

Choosing an n -Pack of Substitutable Products

Edward Fox,^a Laura Norman,^a John Semple^a

^a Cox School of Business, Southern Methodist University, Dallas, Texas 75275

Contact: efox@smu.edu (EF); lenorman@cox.smu.edu (LN); jsemples@smu.edu (JS)

Received: August 28, 2015

Revised: May 13, 2016; August 25, 2016

Accepted: October 17, 2016

Published Online in Articles in Advance:
May 19, 2017

<https://doi.org/10.1287/mnsc.2017.2729>

Copyright: © 2017 INFORMS

Abstract. We develop a framework to model the shopping and consumption decisions of forward-looking consumers. Assuming that the consumer's future utility for each product alternative can be characterized by a standard random utility model, we use dynamic programming to determine the optimal consumption policy and the maximum expected value of consuming any n substitutable products selected while shopping (an n -pack). We propose two models. In the first (*canonical*) model, we assume that an alternative is consumed on each successive consumption occasion and obtain a closed-form optimal policy and a closed-form value function. Given a consumer's preferences for the product alternatives in an assortment, we then show how to identify that consumer's optimal n -pack using a simple swapping algorithm that converges in at most n swaps. In the second (*generalized*) model, we introduce an outside option so that a product alternative need not be consumed on each consumption occasion. We obtain a closed-form value function for the generalized model and show that its optimal n -pack is related to that of the canonical model using a special type of majorization. Additional structural properties and implications of each model are explored, as are other applications.

History: Accepted by Vishal Gaur, operations management.

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/mnsc.2017.2729>.

Keywords: dynamic programming • decision analysis • utility/preference applications • multinomial logit • assortment optimization

1. Introduction

A great deal of research has focused on two key consumer decisions: (i) the shopping, or purchase, decision, which is usually made in store; and (ii) the consumption decision, which is made later at the time of consumption. For consumer packaged goods, the two decisions are inextricably linked, even though they occur at different times and in different places.

Shopping decisions are made for future consumption, which generally occurs over multiple consumption occasions. On each such occasion, only products selected previously while shopping are available to be consumed. The shopping decision therefore creates a set of products for future consumption and so is inherently forward looking. At the same time, the shopping decision *constrains* the product alternatives that are available for future consumption compared to the full assortment in store. Each successive consumption decision may further constrain the product alternatives available, depending on how many units of each product alternative were selected when shopping. If the consumer has only a single unit of a particular product alternative remaining, then consuming it would preclude choosing that alternative on all succeeding consumption occasions. Because consumption decisions—like the shopping decision—affect the expected utility of future consumption, these decisions are also inherently forward looking.

In this paper, we propose a two-stage model of shopping and consumption. In the first stage, the consumer selects a set of n products (n is exogenous), the " n -pack," to be consumed over a horizon having multiple future periods. In the second stage, the n -pack is iteratively consumed over multiple periods, which we model using dynamic programming (DP). The n -pack may include $m \leq n$ product alternatives so multiple units of each alternative may be selected. If one unit is consumed each period and there is no outside option, the horizon has n periods. If an outside option is introduced, the horizon can have any number of periods. The primary purpose of this paper is to investigate and establish the basic analytical properties of the model with and without an outside option.

The model without an outside option is referred to as the *canonical model* because its assumptions mirror those used in the consumer behavior literature for over 25 years, beginning with Simonson (1990). While this model introduces analytical complexities absent from prior work, it yields a closed-form optimal consumption policy and a closed-form value function for any n -pack. The latter function parsimoniously captures the benefits of diversifying the n -pack versus choosing more units of one's most preferred product alternatives (as measured by the alternatives' expected utilities). Analysis of the canonical model yields some interesting and important insights. We find that the

optimal consumption policy does not require the consumer to select the product that offers the highest utility on each consumption occasion; rather, the optimal policy is about matching products with consumption occasions. This matching policy depends on the inventory of available product alternatives and the stochastic (but *not* the deterministic) component of utility for available product alternatives. Consumption decisions therefore do not reveal preferences, per se. Analysis of our model shows that a consumer's optimal n -pack is also the set of products most likely to be consumed over n independent consumption occasions had each consumption choice been made from the full assortment, what Simonson (1990) and Read and Loewenstein (1995) refer to as "sequential choice." Note that sequential choice does not restrict the alternatives that can be selected on each consumption occasion, whereas choosing from a previously selected n -pack most certainly does.

The model with an outside option on each consumption occasion is referred to as the *generalized model*. To our knowledge, no other study of shopping and consumption has incorporated an outside option, which effectively slows the rate of consumption, ceteris paribus. For the generalized model, the value function and the optimal policy are also closed-form but more complex than in the canonical case. We show that the optimal n -pack in the generalized case is at least as diversified as the optimal n -pack in the canonical case. This is made possible by relating the optimal n -pack of the generalized model to the optimal n -pack of the canonical model using a new, specialized type of majorization. Finally, we show that the marginal change in the value of an n -pack decreases as the time horizon increases; *i.e.*, the value function in the generalized model is "concave" in the number of time periods.

For these models, we assume knowledge of the consumer's long-run consumption probabilities for the full assortment of product alternatives in store. This requirement is not as onerous as it might seem, as few products from the full assortment are typically considered (Hauser and Wernerfelt 1990, Roberts and Lattin 1991). We also assume that a consumer's future preferences (utilities) are uncertain and can be described by a standard random utility framework. This is consistent with the work of Guo (2010) and Walsh (1995).

Applied researchers doing behavioral research could use our model as a rational baseline for shopping and consumption decisions when investigating variety seeking (e.g., Simonson 1990, Read and Loewenstein 1995) or state dependence (e.g., Guo 2010). Our closed-form optimal consumption policy would also be useful in structural models of multiproduct shopping. Because our model permits easy estimation of the consumer's valuation of any n -pack, applied researchers

can use those estimates as inputs for other discrete choice models that predict n -pack selection from the modeler's perspective (see Guo 2010 for an example). Indeed, embedding our model within a larger analytical framework—potentially in an operational setting—offers significant application potential.

While this research was not intended to create a decision support tool, manufacturers and retailers could benefit from using our models to develop specific n -packs. Most n -packs are composed of several units of a single product alternative (e.g., 6-packs of beer or carbonated beverages); others are composed of multiple product alternatives from a single manufacturer (e.g., variety packs of single-serve cereals or yogurts). In either case, the n -pack offered may not be optimal for an individual consumer. Our analysis provides a framework to determine how individual consumers or consumer segments would value different n -packs. The value function can be optimized over the space of all possible n -packs to predict the customized n -pack that a given consumer would choose in the shopping stage. Though this optimization problem is inherently difficult to solve using off-the-shelf software, we develop a greedy swapping algorithm that computes the optimal n -pack in at most n swaps.

We have proposed our model as an assortment optimization model at the consumer level. However, it may be applied to other problems. As a case in point, we describe an application for maximizing auction revenues at the end of the paper. Also, while the current paper focuses solely on the analytical aspects of our dynamic models, the authors have already conducted laboratory experiments to confirm predictions the canonical model makes regarding rational consumer behavior. The data collected from these experiments is the basis for a companion paper covering empirical aspects of our models. Some of the data collected from these experiments are used in Sections 3.4 and 3.5.

2. Literature Review

Consumer psychologists and economists have long recognized that preference uncertainty affects consumers' product choices (Pessemier 1978, March 1978, Kreps 1979, Kahneman and Snell 1990). Simonson (1990) was among the first consumer psychologists to study the effect of preference uncertainty on shopping decisions, finding what has come to be known as the diversification bias (cf. Read and Loewenstein 1995). In a series of three experiments, Simonson showed that consumers systematically seek more variety (measured by the absolute number of different product alternatives selected) when choosing an assortment of products for the future compared to choosing each product sequentially at the time of consumption. This research stream has generated additional empirical results. For example, Simonson and Winer (1992) used

scanner panel data to show that increasing the size of a retailer's assortment increases the variety of flavors consumers select. Salisbury and Feinberg (2008) proposed that diversification may involve a rational response to preference uncertainty, in addition to variety seeking. Because our model is based exclusively on utility-maximizing behavior in the presence of preference uncertainty, it provides an appropriate baseline against which to evaluate variety seeking (or positive state dependence) in consumption decisions.

In an article titled "The Lure of Choice," Bown et al. (2003) found that people prefer to preserve options for the future, even when doing so leads to less desirable outcomes. Our canonical model strongly supports this finding. The experiments reported in their study involve two-stage choices, where only a single item is chosen in the second stage. Similar two-stage choice models have been applied to consideration set formation (Hauser and Wernerfelt 1990, Roberts and Lattin 1991) and to choice among retail assortments (Kahn and Lehmann 1991). Like the models we propose herein, these two-stage models specify Gumbel-distributed errors to represent preference uncertainty.

Guo (2010) developed a structural econometric model for consumers' choice of assortments (n -packs). His model allows for consumption flexibility due to future preference uncertainty as well as state dependence; our model addresses only the former. Guo estimated his model on scanner panel data for yogurt purchases. Because consumption data was not available, Guo estimated the consumer's valuation of each assortment (what we call n -packs) using simulation. This involved simulating error streams for each alternative over the consumption horizon and assuming that the consumer selects the alternative offering the highest utility on each consumption occasion. We note that such a consumption policy is plausible but not optimal. Guo found that allowing for both future preference uncertainty and state dependence offers better in- and out-of-sample fits for the scanner panel data than more restricted nested models. However, his parameter estimates indicate positive state dependence—this is the opposite of variety seeking, which is received wisdom in consumer psychology (Simonson 1990, Read and Loewenstein 1995). Guo also found that consumers make consistent multiproduct purchases; that is, they purchase horizontally varied sets of products but purchase similar sets of products over time. In an earlier study, Guo (2006) determined that consumption flexibility, due to preference uncertainty, also affects firm decisions about product variety and pricing. Using a duopoly model, Guo identified the conditions under which consumers purchase multiple competing products. He found that, if consumers have relatively homogeneous preferences, firms can actually make lower

profits by falling into a "flexibility trap" by pricing to attract primary demand.

The work that is closest to ours is due to Walsh (1995). In this paper, the author modeled consumption decisions for assortments with two product alternatives. Both alternatives' future utilities are random, and the problem reduces to an equivalent one in which one alternative has random utility and the other has constant utility (a reduction that only works for assortments with precisely two alternatives). Assuming that consumers are forward looking, Walsh developed dynamic equations that describe optimal consumption behavior and the associated value function. Although the form of the policy and the value function are not available in closed form, Walsh's analysis yielded three interesting findings: (i) consumers may not choose the alternative offering the greatest utility on a particular consumption occasion; (ii) more inventory of an alternative makes it more likely to be selected; and (iii) adding an additional unit to the assortment causes the utility of that assortment to increase by more than the expected utility of the item added. Our canonical model generalizes Walsh's findings (and adds some refinements) while enabling normative predictions for shopping decisions. Further, our generalized model (including an outside option for consumption) demonstrates that the canonical model represents a boundary solution. Compared to Walsh's model, our models (i) apply to n -packs of any size and with any number of product alternatives; (ii) result in a closed-form value function that can be maximized to determine each consumer's optimal n -pack; and (iii) are based on marginal choice probabilities and so can be customized to individual consumers and used for decision support. The trade-off we make is in using the multinomial logit (MNL) framework—deterministic utility plus Gumbel-distributed errors—to model future utilities; Walsh used a general error distribution. Given the ubiquity of MNL in discrete choice and assortment planning models, we feel that this trade-off is justified.

Another related vein of research involves assortment optimization in the revenue management literature. MNL also plays a prominent role in this research. One of the earliest papers in this vein is due to van Ryzin and Mahajan (1999a), who used MNL embedded in the demand model of a newsvendor problem and derived optimal profit functions under several reasonable assumptions. The authors showed that the profit-maximizing assortment is some subset of the most popular products (the most popular products have the highest probability of being selected). The authors used the concept of majorization to derive sufficient conditions that ensure the profits of one category dominate those of another. In a subsequent paper, Talluri and van Ryzin (2004) introduced a dynamic model and developed conditions on the choice probabilities

that ensure the optimal assortment is some contiguous set of the highest fare products (the “nested by fare order” property). They developed necessary and sufficient conditions that once again involve the concept of majorization, and the MNL choice model was shown to satisfy these conditions. More recently, Rusmevichientong and Topaloglu (2012) showed that these results remain valid for MNL in the presence of parameter uncertainty (for the choice probabilities) and a capacity constraint.

In contrast to these papers, our model addresses assortment optimization at the consumer level and not the retailer level. Moreover, our model does not use MNL choice probabilities to capture consumer demand, but instead uses the random utility framework of MNL to capture preference (utility) fluctuations for a given consumer over time. Our model also focuses on the combination of alternatives and quantities that comprise a consumer’s optimal n -pack. Majorization plays an important role in our work as well, but we do not use it as an assumption to prove a theoretical result. Rather, we find that a stronger form of majorization, what we have called “strong majorization,” characterizes the relationship between the optimal solutions of our two main models.

3. Expected Utility of an n -Pack: The Canonical Model

3.1. Assumptions

Consistent with the extant literature, we begin by assuming (in this section) that the consumer selects an alternative from a preselected n -pack on each consumption occasion. There are M distinct product alternatives available in the product category (the full assortment available in store), although only m alternatives are represented in the n -pack ($m \leq n, m \leq M$). The deterministic utility parameters for each alternative are U_i ($i = 1, 2, \dots, M$). These parameters could be a function of many things; however, we take them to be fixed for ease of exposition. On any particular consumption occasion t , the utility that consumer c receives from a particular alternative i is $U_{ci} + \epsilon_{cit}$, where the random errors ϵ_{cit} are assumed to be independent Gumbel distributed with CDF $F(z) = \exp(-e^{-(z-\mu)/\beta})$. The errors account for a variety of unmodeled factors that affect consumption decisions. For example, a consumer might prefer vegetable soup on most consumption occasions but prefer chicken soup when feeling ill—this would be captured in the error term. Like Walsh (1995) and Guo (2010), we assume these errors become known to the consumer at the time of consumption but not before. Given the canonical model’s assumption that one unit is consumed per period, we must have $t = 1, 2, \dots, n$ periods in the consumption horizon.

Without loss of generality, we may assume that the problem has been normalized so that the errors are

standard Gumbel with $\mu = 0$ and $\beta = 1$ (observe that $U_{ci} + \epsilon_{cit} \geq U_{cl} + \epsilon_{clt}$ if and only if $(1/\beta)U_{ci} + (1/\beta) \cdot (\epsilon_{cit} - \mu) \geq (1/\beta)U_{cl} + (1/\beta)(\epsilon_{clt} - \mu)$, but $(1/\beta)(\epsilon_{cit} - \mu)$ is standard Gumbel for all (c, i, t)). The expectation of a standard Gumbel is $E(\epsilon_{cit}) = \int_0^\infty \ln(z)e^{-z} dz$; this is Euler’s constant and denoted by γ . The expected utility of each product is therefore $E(U_{ci} + \epsilon_{cit}) = U_{ci} + \gamma$. Without loss of generality, alternatives are ordered $U_{c1} \geq U_{c2} \geq \dots \geq U_{cM}$. In what follows, we suppress the subscripts on the consumer (c) and the consumption period (t) to improve readability.

An n -pack of substitutable products can be mathematically represented by a vector of integer quantities (k_1, k_2, \dots, k_M) , $k_i \in \mathbb{N}^0 = \{0, 1, 2, \dots\}$, and $\sum_i k_i = n$. Returning to the soup example with vegetable in the vector’s first position and chicken in the second, the consumer might select the pack $(3, 1, 0, \dots, 0)$, which means three cans of vegetable and one can of chicken ($n = 4, m = 2$). Once an n -pack is selected, we assume there is no replenishment of inventory until all units are consumed. The optimal value function, which is the expected utility received by following an optimal consumption policy, is denoted by $V(k_1, k_2, \dots, k_M)$.

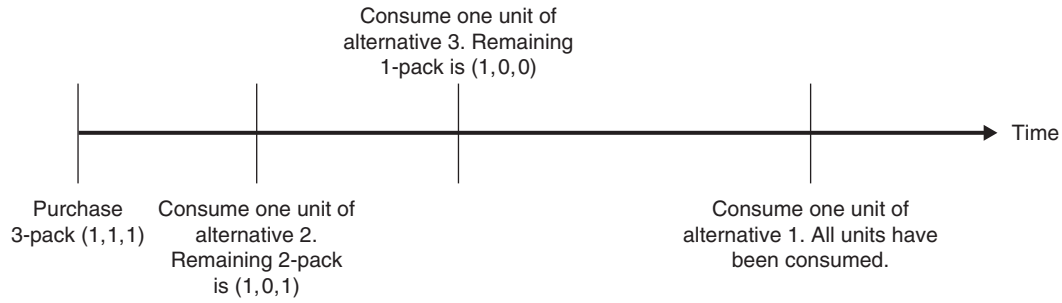
A timeline that illustrates the dynamics for consuming a 3-pack consisting of one unit of alternative 1, one unit of alternative 2, and one unit of alternative 3 is shown in Figure 1. We have assumed the consumption sequence is alternative 2, then alternative 3, then alternative 1. In practice, this sequence would be determined by the optimal consumption policy, which is formally described later in Theorem 1.

3.2. Consuming a Given n -Pack: The Recursion Equation

Recall that the consumption stage is the second stage of our two-stage model (selection/purchase is first, consumption is second), and this stage will involve dynamic programming. Let us first consider the simple case for $M = 2$ product alternatives, labeled 1 and 2 ($U_1 \geq U_2$). The smallest 1-packs for consumption are $(1, 0)$ and $(0, 1)$, and it is clear $V(1, 0) = U_1 + \gamma$ and $V(0, 1) = U_2 + \gamma$. Let us suppose that we have calculated $V(\cdot)$ for all $(n - 1)$ -packs having two or fewer product alternatives. Now consider all n -packs having two or fewer alternatives. The two least diversified n -packs have expected values $V(n, 0) = n \times (U_1 + \gamma)$ and $V(0, n) = n \times (U_2 + \gamma)$. For all remaining n -packs (k_1, k_2) (with $k_1 > 0, k_2 > 0$ and $k_1 + k_2 = n$), we must consider both the current consumption utility, $U_1 + \epsilon_1$ versus $U_2 + \epsilon_2$, and the expected future utility from the remaining items, $V(k_1 - 1, k_2)$ versus $V(k_1, k_2 - 1)$. This means we would (strictly) prefer alternative 1 on the first consumption occasion if and only if

$$U_1 + \epsilon_1 + V(k_1 - 1, k_2) > U_2 + \epsilon_2 + V(k_1, k_2 - 1),$$

Figure 1. Timeline for Consuming the 3-pack (1, 1, 1)



and we would (strictly) prefer alternative 2 if the inequality were reversed. Ties can be broken arbitrarily. To simplify notation, let us define the constant

$$a(k_1, k_2) = U_1 - U_2 + V(k_1 - 1, k_2) - V(k_1, k_2 - 1), \quad (1)$$

so that the optimal policy becomes: choose alternative 1 if $a(k_1, k_2) + \epsilon_1 > \epsilon_2$, otherwise choose alternative 2. Using this optimal policy, we can calculate the expected optimal utility, $V(k_1, k_2)$, as

$$\begin{aligned} V(k_1, k_2) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\epsilon_1 + a(k_1, k_2)} (U_1 + \epsilon_1 + V(k_1 - 1, k_2)) \\ &\quad \cdot \exp(-e^{-\epsilon_1} - e^{-\epsilon_2}) e^{-\epsilon_1 - \epsilon_2} d\epsilon_2 d\epsilon_1 \\ &\quad + \int_{-\infty}^{\infty} \int_{-\infty}^{\epsilon_2 - a(k_1, k_2)} (U_2 + \epsilon_2 + V(k_1, k_2 - 1)) \\ &\quad \cdot \exp(-e^{-\epsilon_1} - e^{-\epsilon_2}) e^{-\epsilon_1 - \epsilon_2} d\epsilon_1 d\epsilon_2 \\ &= \ln(e^{a(k_1, k_2)} + 1) + U_2 + \gamma + V(k_1, k_2 - 1). \end{aligned} \quad (2)$$

Using the recursion in (2), we can determine the value of the n -pack (1, 1). Here, $a(1, 1) = 0$, and so

$$V(1, 1) = U_1 + U_2 + \ln(2) + 2\gamma.$$

Continuing to use (2) in this fashion, we obtain the following valuations:

$$\begin{aligned} V(2, 1) &= 2U_1 + U_2 + 3\gamma + \ln(3) \quad (a = \ln(2)), \\ V(1, 2) &= U_1 + 2U_2 + 3\gamma + \ln(3) \quad (a = -\ln(2)), \\ V(3, 1) &= 3U_1 + U_2 + 4\gamma + \ln(4) \quad (a = \ln(3)), \\ V(2, 2) &= 2U_1 + 2U_2 + 4\gamma + \ln(2) + \ln(3) \quad (a = 0), \\ V(1, 3) &= U_1 + 3U_2 + 4\gamma + \ln(4) \quad (a = -\ln(3)), \text{ etc.} \end{aligned}$$

For a given n -pack having $k_1 \geq 0$ units of alternative 1, the general relationship is

$$V(k_1, k_2) = \ln(n!) - \ln(k_1!) - \ln(k_2!) + k_1 U_1 + k_2 U_2 + n\gamma. \quad (3)$$

Moreover, by the definition of $a(k_1, k_2)$ (see (1))

$$\begin{aligned} a(k_1, k_2) &= U_1 - U_2 + V(k_1 - 1, k_2) - V(k_1, k_2 - 1) \\ &= \ln\left(\frac{n-1}{k_1-1}\right) - \ln\left(\frac{n-1}{k_1}\right) \\ &= \ln\left(\frac{k_1}{k_2}\right) = \ln(k_1) - \ln(k_2). \end{aligned}$$

The optimal policy for consumption reduces to an intuitively appealing condition: “consume alternative 1 if $\epsilon_1 + \ln(k_1) > \epsilon_2 + \ln(k_2)$, otherwise consume alternative 2.” This policy implies that the consumer should not necessarily choose the alternative that maximizes their utility at each consumption occasion, but should instead consider the magnitude of each product’s error term (ϵ_i) adjusted for the quantity of each product on hand ($\ln(k_i)$). Observe that this policy effectively preserves alternatives (probabilistically speaking) for future consumption occasions. For example, a consumer with one unit of alternative 1 and four units of alternative 2 would only consume the last unit of alternative 1 if $-\ln(4) + \epsilon_1 > \epsilon_2$, which occurs with probability 0.2 (see Proposition 1). This policy is consistent with the empirical findings of Bown et al. (2003).

The foregoing dynamic analysis can be generalized to any number of alternatives. (The proof parallels that of the generalized model and is treated there.)

Theorem 1 (Optimal Consumption and Value of an n -Pack, Canonical Model). *Consider an n -pack that includes k_i units of alternative i , $i = 1, \dots, M$. Assume the consumer must select an alternative from their remaining pack on each future consumption occasion. Then the optimal policy for each consumption occasion is to select the alternative that maximizes $\ln(k_i) + \epsilon_i$, and the optimal expected utility (value) for consuming the entire n -pack is*

$$\begin{aligned} V(k_1, k_2, \dots, k_M) &= \ln(n!) - \ln((k_1)! \dots (k_M)!) \\ &\quad + k_1 U_1 + \dots + k_M U_M + n\gamma. \end{aligned} \quad (4)$$

Observe that the linear component $k_1 U_1 + k_2 U_2 + k_3 U_3 + \dots + k_M U_M + n\gamma$ in (4) is simply the expected utility of consuming the entire n -pack if the sequence of consumption were determined in advance. The logarithmic terms $\ln(n!) - \ln((k_1)! (k_2)! \dots (k_M)!)$ together

reflect the additional expected utility of having the freedom to consume products in whatever order one chooses—we will call this additional expected utility a “choice premium.” This premium can also be interpreted as how much value the n -pack provides in terms of hedging against future preference uncertainty.¹ The term $\ln(n!)$ captures the effect of an n -pack’s size, while the term $-\ln((k_1!)(k_2!) \dots (k_M!))$ captures the effects of both variety and inventory. For an n -pack with a fixed number of units n , the choice premium is increased by including more alternatives and/or “flattening” the distribution of alternatives (k_i). The maximum choice premium is $\ln(n!)$, which is realized when there is exactly one unit of n distinct alternatives; the minimum choice premium is zero, which is realized when $k_i = n$ for some alternative i . For any n -pack, then, the ratio of the choice premium to $\ln(n!)$ can be interpreted as the proportion of the available choice premium captured by that n -pack.

The intuition behind this optimal policy can be made clear by considering a simplified case. Let us suppose a given consumer has a 2-pack consisting of one unit of alternative 1 (their favorite) and one unit of alternative 2 (their second favorite), with $U_1 > U_2$. On the first consumption occasion, suppose the observed error terms are e_1 and e_2 with $e_2 > e_1$. Even if $U_2 + e_2 < U_1 + e_1$, alternative 2 represents the better consumption choice. This is because U_1 and U_2 are fixed, and so the *sum* of the realized errors (one now, one later) will ultimately decide the total utility received over both consumption occasions. Because the future error is drawn from the same error distribution, taking the largest error available now is the optimal action. In short, the optimal policy is about matching products with occasion-specific consumption utilities. When more than one unit of inventory is available, an adjustment is necessary. Indeed, if $k_i > 1$ units of an alternative are present, then k_i realizations of ϵ_i must be anticipated over the remaining consumption horizon; this necessarily lowers the bar on the (realized) value of ϵ_i needed to make alternative i the best match. This explains the $\ln(k_i)$ adjustment in the optimal policy.

We can also calculate the probability that a consumer will choose a particular alternative from their remaining n -pack at each consumption occasion.

Proposition 1 (The Proportionality Principle). *The probability of choosing alternative i is $k_i / (\sum_{l=1}^M k_l)$, where k_i is the current quantity of alternative i remaining in the n -pack.*

This proposition follows from rearranging the probability statement and simplifying terms:

$$\begin{aligned} \text{Prob}(i) &= \text{Prob}\{U_i + V(k_1, \dots, k_{i-1}, k_i - 1, k_{i+1}, \dots, k_M) + \epsilon_i \\ &\geq U_j + V(k_1, \dots, k_{j-1}, k_j - 1, k_{j+1}, \dots, k_M) + \epsilon_j \\ &\quad \forall j \neq i\} \end{aligned}$$

$$\begin{aligned} &= \text{Prob}(V(k_1, \dots, k_M) + \ln(k_i) + \epsilon_i \geq V(k_1, \dots, k_M) \\ &\quad + \ln(k_j) + \epsilon_j \quad \forall j \neq i) \\ &= \text{Prob}(\ln(k_i) + \epsilon_i \geq \ln(k_j) + \epsilon_j \quad \forall j \neq i) = \frac{k_i}{\sum_{l=1}^M k_l}. \end{aligned}$$

The last equality follows from the standard logit probability formula with the customary utility parameter “ U_i ” replaced by $\ln(k_i)$.

3.3. Identifying a Consumer’s Optimal n -Pack

Given a consumer’s utility parameters U_i (as can be estimated from purchase histories or using preference elicitation methods), the value function in (4) can then be optimized over all possible integer quantities (k_1, k_2, \dots, k_M) ($k_i \geq 0, \sum_{i=1}^M k_i = n$) to obtain the consumer’s *optimal n -pack*, $(k_1^*, k_2^*, \dots, k_M^*)$. The optimal pack represents the solution to the first stage (the selection/shopping stage) of our two-stage problem (selection and then consumption). Figures 2(a) and 2(b) show the optimal n -packs of sizes $n = 2$ and $n = 3$. The optimal n -packs vary by region, depending on *differences* in their ranked utilities; this is because translating all utilities by a constant translates all n -pack values by a constant as well. We assume that the utilities are ordered so that $U_1 \geq U_2 \geq U_3$, fixing $U_1 = 0$ for identification purposes. For $n = 2$, the distribution of product utilities is captured in the difference $U_1 - U_2$, which we plot on the horizontal axis. Figure 2(a) shows that the 2-pack (1, 1) is optimal in the region $0 \leq U_1 - U_2 \leq \ln(2)$,

Figure 2(a). Optimal n -Pack by Region for $n = 2$





Figure 2(b). Optimal n -Pack by Region for $n = 3$



where the relative preference for U_1 is weaker, while $(2, 0)$ is optimal in the region $U_1 - U_2 > \ln(2)$, where the relative preference for U_1 is stronger. For $n = 3$, the distribution of utilities depends on both $U_1 - U_2$ and $U_2 - U_3$; the latter is plotted on the vertical axis. Figure 2(b) shows that the 3-pack $(1, 1, 1)$ is optimal when the relative differences in utility, both $U_1 - U_2$ and $U_2 - U_3$, are sufficiently small; $(2, 1, 0)$ is optimal if $U_1 - U_2$ is sufficiently small but $U_2 - U_3$ is sufficiently large; and $(3, 0, 0)$ is optimal if $U_1 - U_2$ is sufficiently large. Observe that, without knowing a consumer's particular utility parameters, it would be impossible to determine if a consumer's n -pack selection represents a rational decision or not. To simplify the exposition and avoid set notation, we henceforth assume the optimal n -pack is unique. When this is not the case, our subsequent results apply to each optimal solution or the set of optimal solutions (as appropriate).

The value function is separable and concave; however, the optimization is done over the lattice points of a scaled simplex ($\sum_{i=1}^M k_i = n, k_i \in \mathbb{N}^0$). This problem is less than ideal for many solvers, and we experienced significant difficulties trying to solve this using off-the-shelf software (in large part because many of the k_i^* are precisely zero, and the optimizer would terminate if these slipped into negative territory). Fortunately, it is quite easy to solve this problem using a simple optimization algorithm based on swaps that is guaranteed to produce a global optimal solution in at most n swaps. This algorithm is described next.

Suppose we have an incumbent feasible solution k_1, k_2, \dots, k_M ($\sum_{i=1}^M k_i = n, k_i \in \mathbb{N}^0$) and want to improve it. One way is to increase a nonnegative k_i by one unit and decrease a currently positive k_j by one unit. The net change in the objective function caused by this one-unit swap is

$$U_i - \ln(k_i + 1) - U_j + \ln(k_j), \quad (5)$$

and this suggests the following greedy technique. Interpreting $\ln(0) = \lim_{a \rightarrow 0^+} \ln(a) = -\infty$, calculate the optimal indices

$$i^* = \arg \max_i \{U_i - \ln(k_i + 1), i = 1, \dots, M\}, \quad (6)$$

$$j^* = \arg \min_j \{U_j - \ln(k_j), j = 1, \dots, M\}. \quad (7)$$

Ties in the maximum can be broken by selecting the alternative with the smallest index; ties in the minimum can be broken by selecting the alternative with the largest index. Then define the difference

$$\Delta \equiv U_{i^*} - \ln(k_{i^*} + 1) - U_{j^*} + \ln(k_{j^*}). \quad (8)$$

If $\Delta > 0$, then increase k_{i^*} by one unit and decrease k_{j^*} by one unit; if no such combination exists, then stop. Observe that if $\Delta > 0$, the objective function increases

by a strictly positive amount Δ after each swap. Otherwise, $\Delta \leq 0$ and we must have achieved the stopping condition

$$\begin{aligned} & \max_i \{U_i - \ln(k_i + 1), i = 1, \dots, M\} \\ & \leq \min_j \{U_j - \ln(k_j), j = 1, \dots, M\}. \end{aligned} \quad (9)$$

Theorem 2 (Optimality of the Swapping Algorithm). *Given any starting solution (k_1, k_2, \dots, k_M) with $\sum_{i=1}^M k_i = n$, $k_i \in \mathbb{N}^0$, the swapping procedure described above converges to an optimal n -pack in at most n swaps.*

The conditions (9) represent necessary and sufficient conditions for an n -pack to be optimal, and they can be exploited in several ways. One of these is the connection between the optimal (n)-pack and the optimal ($n + 1$)-pack, which is described next.

Theorem 3. *Let $(k_1^*, k_2^*, \dots, k_M^*)$ represent the consumer's optimal n -pack. Then the consumer's optimal $(n + 1)$ -pack is obtained by identifying the alternative i that maximizes $U_i - \ln(k_i^* + 1)$ and increasing that alternative by one unit.*

The latter theorem not only means we can build larger optimal packs from smaller optimal packs, but it also identifies the marginal unit that should be added to any optimal n -pack. For a retailer looking to encourage a consumer to buy an additional unit, this information would be extremely useful.

The optimality conditions can be used to obtain insights into the structural properties of the optimal n -pack as well. Some of these properties are included in the following theorem, whose proof is contained in the online appendix. (Recall $\ln(0) = \lim_{a \rightarrow 0^+} \ln(a) = -\infty$.)

Theorem 4. *The optimal n -pack, $(k_1^*, k_2^*, \dots, k_M^*)$, satisfies the following:*

- (i) If $k_j^* = 0$, then $k_i^* = 0$ for $U_i < U_j$;
- (ii) $k_1^* \geq k_2^* \geq k_3^* \geq \dots \geq k_M^*$;
- (iii) $(k_i^* + 1)/k_j^* \geq \exp(U_i - U_j) \forall i, j$.

The first condition implies the consumer's optimal n -pack is a contiguous set of their favorite alternatives. This property is reminiscent of the assortment optimization results of van Ryzin and Mahajan (1999b) from the retailer's perspective, which may not be surprising given that our problem is an assortment optimization problem from the consumer's perspective. The second condition requires monotonicity in quantities; higher quantities go with higher utilities. The last property demonstrates that the ratio $(k_i^* + 1)/k_j^*$ must exceed the ratio $(\exp(U_i)/\exp(U_j))$ for any two products i and j ; we recall that $\exp(U_i)/\exp(U_j)$ is the ratio of i and j 's choice probabilities in a logit framework. These properties will be used in our subsequent analyses.

3.4. The Optimal n -Pack and Sequential Choice

The literature on multi-item set selection often uses sequential choice experiments (or simply “sequential choice”) as a benchmark for measuring variety. Sequential choice means the consumer is allowed to choose any alternative from the full assortment (i.e., every alternative in a category that could be selected from a store) on each consumption occasion. This wait-and-see approach means the consumer can observe the random component of utility (ϵ_i) for every alternative in the full assortment immediately before making a consumption decision. The consumer thus maximizes their utility on every consumption occasion and cannot obtain any greater utility than this when consuming n items on n consecutive occasions. We show next that the optimal n -pack selected *a priori* is the same as the most probable set of n items chosen sequentially.

As before, we describe each n -pack using an M dimensional vector (k_1, k_2, \dots, k_M) where $k_i \in \mathbb{N}^0$ represents the integer quantity of alternative i and $\sum_{i=1}^M k_i = n$. The optimal n -pack, denoted by $(k_1^*, k_2^*, \dots, k_M^*)$, maximizes $V(k_1, k_2, \dots, k_M)$ in (4). Consider the M dimensional vector (x_1, x_2, \dots, x_M) where x_i represents the quantity of alternative i consumed sequentially ($\sum_{i=1}^M x_i = n$). In contrast to an n -pack, the vector (x_1, x_2, \dots, x_M) is not selected in advance but rather constructed over n successive consumption occasions; it represents the “final tally” for each alternative after the n th consumption occasion. Define $(x_1^*, x_2^*, \dots, x_M^*)$ to be the vector of quantities that is *most likely* to be consumed when choosing alternatives sequentially from the full assortment.

Proposition 2 (Optimal n -Pack vs. Sequential Choice). *Let $(k_1^*, k_2^*, \dots, k_M^*)$ be the optimal n -pack for a consumer. For the same consumer, let $(x_1^*, x_2^*, \dots, x_M^*)$ be the set of size n that is most likely to be chosen sequentially from the full assortment on n consecutive consumption occasions. Then $(k_1^*, k_2^*, \dots, k_M^*) = (x_1^*, x_2^*, \dots, x_M^*)$.*

To see why the proposition is true, take $U_1 \geq U_2 \geq \dots \geq U_M$ and translate (shift) all U_i by a suitable constant so that $\sum_{i=1}^M e^{U_i} = 1$. This translation merely affects the additive constant ($n\gamma$) of the value function, which has no bearing on the ordering of n -pack values and thus can be ignored. The marginal probability of selecting alternative i from the full assortment is therefore $p_i = e^{U_i} / (e^{U_1} + e^{U_2} + \dots + e^{U_M}) = e^{U_i}$. Note that $\ln(p_i) = U_i$ due to our translation. Then the probability of consuming the set (x_1, x_2, \dots, x_M) via sequential selection is

$$\Pr(x_1, x_2, \dots, x_M) = \frac{n!}{\prod_{i=1}^M x_i!} \prod_{i=1}^M p_i^{x_i}. \quad (10)$$

We then observe that

$$(x_1^*, x_2^*, \dots, x_M^*) = \arg \max_{x_i \in \mathbb{N}^0, \sum_{i=1}^M x_i = n} \frac{n!}{\prod_{i=1}^M x_i!} \prod_{i=1}^M p_i^{x_i}$$

$$\begin{aligned} &= \arg \max_{x_i \in \mathbb{N}^0, \sum_{i=1}^M x_i = n} \ln \left[\frac{n!}{\prod_{i=1}^M x_i!} \prod_{i=1}^M p_i^{x_i} \right] \\ &= \arg \max_{x_i \in \mathbb{N}^0, \sum_{i=1}^M x_i = n} \left\{ \ln(n!) - \ln \left(\prod_{i=1}^M x_i! \right) \right\} \\ &\quad + \sum_{i=1}^M x_i U_i = (k_1^*, k_2^*, \dots, k_M^*). \quad (11) \end{aligned}$$

A common finding in the consumer psychology literature is that sets consumed via sequential choice (the (x_1, x_2, \dots, x_M)) typically exhibit less variety than pre-selected n -packs of the same size (e.g., Simonson 1990, Read and Loewenstein 1995, Salisbury and Feinberg 2008). If consumers select their n -packs optimally, then this would naturally be the case if the *most likely* set selected sequentially exhibited more variety than most other sets selected sequentially (probabilistically speaking). For example, consider a hypothetical situation where a utility-maximizing consumer likes $M = 3$ alternatives equally well. For this consumer, the choice probability (or choice frequency) for each alternative is $\frac{1}{3}$, and the optimal 3-pack is $(1, 1, 1)$, which solves (11). However, there are $3 \times 3 \times 3 = 27$ equally probable permutations that this consumer could consume via sequential selection, and only six of these permutations include all three product alternatives. Therefore, this same consumer selecting alternatives sequentially would naturally consume less than three different product alternatives with probability $21/27$.² Thus, while it might appear that this consumer has selected a 3-pack with too much variety (compared to what would be consumed via sequential choice), this “variety asymmetry” would be quite rational.

To see if this asymmetry was more than just a theoretical possibility, we surveyed 168 business students (61 MBA students and 107 BBA students) and asked them to report the relative frequency with which they would consume their top three snack alternatives (drawn from a larger list of approximately 15 snack alternatives available in local vending machines). Since choice frequencies were only recorded for their top three snacks, we can only compute results for the case $n = M = 3$, identical to our earlier hypothetical scenario. For each 3-tuple of self-reported choice frequencies (favorite, second favorite, third favorite), we can calculate each student’s optimal 3-pack and thus their optimal number of alternatives. Having done this for the sample of 168 students, we found that 26 students had an optimal 3-pack with exactly one alternative; 80 students had an optimal 3-pack with exactly two alternatives; and 62 students had an optimal 3-pack with exactly three alternatives. Using the same self-reported choice frequencies, we found that in 102 of the 168 cases (60.7%), the probability of a student consuming less variety than their optimal 3-pack in a sequential

choice experiment would be greater than their probability of consuming more. Additionally, the average probability of a student consuming less variety than their optimal 3-pack over all 168 cases was computed to be 0.383, whereas the average probability of consuming more variety was computed to be 0.151. This offers additional support for the conjecture that consuming less variety in sequential choice experiments may be a consequence, in part, of probabilistic principles stemming from rational decision making. Additional work is under way to rigorously test this and other conjectures regarding variety.

3.5. Robustness of the Model: Assessing the Impact of the Gumbel Assumption

We conducted several numerical experiments to ensure that our results were not overly dependent on the assumption of a Gumbel distribution. We provide a summary of the results here; the reader is referred to Online Appendix B for details.

Two additional error distributions were selected, the uniform and the normal. In the first numerical experiment, we analyzed 3-packs based on utilities calculated from actual choice data. In the second experiment, we analyzed 6-packs based on utilities calculated from simulated choice probabilities. In total, 106 test problems were analyzed, 6 involving 3-packs and 100 involving 6-packs. In general, there were virtually no meaningful discrepancies in valuations of n -packs that consumers would actually choose. Consequently, the error distribution appears to have little if any impact on a consumer's valuation of their most preferred n -packs (say the ratings for their top 10–20 n -packs). There were some discrepancies in valuations for problems that included alternatives the consumer would rarely (if ever) select; i.e., test problems that included one or more "unpopular alternatives" with choice probabilities approaching zero. In such cases, the corresponding utility parameter in the normal and Gumbel models becomes unbounded from below, whereas the utility parameter in the uniform model does not (it is always bounded). For this reason, the valuations for the Gumbel and normal distributions tended to track each other closely for all n -packs, whereas the valuations for the uniform tended to diverge for those n -packs that included extremely unpopular alternatives. Because unpopular n -packs are implicitly ignored when determining a consumer's optimal n -pack, this has no impact on our results.

4. Expected Utility of an n -Pack: The Generalized Model

4.1. The Optimal Value Function

In this section, we assume the consumer may select the outside option on any consumption occasion and thus

reject all items remaining in their pack. The introduction of an outside option effectively allows for different consumption rates. It does, however, add complexity compared to the canonical case.

We will again use n to denote the number of total units in the n -pack, and M to denote the total number of distinct alternatives available. Let $k_i \in \mathbb{N}^0 = \{0, 1, 2, 3, \dots\}$ represent the number of units of alternative i in the n -pack. The utility parameter for each alternative is denoted by U_i for $i = 0, 1, \dots, M$ (note that we include U_0 , the utility of the outside option). The number of consumption occasions is denoted by t , which is also the number of time periods in our dynamic analysis, and the value function with t periods to go is denoted by $V_t(k_1, k_2, \dots, k_M)$, which means consumption periods in the generalized model are numbered *backward* (as is frequently done in dynamic programming models). The value function in the terminal (salvage) period (period 0) is $V_0(k_1, k_2, \dots, k_M) = 0$. One can think of this as an n -pack becoming worthless if its expiration date is reached without having been consumed.

The "no consumption" option is represented by the subscript zero, and we can represent the set that is ultimately consumed (or "realized") after t consumption occasions by an $M + 1$ dimensional *consumption vector* (x_0, x_1, \dots, x_M) , $x_i \in \mathbb{N}^0$. As was the case in Section 3.4, the consumption vector is simply the "final tally" of units consumed for each alternative, including the number of times the no consumption option was invoked (this is recorded in the vector's first position). Define the index set of vectors

$$I_t(y_0, y_1, y_2, \dots, y_M) = \left\{ (x_0, x_1, \dots, x_M) : \sum_{i=0}^M x_i = t; 0 \leq x_i \leq y_i, x_i \in \mathbb{N}^0, i = 0, 1, \dots, M \right\} \quad (12)$$

Observe that the set I_t requires $M + 1$ inputs $(y_0, y_1, y_2, \dots, y_M)$ that serve as upper bounds on all possible $M + 1$ dimensional consumption vectors for a pack having y_i units of alternative i . As an example, suppose there are only $M = 2$ products available, alternative 1 and alternative 2. A 3-pack having two units of alternative 1 and one unit of alternative 2 would lead to

$$\begin{aligned} (\text{four periods}) \quad I_4(4, 2, 1) &= \{(4, 0, 0); (3, 1, 0); (3, 0, 1); \\ &\quad (2, 1, 1); (2, 2, 0), (1, 2, 1)\}; \\ (\text{three periods}) \quad I_3(3, 2, 1) &= \{(3, 0, 0); (2, 1, 0); (2, 0, 1); \\ &\quad (1, 1, 1); (1, 2, 0); (0, 2, 1)\}; \\ (\text{two periods}) \quad I_2(2, 2, 1) &= \{(2, 0, 0); (1, 1, 0); (1, 0, 1); \\ &\quad (0, 1, 1); (0, 2, 0)\}; \\ (\text{one period}) \quad I_1(1, 2, 1) &= \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}. \end{aligned}$$

Observe that in each of these three sets, the first input (y_0) is taken to equal the number of consumption occasions. This is appropriate since y_0 is the upper bound on the number of times the “no consumption” option could be invoked, which is equal to the number of consumption periods, t . While the number of terms can be quite large, it is bounded independently of the number of consumption occasions t . Indeed, there are at most $\prod_{i=1}^M (k_i + 1)$ elements in $I_t(t, k_1, k_2, \dots, k_M)$, which corresponds to the number of distinct subsets of (k_1, k_2, \dots, k_M) padded by the appropriate number of “outside option” selections to bring the total number of selections to t . This upper bound is obtained for all $t \geq n$.

Theorem 5 (Optimal Value of an n -Pack, Generalized Model). *At each consumption occasion, assume the consumer can choose a product from the n -pack or select an outside option. The optimal expected value function over t consumption periods (assuming an optimal policy is followed each period) is*

$$V_t(k_1, k_2, \dots, k_M) = \ln \left[\sum_{(x_0, x_1, \dots, x_M) \in I_t(t, k_1, k_2, \dots, k_M)} \frac{t!}{x_0! x_1! \dots x_M!} e^{\sum_{j=0}^M x_j U_j} \right] + t\gamma. \quad (13)$$

The optimal policy at each consumption occasion t is to select, among the available alternatives, the one that maximizes current utility plus expected utility-to-go; i.e., the one that maximizes $U_0 + \epsilon_0 + V_{t-1}(k_1, k_2, \dots, k_M)$, $U_j + \epsilon_j + V_{t-1}(k_1, k_2, k_j - 1, \dots, k_M)$ for $k_j > 0$. Unlike the canonical version, there is no additional simplification in the optimal policy.

While the value function is somewhat complicated, it can be simplified under the assumption $t \geq n$, which we would expect to hold in practice. As noted earlier, there are a constant $\prod_{i=1}^M (k_i + 1)$ terms in the summation of (13). Additionally, defining the shifted parameters $U'_i = U_i - U_0$ (so that the outside option has utility parameter $U'_0 = 0$), the value function (13) can be expressed as

$$V_t(k_1, k_2, \dots, k_M) = \ln \left[\sum_{\substack{x_i \leq k_i, i \geq 1 \\ x_0 \geq t - \sum_{i=1}^M x_i}} \frac{t!}{(t - \sum_{i=1}^M x_i)!} \cdot \frac{e^{\sum_{j=1}^M x_j U'_j}}{x_1! \dots x_M!} \right] + t(\gamma + U_0). \quad (14)$$

This simplified form is easier to manipulate and is used extensively in Theorem 7.

The terms in the value function (13) generalize the probability interpretation established for the canonical model in Section 3.4. There, we established that the value function for the n -pack (k_1, k_2, \dots, k_M) in the canonical model could be equated to the log-probability of consuming the same set of products in a sequential choice experiment; i.e., an experiment in which the consumer can select the product from the

full assortment that is most preferred on each consumption occasion. This further implies that the utility-maximizing n -pack selected a priori—the consumer’s “optimal n -pack”—is also the set of product alternatives most likely to be consumed in a sequential choice experiment. As in the canonical model, we assume that all utility parameters have been translated so that $\sum_{i=0}^M e^{U_i} = 1$, which we recall simply alters the additive constant γ used in the value function (13). The terms in the summation in (13) then represent the cumulative probability of consuming the n -pack (k_1, k_2, \dots, k_M) or any subset thereof in a sequential choice experiment. (If a subset of the n -pack of size $n' < n$ is consumed, the outside option must have been selected exactly $t - n'$ times.) Moreover, maximizing the value function in the generalized model is equivalent to maximizing the probability of consuming the n -pack (k_1, k_2, \dots, k_M) or any of its subsets in a sequential choice experiment. The main difference in the generalized model is that subsets of the original n -pack must be included in the probability statement because the entire n -pack need not be consumed within t time periods.

The position of U_0 relative to the products U_1, U_2, \dots, U_M has considerable impact on the model (canonical or generalized) that is most appropriate. In applications where the utilities for the products in the n -pack are considerably greater than the utility for the outside option, one would expect the canonical model to work well. This is because the outside option has virtually no chance of being selected as the preferred alternative (unless the pack is exhausted). This could be the case for many products that are consumed on a regular schedule; e.g., cereal in the morning. But for product categories where the utility of the outside option is greater than or equal to the utilities of products in the n -pack, the outside option becomes a viable alternative. One would expect categories of less frequently consumed goods to fit this scenario. The natural question is, then, “How does the optimal n -pack for the generalized model compare to the optimal n -pack for the canonical model?” This is explored next.

4.2. Variety, Consumption Horizons, and the Outside Option

Imagine two consumers, A and B, both of whom like exactly three types of wine: Chardonnay, Merlot, and Cabernet. Both consumers prefer Chardonnay 70% of the time, Merlot 15% of the time, and Cabernet 15% of the time. However, consumer A enjoys a bottle every evening, whereas consumer B enjoys a bottle about once a week. If we assume the opportunity to consume wine presents itself every evening, then consumer A has a very small value for U_0 and thus a high usage rate for wine, whereas consumer B has a much higher value for U_0 and thus a low usage rate for wine. Given this information, which 3-pack of wine should each

Table 1. Optimal 3-Pack for the Wine Example (Consumer B)

	Horizon (t)							
	3	7	14	21	28	35	42	100
Optimal n -pack	(1, 1, 1)	(1, 1, 1)	(2, 1, 0)	(2, 1, 0)	(2, 1, 0)	(2, 1, 0)	(3, 0, 0)	(3, 0, 0)

consumer buy? (We ignore the fact consumer A would probably tend to buy a larger pack size.)

The fact that consumer A always chooses to consume wine is evidence of an intrinsically low value for U_0 , one that is exceeded by U_1 , U_2 , and U_3 . We would expect this consumer to always prefer wine compared to the outside option. Indeed, for any $t \geq n$, as $U_0 \rightarrow -\infty$, the value function of the generalized model (13) can be well approximated by the value function of the canonical model (up to an additive constant). This is because the dominant term in (13) for any n -pack becomes $(t! / ((t-n)! k_1! \dots k_M!)) e^{\sum_{j=1}^M k_j U_j} e^{(t-n)U_0}$, with the remaining $[\prod_{i=1}^M (k_i + 1)] - 1$ terms containing additional powers of e^{U_0} , which diminishes their contribution to V_t as $U_0 \rightarrow -\infty$. In consumer A's case, the utility-maximizing 3-pack is therefore three bottles of Chardonnay, which is the same as in the canonical model.

The fact that consumer B chooses to consume wine infrequently is evidence of an intrinsically higher value for U_0 . Given they select the outside option 6/7 of the time, their value for U_0 exceeds the values for U_1 , U_2 , and U_3 (we used $U_0 = -0.1544$, $U_1 = -2.3011$, $U_2 = U_3 = -3.8415$ so that $\sum_{i=0}^3 e^{U_i} = 1$ in our subsequent calculations). Nevertheless, for large values of t , the results for consumer B are identical to those of consumer A. This is because, as $t \rightarrow \infty$, the dominant term in (13) is again $(t! / ((t-n)! k_1! \dots k_M!)) e^{\sum_{j=1}^M k_j U_j} e^{(t-n)U_0}$, with the remaining $[\prod_{i=1}^M (k_i + 1)] - 1$ terms multiplied by at least an additional factor of $1/(t-n+1)$, which diminishes their contribution to V_t as $t \rightarrow \infty$. Thus, the infinite horizon results are identical to those of the canonical model too.

However, given smaller values of t for consumer B, the optimal 3-pack is not simply three bottles of Chardonnay. Table 1 tracks the optimal 3-pack for different horizon lengths, t . The intuition behind the increased variety for smaller values of t can be illustrated by a simple thought experiment. Suppose that the consumer could select a 3-pack for only one ($t = 1$) consumption opportunity. It should be clear that the optimal 3-pack for the case $t = 1$ would be for the consumer to have exactly one unit of their three favorite products; there is no advantage to having multiple units of any product when $t = 1$, and so maximizing variety maximizes the number of independent draws, which in turn maximizes the chances of getting the best match for this single-period consumption occasion. The case $t = 1$ captures what is typically known in

dynamic programming as the “end of horizon effect.” Over slightly longer horizons, carrying greater variety than (3, 0, 0) would still be optimal to address these end-of-horizon effects. Nevertheless, end-of-horizon effects typically dissipate over sufficiently long horizons, and this dynamic model is no different. As the horizon lengthens, there are sufficient opportunities to match the right item with the right consumption occasion, and the optimal n -pack converges to that of the canonical model.

The relationship between the optimal n -pack in the generalized model and the optimal n -pack in the canonical model can be characterized more precisely. The precise result is stated in the following theorem (where alternatives are again labelled so that $U_1 \geq U_2 \geq \dots \geq U_M$). The condition $t \geq n$ is included to simplify the proof; otherwise, the number of terms in the value function would also depend on t .

Theorem 6. For any horizon $t \geq n$, denote the optimal n -pack of the canonical model by $(k_1^*, k_2^*, \dots, k_M^*)$ and that of the generalized model by $(q_1^*, q_2^*, \dots, q_M^*)$. Then it cannot happen that $q_j^* > k_j^*$ and $q_i^* < k_i^*$ for $j < i$.

The theorem requires that the components of the optimal vectors $k^* = (k_1^*, k_2^*, \dots, k_M^*)$ and $q^* = (q_1^*, q_2^*, \dots, q_M^*)$ follow a strict pattern. Assuming the vectors are not identical, then there is an index, say c , where a change occurs; for $i \leq c$, we must have $k_i^* \geq q_i^*$ (with at least one strict inequality); and for $i > c$, we must have $k_i^* \leq q_i^*$ (with at least one strict inequality). The theorem permits the possibility $k^* = (5, 2, 2, 0, 0)$ and $q^* = (3, 3, 3, 0, 0)$, but not the possibility $k^* = (5, 2, 2, 0, 0)$ and $q^* = (3, 3, 1, 1, 1)$. In the first case, the k_i^* and q_i^* have the componentwise relationship $(>, <, <, \leq, \leq)$, and thus the change in inequalities occurs at $c = 1$. In the second case, the componentwise relationship is $(>, <, >, <, <)$, and no change point c exists. One should think of the index c as a turning point.

The structural relationship between q^* and k^* is closely related to the concept of *majorization*. We say a vector $x \in \mathbb{R}^M$ majorizes (or dominates) a vector $y \in \mathbb{R}^M$ (written $x > y$) provided (i) $\sum_{i=1}^l x_{(i)} \geq \sum_{i=1}^l y_{(i)}$ for $l = 1, 2, \dots, M$, where the notation $z_{(i)}$ refers to the i th largest value in the vector z , and (ii) $\sum_{i=1}^M x_i = \sum_{i=1}^M y_i$. The concept of majorization is weaker than the condition posed in Theorem 6. Consider our previous example involving $k^* = (5, 2, 2, 0, 0)$ and $q^* = (3, 3, 1, 1, 1)$; it is clear that $k^* > q^*$ even though this pair does not satisfy the relationship described by the theorem. The contrast

between the two concepts can be sharpened by looking at the difference in partial sums. Whereas majorization requires that $\sum_{i=1}^l x_{(i)} - \sum_{i=1}^l y_{(i)} \geq 0$ for $l = 1, 2, \dots, M$, the condition $k_i^* \geq q_i^*$ for $i \leq c$ and $k_i^* \leq q_i^*$ for $i > c$ requires that $\sum_{i=1}^l x_{(i)} - \sum_{i=1}^l y_{(i)}$ is nonnegative and nondecreasing for $l \leq c$ and nonincreasing for $l > c$. If we define a piecewise linear function $g(l)$ whose value at the integers is $g(l) = \sum_{i=1}^l x_{(i)} - \sum_{i=1}^l y_{(i)}$ ($l = 1, 2, \dots, M$), majorization requires g to be nonnegative, whereas Theorem 6 requires g to be nonnegative and quasi-concave with a maximum at the change point c . We could find no reference in the literature to this stronger form of majorization, which admits further generalizations (e.g., the condition $x_{(i)} - y_{(i)} \geq x_{(i+1)} - y_{(i+1)}$ would make $g(l)$ concave). We refer to the type of majorization given in Theorem 6 as *strong majorization*.

Definition 1 (Strong Majorization). We say a vector $x \in \mathbb{R}^M$ strongly majorizes a vector $y \in \mathbb{R}^M$ and write $x \succ^s y$ provided:

1. $\sum_{i=1}^M x_i = \sum_{i=1}^M y_i$
2. For some index $c < M$, we have $x_{(i)} \geq y_{(i)}$ for $i \leq c$ and $y_{(i)} \geq x_{(i)}$ for $i > c$.

In the language of strong majorization, Theorem 6 states that for any $t \geq n$, we must have $(k_1^*, k_2^*, \dots, k_M^*) \succ^s (q_1^*, q_2^*, \dots, q_M^*)$. Observe that this condition implies the number of different alternatives included in the optimal n -pack of the generalized model must be at least as large as the number of different alternatives included in the optimal n -pack of the canonical model.

We sought to generalize Theorem 6 by comparing the optimal solutions to (13) between successive time periods. In numerous numerical examples, we found the optimal solution to the $(t + 1)$ period problem strongly majorized the optimal solution to the t period problem. We conjecture that this is always the case. However, we could not find an analytical proof for this condition. Unlike Theorem 6, whose proof exploits the simple optimality conditions (9) associated with the canonical n -pack, the solution to the generalized model is not characterized by simple optimality conditions. Even the necessary conditions for local optimality, which are based on localized unit swaps, are very complex.

Another property that was observed in all of our numerical examples reflects how the value function changes over time for a particular n -pack. Fortunately, this property has an analytical proof (although ours is surprisingly complicated). The condition $t \geq n + 1$ is used just as $t \geq n$ was used in the previous theorem; it implies the number of terms included in the summation of (13) is independent of t (see also (14)).

Theorem 7 (Diminishing Marginal Value). Consider any n -pack $K = (k_1, k_2, \dots, k_M)$ and any time period $t \geq n + 1$. Then $V_{t+1}(K) - V_t(K) \leq V_t(K) - V_{t-1}(K)$.

The property of diminishing marginal value makes intuitive sense. Because the optimal value is based on matching the product with the consumption occasion, adding additional periods should not yield proportional gains in value. Adding additional periods means we are searching for better opportunities in the right-hand tail of the error distribution, and better payoffs should become increasingly difficult to obtain as time increases. Dynamic models in revenue management often require this type of structure, and so the property is an important one if the model is to find additional applications in this area.

5. Summary and Future Research

We have proposed a utility-maximizing model based on consumers' long-run consumption preferences to estimate the value they can expect to receive from an n -pack of substitutable products. Our canonical model predicts that (i) strategic consumers will choose different product alternatives in proportion to their available inventory and (ii) the total value consumers derive from an n -pack increases in the pack's utility parameters but decreases as the distribution of products within the n -pack becomes more concentrated. This result could explain the seemingly excessive variety that has been observed in behavioral experiments on n -pack selection for future consumption. Our generalized model demonstrates that the inclusion of an outside option (effectively reducing the consumption rate), which to our knowledge has never been done, would lead to even more variety in n -pack alternatives and less concentration in n -pack quantities.

Our model assumes that both the consumption utilities U_i and the distribution of stochastic errors ϵ_i are stationary, but this might not always be the case. Allowing for nonstationarity in consumption utilities (i.e., variety seeking, state dependence) or in the stochastic error distribution (i.e., learning) might lead to new results and insights. Our model also assumes that there is no discounting of future utilities, so temporal discounting is another possible area for future research. The direct approach, introducing a discount factor on the expected "value-to-go" function, sacrifices the virtue of a closed-form value function. However, it may still be possible to analyze the value function implicitly or to introduce a different discount mechanism that preserves the closed-form solution. Our model also does not currently allow for any replenishment of inventory by the consumer. Could replenishment be included as yet another choice at each consumption occasion? This would allow the consumer to obtain an alternative they strongly prefer that is not currently available in what remains of their n -pack. Is it possible to obtain an optimal replenishment policy?

There are also many ways to extend our analysis to include various operational considerations. Perhaps the most obvious extension is to assume that the consumption utility parameters U_i are functions of product attributes, such as price. Still another useful extension would be to investigate how n -pack valuation affects consumers' willingness to pay. For example, many retailers implicitly offer the option to purchase an n -pack including only a single product alternative at a low price per unit, or purchase single units of different product alternatives at a higher price per unit. For example, a 6-pack of a single brand/type of beer might cost \$8.99 (\approx \$1.50/beer) while purchasing different beers individually might cost \$1.99/beer. Depending on the difference between consumers' valuation of their optimal n -pack and an n -pack with only their favorite alternative, the retailer may be able to price a "build your own 6-pack" option to extract additional revenues while also increasing consumers' utility. Pricing n -packs and designing promotions that target individual consumers are natural applications for this type of model.

As noted in the introduction, there are some other applications of our model that would require very little additional work. One such application involves auctioning a set of n related products. In this case, the utility parameters U_i are replaced by the expected maximum bid price for alternative i , say W_i . Calculation of the W_i might require a separate model to account for the number of bidders and other factors. The price actually bid for alternative i at auction would then be modeled by $W_i + \epsilon_i$, where we assume the W_i have been normalized (rescaled) so that the error term ϵ_i has a standard Gumbel distribution. A set of n related products having k_i units of alternative i is then put up for auction, and bidders submit individual bids for any or all alternatives they are interested in. At the conclusion of bidding, the maximum bid for alternative i , b_i , is noted, and the winning bid for a single unit of a single alternative is revealed. The winning bid/alternative would be determined by $\max_i(\ln(k_i) + e_i)$ where $e_i = b_i - W_i$. Observe this is the optimal policy in Theorem 1, where now it serves to optimally match the product with the auction. One unit of the winning alternative is removed from the set, and the remaining set of $n - 1$ units is again put up for auction. After n such auctions, all units would be sold. The expected revenue received from this type of auction would be given by the value function in Theorem 1, which exceeds the expected revenue of selling each item in a series of individual auctions. While there are potential model assumption violations that would need to be addressed in practice (independence of the error terms, etc.), this type of auction provides several interesting possibilities. For example, because only one unit is sold at a time, bidders can safely bid on a subset of products without

risking an excessive payout; the auction house would not need to accept poor bids for a particular alternative as long as there are better bidding results for another alternative; in the earlier auction rounds, valuable bidding data is obtained on all unsold alternatives.

We have treated the size of the n -pack as exogenous. However, one could extend our model to include unit price and/or a disutility term for storage, and the optimization would then determine the appropriate pack size along with the optimal pack. Given the simple structure of the value function in the canonical model, the optimization problem thus created might be relatively straightforward to analyze.

Acknowledgments

This work was done while L. Norman served as the associate director of the JC Penney Center for Retail Excellence, Cox School of Business, Southern Methodist University.

Endnotes

¹ Although the choice premium would appear to be fixed, whereas the expected utilities U_i are subject to changes in scale and location, this is not an issue. First, recall that the scale of the utilities has been normalized to make the error terms standard Gumbel. Second, for two n -packs (x_1, x_2, \dots, x_M) and (y_1, y_2, \dots, y_M) , $V(x_1, x_2, \dots, x_M) - V(y_1, y_2, \dots, y_M)$ is independent of any shifts in location of the U_i . Thus, when comparing two n -packs via the optimal expected utility function $V(\cdot)$, only differences in their respective (normalized) utilities and differences in their choice premiums matter.

² Of the 27 outcomes, 18 include two different product alternatives and 3 include only a single product alternative, 1, 2 or 3.

References

- Bown NJ, Read D, Summers B (2003) The lure of choice. *J. Behav. Decision Making* 16(4):297–308.
- Guo L (2006) Consumption flexibility, product configuration, and market competition. *Marketing Sci.* 25(2):116–130.
- Guo L (2010) Capturing consumption flexibility in assortment choice from scanner panel data. *Management Sci.* 56(10):1815–1832.
- Hauser JR, Wernerfelt B (1990) An evaluation cost model of consideration sets. *J. Consumer Res.* 16(4):393–408.
- Kahn BE, Lehmann DR (1991) Modeling choice among assortments. *J. Retailing* 67(3):274–299.
- Kahneman D, Snell J (1990) Predicting utility. Insights in Decision Making: A Tribute to Hillel J. Einhorn. Hogarth RM, ed. (University of Chicago Press, Chicago), 295–310.
- Kreps DM (1979) A representative theorem for "preference for flexibility." *Econometrica* 47(3):565–577.
- March JG (1978) Bounded rationality, ambiguity, and the engineering of choice. *Bell J. Econom.* 9(2):587–608.
- Pessemier EA (1978) Stochastic properties of changing preferences. *Amer. Econom. Rev.* 68(2):380–385.
- Read D, Loewenstein G (1995) Diversification bias: Explaining the discrepancy in variety seeking between combined and separated choices. *J. Experiment. Psych. Appl.* 1(1):34–49.
- Roberts JH, Lattin JM (1991) Development and testing of a model of consideration set composition. *J. Marketing Res.* 28(4):429–440.
- Rusmevichientong P, Topaloglu H (2012) Robust assortment optimization in revenue management under the multinomial logit choice model. *Oper. Res.* 60(4):865–882.
- Salisbury LC, Feinberg FM (2008) Future preference uncertainty and diversification: The role of temporal stochastic inflation. *J. Consumer Res.* 35(2):349–359.

- Simonson I (1990) The effect of purchase quantity and timing on variety-seeking behavior. *J. Marketing Res.* 27(2):150–162.
- Simonson I, Winer RS (1992) The influence of purchase quantity and display format on consumer preference for variety. *J. Consumer Res.* 19(1):133–138.
- Talluri K, van Ryzin G (2004) Revenue management under a general discrete choice model of consumer behavior. *Management Sci.* 50(1):15–23.
- van Ryzin G, Mahajan S (1999a) On the relationship between inventory costs and variety benefits in retail assortments. *Management Sci.* 45(11):1496–1509.
- van Ryzin GJ, Mahajan S (1999b) Retail inventories and consumer choice. *Quantitative Models for Supply Chain Management* (Springer, New York), 491–551.
- Walsh JW (1995) Flexibility in consumer purchasing for uncertain future tastes. *Marketing Sci.* 14(2):148–165.