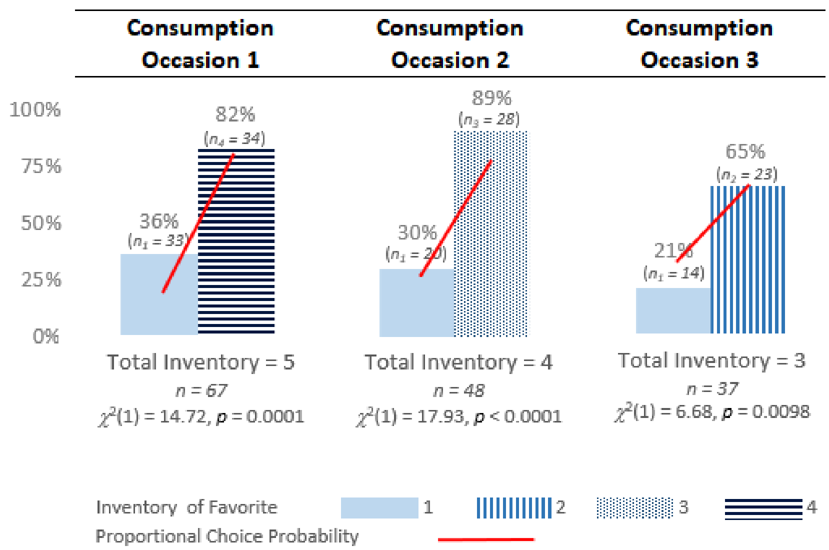

Figure 1 shows, from left to right, the sample proportion of participants who chose their favorite on the first, second, and third consumption occasions. We did not assess the last two consumption occasions because they offered no information about the relationship between inventory and consumption. The red line shows the theoretical (i.e., proportional) probability of choosing the favorite, based on the inventory of the two product alternatives available. Note that, if consumers are myopic, we would not expect the bar heights to be significantly different because consumption of the favorite would be driven by preference, not inventory. A visual inspection clearly shows that the data are not consistent with myopic behavior. On the first consumption occasion, all 67 participants chose from a set that included either one or four units of their favorite product alternative with the remainder being their second-favorite. For this consumption occasion, we find strong evidence of a relationship between inventory and consumption choice of the favorite $\left(\chi^2\,(1) = 14.72, p = 0.0001\right)$. On the second consumption occasion, the 48 participants who still had a choice (and completed the task) chose from a set that included either one or three units of their favorite with the remainder being their second-favorite. For this consumption occasion, we find again strong evidence of a

this manipulation was not included in the two longitudinal experiments.

relationship between inventory and consumption choice $\left(\chi^2(1) = 17.93, p < 0.0001\right)$. On the third consumption occasion, the 37 participants who still had a choice (and completed the task) chose from a set with either one or two units of their favorite alternative with the remainder being their second-favorite. For this consumption occasion, we find again strong evidence of a relationship between inventory and consumption choice $\left(\chi^2(1) = 6.68, p = 0.0098\right)$. Across the three consumption occasions, those who had only a single unit of their favorite in inventory were significantly less likely to consume it than those who had multiple units of their favorite in inventory. Further, participants were more than twice as likely to choose their favorite if they had multiple units in inventory (vs. a single unit) on all three consumption occasions. Overall then, Experiment 1's consumption choices provide strong support for Proposition 2's inventory-based (i.e., forward looking) consumption policy.

5.3 Longitudinal Experiment 2

Recall that some previous studies characterized diversification in multi-product choice as variety-seeking (Simonson 1990, Read and Loewenstein 1995). Yet satiation, the rationale for variety-seeking, applies to *consumption choices* as opposed to the multi-product choices that precede consumption (Raju 1980, McAlister and Pessimier 1982). In order for variety seeking to affect multi-product choices, consumers would have to anticipate variety seeking in their downstream consumption choices. In Experiment 2, we evaluate variety-seeking as an alternative explanation for diversification of multi-product choices by testing whether consumption choices show evidence of variety-seeking. Specifically, we test for a relationship between consecutive consumption choices. The design of the second phase consumption task in Experiment 1 did not permit such a test. In that experiment, every consumption choice was made between a favorite product alternative with multiple units in inventory and another with a single unit in inventory. If the product alternative with a single unit in inventory was chosen, it could not be chosen again. As a result, every consumption choice on the second and third consumption occasions was necessarily preceded by consumption of the alternative with multiple units in inventory. This confound precluded a test of consecutive consumption choices.

In Experiment 2, participants again received a set of five snacks split between participants' favorite and second-favorite product alternatives. In this experiment however, participants received either one, two, three, or four units of their favorite—not one or four units as in the previous experiment—with the remainder being their second-favorite. This design allowed us to test variety-seeking in consumption choices while also evaluating the applicability of Proposition 2's inventory-based consumption policy to a wider range of inventory allocations. Note that Experiment 2 was otherwise unchanged from Experiment 1.

A total of 201 participants completed the first phase questionnaire; 192 provided data for the second phase. Figure 2 shows consumption choices of the favorite on Experiment 2's first, second, and third consumption occasions. On the first consumption occasion, 192 participants chose from a five-product set that included one, two, three, or four units of their favorite product alternative with the remainder being their second-favorite. For this consumption occasion, we find strong evidence of a relationship between inventory and consumption choice of the favorite $\left(\chi^2(3) = 35.68, p < 0.0001\right)$. On the second consumption occasion, the 166 participants who still had a choice (and completed the task) chose from a set with one, two, or three units of their favorite alternative with the remainder being their second-favorite. For this consumption occasion, we find again strong evidence of a relationship between inventory and consumption choice $\left(\chi^2(2) = 23.96, p < 0.0001\right)$. On the third consumption occasion, the remaining 128 participants who still had a choice (and completed the task) chose from a set with either one or two units

Figure 2: Experiment 2 - Percent Favorite Chosen vs. Favorite Inventory

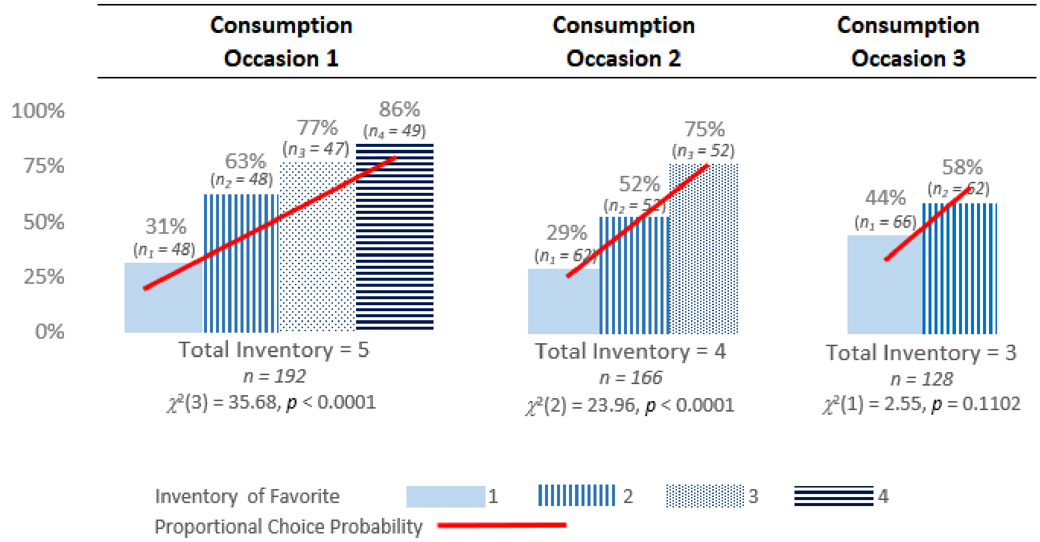

of their favorite with the remainder being their second-favorite. For this consumption occasion, we do not find sufficient evidence to confirm a relationship between inventory and consumption choice $\left(\chi^2(1) = 2.55, p = 0.1102\right)$, although the pattern of choices is consistent with Proposition 2's inventory-based consumption policy. It is important to note that participants were more likely to choose their favorite when they had more inventory of that favorite for every inventory configuration on every consumption occasion—additional strong support for Proposition 2 and *RMPC* theory's inventory-based consumption policy.

The inventory-based policy that we observe in the data may be intuitive, but is by no means the only potentially intuitive consumption policy. As noted earlier, myopic consumers might consume their favorite in proportion to its long-run choice probability. Alternatively, indifferent consumers might consume available product alternatives with equal probability. We observed neither of these patterns; rather we observed a consumption policy consistent with forward-looking consumers maximizing the future value of the set.

In Experiment 2, we extended the manipulation of beginning inventory levels to test for variety seeking in consumption choices. Figure 4 shows consumption choices for the second, third, and fourth consumption occasions, together with lagged choices (i.e., the previous consumption occasion's choice). Satiation, and so variety seeking, would imply a negative relationship between successive consumption choices as the consumer's utility for the product and its attributes diminishes. We find no such relationship. Specifically, tests of independence for the second $\left(\chi^2(1) = 0.85, p = 0.3558\right)$, third $\left(\chi^2(1) = 0.12, p = 0.7257\right)$, and fourth $\left(\chi^2(1) = 0.23, p = 0.6324\right)$ consumption occasions provide scant evidence of satiation in participants' consecutive consumption choices. The lack of variety seeking evidence stands in contrast to the compelling evidence of inventory-based consumption choices.

Following Experiment 2's consumption phase, we followed up with participants to determine which factors had affected their consumption choices. Specifically, participants were asked: "Looking back on only those days when you could choose between your favorite and second-favorite snacks, which of the following factors affected your choices?" Stochastic *preference* was identified as "which snack I felt like eating the most on that day." *Inventory* was identified as "the number of each product that was available in my box on that day." *Satiation* was identified as "which snack I had eaten recently and so was 'getting tired of.'" Participants evaluated all three factors

Figure 3: Experiment 2 - Percent Favorite Chosen vs. Lagged Favorite/Non-Favorite

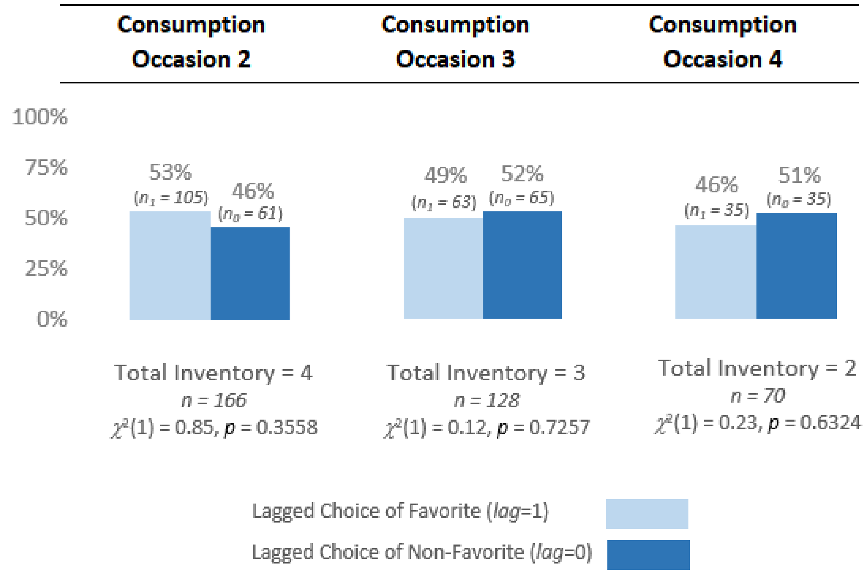

using a 7-point agreement scale with 1 = "Strongly Disagree" and 7 = "Strongly Agree." For the 177 participants who responded, the mean response for current *preference* was 5.71 (indicating agreement), the mean response for *inventory* was 5.14 (indicating somewhat lesser agreement), and the mean response for *satiation* was 4.23 (indicating neither agreement nor disagreement). Paired t-tests clearly show that both current *preference* ($t\,(175) = 7.75, p < 0.0001$) and *inventory* ($t\,(175) = 5.58, p < 0.0001$)were more important than *satiation* when making consumption choices. It is worth noting that *RMPC* theory holds that both current preference and inventory will be considered when making consumption choices. Satiation, and hence variety seeking, is an entirely unrelated rationale that is not supported by the data.

In summary, the two longitudinal experiments provide clear and consistent support for Proposition 2 and *RMPC* theory's inventory-based optimal consumption policy. The second longitudinal experiment does not support variety seeking as an alternative explanation for consumption choices.

6 Panel Data Evidence

Proposition 3 states that the variety included in a consumer's optimal multi-product choice set will decrease with consumption rate. Rather than testing this proposition by attempting to manipulate consumption rate, we chose instead to use observational data. Following Simonson and Winer (1992), we use panel data purchases of single-serve yogurt products. We use syndicated panel data from two major metropolitan markets (one in the midwestern US totaling 1707 households, the other in the southeastern US totaling 1031 households) with consumer packaged goods purchases made during a 24-month period between September 2004 and September 2006. We limit the data to single-member households in order to avoid confounding multi-product choice, as it applies to an individual consumer, with intra-household preference heterogeneity. This screen yielded 445 = 308 + 137 single-person households from the two markets. The dataset is partitioned so that the first 18 months serves as a calibration period while the final six months is used for estimation. We required that households had made at least 10 yogurt purchases during the calibration period, then purchased again during the estimation period. The final dataset includes 70 single-member households that purchased a total of 8,670 single-serve yogurt cups on 1,611 shopping trips during the calibration period, then purchased 2,376 cups on 443 shopping trips during the estimation period. Although this dataset is small, it is constructed purposefully to avoid intra-household heterogeneity and so provide a clean test of Proposition 3.[15]

Data from the calibration period were used for two purposes. The first was to determine panelists' long-run consumption preferences. Those preferences were determined by *UPC* because the variety and ambiguity of flavors (e.g., white chocolate strawberry, cherry vanilla creme, pina colada, cookies & creme, apricot mango, lemon meringue, key lime pie, mixed berry) did not allow for a parsimonious attribute decomposition. From consumption preferences, we developed household-level utilities for *UPC*s. The second use of calibration data was to calculate consumption rates. Panelists did not record their consumption—such data are rare—so we we estimated consumption rates, assuming that all yogurt purchases made during the calibration period were consumed. Initially, we conjectured that each day presented a consumption opportunity. Interestingly, we found that one panelist consumed 1.328 units/day (recall that each unit is a single serving), buying yogurt on 114 shopping trips during the calibration period. All other panelists consumed an average of less than one unit/day.

Data from the calibration period were used to assess the relative variety of yogurt purchases. Consistent with *RMPC* theory of shopping and consumption, we assumed that n, the number of units chosen on a given trip, is exogenous.[16] Variety was measured as the number of product alternatives m in the chosen set. Clearly, m depends on the set size n ($m \leq n$). To control for this dependency, we took advantage of the fact that the base *RMPC* model (with no outside option) is actually the limiting case of the more general model (with utility of the outside option set to -∞) for a given set size n. To evaluate the observed variety, m, we therefore compared it to the variety of the base *RMPC* model's optimal set of the same size, m^{opt}, which implicitly assumes the maximum consumption rate. Using (3.1) to compute set valuations, we determined m^{opt} from the base *RMPC* model for every panelist and every set size n, which we then used to determine the relative variety of observed purchases. The relative variety measure for yogurt purchases is the proportional difference between observed and optimal variety, $D = \frac{m-m^{opt}}{m^{opt}}$.[17]

Proposition 3 states that the variety included in a consumer's optimal choice set is decreasing in consumption rate; however, it does not specify a functional form for that relationship. We therefore estimated nonparametric correlations—Spearman's Rank Correlation and Kendall's Tau—as well as Pearson's R. An important characteristic of the data is that relative variety changed across a panelist's yogurt purchases, but the panelist's consumption rate did not. We therefore include subscripts for trip t and household h in the remaining exposition. Household consumption rate was measured in $Units/Day_h$, but its inverse $Days/Unit_h$ was also analyzed. The proportional difference in variety for household h on trip t, D_{ht}, was used to compute trip-level correlations; the household's average across trips, $\overline{D}_h$, was used to compute correlations at the household-level.

Table 5 shows that all correlations have the expected sign—negative for $Units/Day_h$ and positive for $Days/Unit_h$. All nonparametric correlations are significant at the 0.05 level, and are uniformly higher in magnitude for household-level correlations than for trip-level correlations. Interestingly, Pearson's R was higher in magnitude for $Days/Unit_h$ than for $Units/Day_h$, suggesting that the relationship between relative variety and consumption rate was more linear for the former

[15]Two panelists that met the screening criteria were omitted from the dataset because they consistently made exceptionally large purchases—up to 146 cups on a single trip. Given the perishable nature of yogurt, such purchases were clearly not intended for personal consumption.

[16]We would expect higher consumption rates to be associated with larger n; however, the relative variety measure we used prevents this association from affecting our analysis.

[17]Using $\frac{m-m^{opt}}{m^{opt}}$ rather than, say, $m - m^{opt}$ enabled us to avoid scaling issues.

Table 5: Correlations of Relative Variety and Consumption Rate

	Trip-Level Correlations (D_{ht})		Household-Level Correlations ($\overline{D}_h$)	
	Units/Day$_h$	Days/Unit$_h$	Units/Day$_h$	Days/Unit$_h$
Spearman's Rank	-0.1707	0.1707	-0.3303	0.3303
	(0.0003)	(0.0003)	(0.0052)	(0.0052)
Kendall's Tau	-0.1249	0.1249	-0.2107	0.2107
	(0.0003)	(0.0003)	(0.0116)	(0.0116)
Pearson's R	-0.1282	0.2239	-0.1625	0.2632
	(0.0069)	(<.0001)	(0.1788)	(0.0277)

p-values are shown in parentheses

than the later. To further evaluate the functional form of the relationship between relative variety and consumption rate, we conducted a hierarchical Bayesian analysis (see Appendix B). Beyond insights into functional form, that analysis confirmed the relationship between consumption rate and relative variety that is evident in the correlations reported in Table 5. Taken together, the correlations and hierarchical Bayesian analysis provide clear support for Proposition 3.

7 Concluding Remarks

This paper presents empirical evidence of rationality in multi-product choices, rather than a bias to include too much variety. The first proposition we tested deals with the inherent tradeoff between the intrinsic utility of products in the set and the consumption flexbility afforded by the structure of the set. In a large-scale experiment, we found that this tradeoff explains consumers' multi-product choices better than other models. In particular, we found that including consumption flexibility in a model with intrinsic utility results in vastly higher validation hit rates than a model with intrinsic utility alone. This tradeoff model could be used for a number of practical applications, including the design and pricing of variety packs for consumer markets and the pricing of customized variety packs for individual consumers. Demonstrating the explanatory power of this tradeoff, in particular the incremental explanation provided by consumption flexibility, is a primary contribution of this paper. Estimating the rational benefit of variety and consumption flexibility in multi-product choice is another important contribution.

The second proposition we tested considers how multi-product sets will be consumed. The consumption policy implies that consumers will balance their inventory of product alternatives probabistically over consumption occasions. We found strong evidence that consumption choices are consistent with this policy. In two longitudinal experiments, we observed a pattern consistent with inventory-based consumption; more generally, consistent with consumers probabilitistically matching available product alternatives to consumption occasions in a way that extracts the maximum utility from their multi-product set. Support for this second proposition also offers additional indirect support for the first proposition inasmuch as the consumption policy underpins the optimality of the tradeoff between intrinsic utility and consumption flexibility.

The third proposition that we tested predicts that the multi-product choices of consumers with lower consumption rates will include more variety than the multi-product choices of consumers with higher consumption rates. This prediction can be explained in terms of competition. In cases where a category is seldom consumed (implying an attractive outside option), more variety in multi-product choice is the best strategy to overcome that outside option and 'win' a consumption occasion. For example, if one rarely drinks wine, having two bottles of red or two bottles of white is likely inferior to having one bottle of each. The set with greater variety offers a higher probability of including an alternative that one wants when the mood strikes. We found empirical support for this proposition using yogurt purchases in multi-market panel data.

In summary, we have presented evidence that uniformly supports a rational theory of multi-product choice based on the maximization of expected future utility. This theory represents a compelling alternative to existing theories of multi-product choice that explain observed diversification in terms of bias (i.e., diversification bias) or the expectation of satiation (i.e., variety-seeking).

This paper has the potential to generate a number of future research opportunities. Testing this normative theory of multi-product choice, along with variety-seeking and diversification bias, is likely to find conditions under which different models of behavior apply. For example, consumption rate might moderate the effect of variety seeking in multi-product choice. Based on once- or twice-weekly consumption of snack products (cf. Simonson 1990), we found evidence of forward-looking consumption choices but not variety seeking. This might change for more frequent consumption. Exploring boundary conditions of different multi-product choice models could also be a fruitful avenue for future research. Consumption rate is logically related to constructs like involvement, familiarity, and expertise in a product category, yet might affect multi-product choices in different ways than these constructs suggest. Finally, studies of consumption have generally been limited—perhaps because consumption data are difficult to obtain. Gathering and analyzing observational data of consumption in the context of multi-product choice could be another potentially fruitful avenue for future research.

References

[1] A. Alptekinoglu and J. H. Semple. Consuming a set of substitutable products using a general random utility framework. 2018.

[2] M. Ben-Akiva and S. R. Lerman. *Discrete Choice Analysis*. The MIT Press, Cambridge, MA, 1985.

[3] J. P. Dube. Multiple discreteness and product differentiation: Demand for carbonated soft drinks. *Marketing Science*, 23(1):66–81, 2004.

[4] Dynata. https://www.dynata.com, 2022.

[5] E. J. Fox, L. E. Norman, and J. H. Semple. Choosing an n-pack of substitutable products. *Management Science*, 64(5):2366–2379, 2018.

[6] L. Guo. Capturing consumption flexibility in assortment choice from scanner panel data. *Management Science*, 56(10):226–40, 2010.

[7] B. E. Kahn and R. K. Ratner. *Variety for the Sake of Variety? Diversification Motives in Consumer Choice*, pages 102–21. Routledge, London, 2005.

[8] D. Kahneman and J. Snell. *Predicting Utility*, pages 295–310. University of Chicago Press, Chicago, 1990.

[9] J. Kim, G. M. Allenby, and P. E. Rossi. Modeling consumer demand for variety. *Marketing Science*, 21(3):229–50, 2002.

[10] J. Kim, G. M. Allenby, and P. E. Rossi. Product attributes and models of multiple discreteness. *Journal of Econometrics*, 38(1):208–230, 2007.

[11] D. T. Kollat and R. P. Willett. Consumer impulse purchasing behavior. *Journal of Marketing Research*, 4(1):21–31, 1967.

[12] S. Lee and G. M. Allenby. Modeling indivisible demand. *Marketing Science*, 33(3):364–81, 2014.

[13] L. McAlister and E. Pessemier. Variety seeking behavior: An interdisciplinary review. *Journal of Consumer Research*, 9(3):311–22, 1982.

[14] D. McFadden. *Modelling the Choice of Residential Location*, volume 1, pages 75–96. North Holland, 1978.

[15] P. S. Raju. Optimum stimulation level: Its relationship to personality, demographics, and exploratory behavior. *Journal of Consumer Research*, 7(3):272–82, 1982.

[16] D. Read and G. Loewenstein. Diversification bias: Explaining the discrepancy in variety seeking between combined and separated choices. *Journal of Experimental Psychology: Applied*, 1(1):34–49, 1995.

[17] T. J. Richards, M. I. Gomez, and G. Pofahl. A multiple-discrete/continuous model of price promotion. *Journal of Retailing*, 88(2):206–25, 2012.

[18] L. C. Salisbury and F. M. Feinberg. Future preference uncertainty and diversification: The role of temporal stochastic inflation. *Journal of Marketing Research*, 35(2):349–59, 2008.

[19] I. Simonson. The effect of purchase quantity and timing on variety-seeking. *Journal of Marketing Research*, 25(2):150–62, 1990.

[20] I. Simonson and R. S. Winer. The influence of purchase quantity and display format on consumer preference for variety. *Journal of Consumer Research*, 19(1):133–8, 1992.

[21] D. J. Spiegelhalter, N. G. Best, B. P. Carlin, and A. van der Linde. Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society Series B (Methodological)*, 64(4):583–639, 2002.

[22] John W. Walsh. Flexibility in consumer purchasing for uncertain future tastes. *Marketing Science*, 14(2):148–65, 1995.

Appendix A: Short Proof of the Value Function and Optimal Policy

Proof. The proof is by induction. The standard zero-mean Gumbel ϵ_i has c.d.f. $F(x) = exp\left(-exp\left(-x-\gamma\right)\right)$ where γ is the Euler-Mascheroni constant. The value formula (3.1) is trivially true for $n = 1$. Assume the truth of (3.1) for the case (n-1). Let $k^T = (k_1, \ldots, k_M)$ with $\sum_{i=1}^{M} k_i = n-1$. Then the truth of the result for n-1 implies $V(k) = ln\left(\frac{(n-1)!}{\prod_{i=1}^{M} k_i!} \cdot exp\left(k^T u\right)\right)$. We now use the expected value formula for $\max_{i=1,2,\ldots,m}(W_i + \epsilon_i)$ where W_i are given parameters. If there are m distinct alternatives to choose from each having expected utility W_i, then the expected value of the best (maximum) choice is

$$E\left(\underset{i=1,\ldots,m}{Max}\{W_i + \epsilon_i\}\right) = ln\left[\sum_{i=1}^{m} e^{W_i}\right]. \tag{.1}$$

For the case of n items, assume *wlog* that there are m distinct alternatives and that they are the first m alternatives of the M possible alternatives. Let $k^T = (k_1, \ldots, k_m)$ with $\sum_{i=1}^{m} k_i = n$. The expected utility of alternative i is u_i, and $u^T = (u_1, \ldots, u_m)$. The expected value of the set is then

$$\begin{aligned}
V(k) &= E\left(\underset{i}{Max}\{V(k-e_i) + u_i + \epsilon_i\}\right) \\
&= ln\left(\sum_{i=1}^{m} exp\left(V(k-e_i) + u_i\right)\right) \qquad (by\ (.1)\ with\ W_i = V(k-e_i) + u_i) \\
&= ln\left(\sum_{i=1}^{m} \frac{(n-1)!(k_i)}{\prod_{j=1}^{m} k_j!} \cdot exp\left((k-e_i)^T u + u_i\right)\right) \qquad (by\ the\ induction\ hypothesis) \\
&= ln\left(\frac{n!}{\prod_{j=1}^{M} k_j!} \cdot exp\left(k^T u\right)\right).
\end{aligned}$$

(Note that we use the fact that $0! = 1$ to continue the formula to the remaining $M-m$ alternatives in the last step; a similar continuation applies to $k^T u$.) Moreover, the optimal policy, to select the alternative maximizing $ln\,(k_i) + \epsilon_i$, is readily apparent:

$$\begin{aligned}
\underset{i}{Max}\{V(k-e_i) + u_i + \epsilon_i\} &= \underset{i}{Max}\left\{ln\left(\frac{(n-1)!(k_i)}{\prod_{j=1}^{m} k_j!} \cdot exp\left((k-e_i)^T u\right)\right) + u_i + \epsilon_i\right\} \\
&= \underset{i}{Max}\left\{ln\left(\frac{(n-1)!(k_i)}{\prod_{j=1}^{m} k_j!} \cdot exp\left(k^T u\right)\right) + \epsilon_i\right\} \\
&= ln\left(\frac{(n-1)!}{\prod_{j=1}^{m} k_j!} \cdot exp\left(k^T u\right)\right) + \underset{i}{Max}\{ln\,(k_i) + \epsilon_i\}
\end{aligned}$$

□

Appendix B: Hierarchical Bayesian Analysis of Relative Variety

The combination of a time invariant household-level predictor ($Units/Day_h$ or $Days/Unit_h$) and a time varying household-level response variable (D_{ht}) led us to model the data using a hierarchical random coefficients model. The first-level equation was specified

$$D_{ht} = \theta_h + \xi_{ht}$$

and the hierarchical equation was specified

$$\theta_h = \delta + \gamma \cdot f\,(Units/Day_h) + \zeta_h,$$

where $f(\cdot)$ is a monotonic transformation to allow for flexibility in functional form. The parameter of interest is γ, which captures the relationship between the relative variety of a choice set, D_{ht}, and the transformed consumption rate, $f\,(Units/Day_h)$. The resulting model was estimated in a hierarchical Bayesian framework with minimally informative priors so that the posterior estimates were driven by the data. For each model, the 25000 Markov Chain Monte Carlo iterations converged quickly after a short burn-in period and autocorrelation proved acceptable (we also "thinned the line"), resulting in a stable posterior distribution for γ.

Table 6: Relative Variety Model Estimates

Model	Predictor	Posterior Predictions			Posterior Distribution of γ		
		DIC	MAD	MSE	Mean	$Pr(\gamma < 0)$	(2.5% , 97.5%)
Fixed Intercept	N/A	364.26	0.267	0.132	N/A	N/A	N/A
Random Coefficient (γ=0)	N/A	245.33	0.267	0.133	N/A	N/A	N/A
Random Coefficient	$Units/Day_h$	245.24	0.267	0.134	-0.2107	0.909	(-.5232,.0970)
Random Coefficient	$exp(Units/Day_h)$	244.80	0.270	0.136	-0.1000	0.870	(-.2747,.0857)
Random Coefficient	$Days/Unit_h$	245.55	0.258	0.127	0.0145	0.016	(.0016,.0274)
Random Coefficient	$ln(Days/Unit_h)$	244.73	0.261	0.130	0.0923	0.024	(.0006,.1838)

Table 6 shows estimates for the "Random Coefficients" model specified by the two equations above, with selected monotonic transformations of the independent variable. The table also shows estimates for two nested models models: (i) "Fixed Intercept" (i.e., $\theta_1 = \theta_2 = \ldots = \theta_H = \theta$), and (ii) "Random Coefficients ($\gamma = 0$)." To compare models, we used the Deviance Information Criterion [DIC], a Bayesian analog of AIC (see Spiegelhalter, et al. 2002). To assess model predictions, we bootstrapped the estimation dataset, then compared actual values of D_{ht} with predicted values using Mean Absolute Deviation [MAD] and Mean Squared Error [MSE]. Four random coefficients models are reported in the table, reflecting different transformations of $Units/Day_h$. We selected $Units/Day_h$ (no transformation) and $Days/Unit_h$ (inverse transformation), along with two monotonic transformations that generally fit the data better: $exp\,(Units/Day_h)$ and $ln\,(Days/Unit_h)$.

The two nested models, "Fixed Intercept" and "Random Coefficient ($\gamma = 0$)," differ greatly from one another in terms of fit—"Random Coefficient ($\gamma = 0$)" has a far lower DIC than "Fixed Intercept" (lower is better)—but are very similar in predictive accuracy. In terms of fit, the "Random Coefficient ($\gamma = 0$)" model has only a slightly higher DIC than three of the four full "Random

Coefficient" models, and actually has a slightly lower DIC than the fourth. In terms of predictive accuracy, the "Random Coefficient ($\gamma = 0$)" model offers posterior predictions similar to the four full "Random Coefficient" models. We therefore conclude that unmodeled individual differences explain much more variation in D_{ht} than consumption rate does. On the other hand, the functional form of the relationship between consumption rate and variety matters for predictive accuracy. The "Random Coefficient" model using $ln\left(Days/Unit_h\right)$ as the predictor has the lowest DIC and, like the "Random Coefficient" model using $Days/Unit_h$ as the predictor, offers more accurate predictions than the other models. The superior predictive accuracy of the two models using $f\left(Days/Unit_h\right)$ as the predictor is consistent with the nonparametric correlations reported above, where the proportional difference in variety was more highly correlated with $Days/Unit$ than with $Units/Day$. Taken together, these results suggest that the relationship between relative variety and usage rate be specified as a function of $Days/Unit$. A more extensive exploration of functional form is left for future research.

Returning to the preferred model, the "Random Coefficient" model using $ln\left(Days/Unit_h\right)$ as the predictor, the parameter estimate for γ is positive and significant. The mean estimate is 0.092 and, based on the posterior *cdf,* $Pr\left(\gamma > 0\right) = 0.024$. For the "Random Coefficient" models using different predictors, the posterior estimate of γ is always in the expected direction—positive for $f\left(Days/Unit_h\right)$ and negative for $f\left(Units/Day_h\right)$, though the posterior estimates of γ are only significant for models using $f\left(Days/Unit_h\right)$ as the predictor.