Do Consumers Hedge Against Future Preference Uncertainty? Testing a Theory of Shopping and Consumption

August 30, 2017

Abstract

This paper presents empirical evidence from two behavioral experiments and from multi-market panel data that together support a normative theory of multi-product choice and consumption. The experiments show that sets of hedonic products are chosen in a way that maximizes total expected future utility, given uncertainty about future preferences. The chosen set can therefore be interpreted as a utility-maximizing hedge. To maximize expected utility, products in the set must be consumed strategically. We find evidence of the strategic consumption pattern predicted by the theory, with consumers effectively preserving variety for future consumption occasions. A boundary condition for strategic consumption is also identified. Finally, analysis of yogurt purchases made by single-person households in two geographic markets shows that shoppers with lower consumption rates choose relatively more variety than shoppers with higher consumption rates do. Taken together, these results provide support for the normative theory of multi-product shopping and consumption.

1 Introduction

People generally shop at grocery stores in order to purchase products for future consumption. They buy enough to justify the “fixed cost” of shopping (cf. Bell, Ho, and Tang, 1998), although perishability or other holding costs may limit purchase quantities. Products purchased in store are held in inventory in the shopper’s pantry, refrigerator, or freezer until they are consumed. Each consumption choice reduces overall inventory and the quantity of the chosen alternative. As a result, consuming hedonic products involves a dynamic sequence of interrelated choices that ultimately must be reflected in the set chosen while shopping.

In this paper, we test a normative theory of shopping and consumption proposed by Fox, Norman, and Semple (2017), henceforth called *FNS*. They introduced and analyzed two models, both comprising two stages: (i) the shopping stage, in which a set of n products is chosen for future consumption; and (ii) the consumption stage, in which products are consumed one-at-a-time from the previously chosen set. Importantly, the sequence in which products are chosen in the consumption stage is not fixed—consumers are free to choose whichever product maximizes their current and future (expected) utility. The canonical n-pack model [*CNPM*] assumes that, in the consumption stage, a product is chosen on each of n consecutive consumption occasions. The generalized n-pack model [*GNPM*] incorporates an outside option so that, on each consumption occasion, the consumer may choose a product or choose not to consume. This model represents the situation in which different categories of products compete for the consumer’s attention, and is appropriate to analyze consumers with different usage rates in the category of interest. In fact, previous analysis of this model showed that consumption rate may affect the set of products chosen in the shopping stage.

For over 25 years, research in this area has focused primarily on “diversification bias” (Simonson, 1990; Read and Loewenstein, 1995), the finding that consumers who select a set of products for future consumption generally include more variety than if they had chosen from a full assortment at the time of consumption. The difference in variety has been attributed to consumers’ inability to accurately forecast what they will want and when they will want it. Kahneman and Snell (1990) called it "a mistake, which they could perhaps avoid by a serious attempt to predict their tastes on each of these weeks [*consumption occasions*] separately" (p. 304). In contrast, *FNS* assumed that consumers recognize their inability to forecast future preferences and so resort to hedging. The best set is therefore the one that maximizes expected utility over the consumption horizon; that is, the optimal hedge against future preference uncertainty. To illustrate how the two theories lead to different outcomes, consider a consumer selecting a set of three products from an assortment of three product alternatives, $\{A,B,C\}$. Assume that the consumer makes a multinomial choice but has the same intrinsic utility for each alternative. If this consumer were to choose a product on three consecutive consumption occasions, each from the full assortment of alternatives, those choices would most likely include just two different alternatives. This is because, of the 27 ($= 3^3$) equally probable ordered consumption sequences, 18 include exactly two alternatives (e.g., $[A,B,A]$, $[B,C,C]$). If the consumer could forecast their three future consumption preferences with perfect accuracy, that set would therefore be limited to two product alternatives with probability 2/3 (= 18/27). However, if the same consumer were hedging against future preference uncertainty, the set providing the best hedge is the one that includes a single unit of each alternative. This intuitive choice will be confirmed by the dynamic models presented in §3.

To evaluate the notion that consumers hedge against future preference uncertainty, we test three propositions based on the *FNS* models. The first involves the value function of the *CNPM*. This is a value function in closed-form, unique to each consumer, that may be used to determine the consumer's valuation of every possible set. If consumers are hedging, they should select sets with higher valuations. Our first and second laboratory experiments provided data for this test. The second proposition focuses on consumption choices. The optimal consumption policy of the *CNPM* is a dynamic decision rule that considers the logarithm of each alternative's current inventory along with the alternative's current utility when selecting a product for consumption (a proof of the precise rule is provided in the

Appendix). This state-based rule allows for many different consumption sequences and should be contrasted with the forecast-based (fixed) sequences common in prior research. Both of our laboratory experiments provided data to test this dynamic decision rule, although the second experiment proved more diagnostic. The third proposition concerns whether consumers with lower consumption rates choose sets with relatively more variety (i.e., more product alternatives). This arguably counter-intuitive proposition is based on analysis of the *GNPM*, where greater variety in the chosen set was shown to be necessary to compete with an attractive outside option (implied by a lower consumption rate). We elaborate on the underlying intuition in §3.2. Two years of multi-market panel data for yogurt purchases were used to test this proposition.

2 Literature Review

Three literature streams have addressed consumers' choices of multiple substitutable products. Research in consumer psychology has compared the product variety of a set chosen for future consumption (termed "simultaneous choice") with the variety of products chosen on successive consumption occasions (termed "sequential choices"). Observe that simultaneous choice constrains subsequent consumption to the product alternatives remaining in the chosen set as it is depleted. Sequential choices are not so constrained; each consumption choice is made from the full assortment of product alternatives. The primary finding in this literature is that simultaneous choice of a given set generally includes more varied product alternatives than the same number of sequential choices (e.g., Simonson, 1990; Simonson and Winer, 1992; Read and Loewenstein, 1995). Researchers argued that this "diversification bias" (cf. Read and Loewenstein 1995) was due to poor forecasting, 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 the 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. 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. Read and Loewenstein (1995) 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). Though Read and Loewenstein (1995) did not find diversification to be be a rational hedge against uncertainty, others did. Notably, Salisbury and Feinberg (2008) used simulation studies to show that the extent of diversification should depend on the level of future preference uncertainty, along with the relative attractiveness of product alternatives and uncertainty about their attractiveness.

Three recent econometric studies focused on 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. 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.

Walsh (1995) proposed and analyzed the first dynamic model of shopping and spending. 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 units to a set added more to 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.

3 Normative Theory of Shopping and Consumption

The normative theory of shopping and consumption that we test is based on the two models analyzed in *FNS*: (i) the *CNPM*, which defines consumption occasions based on actual consumption of an item from the chosen set, and (ii) the *GNPM*, which includes an outside option to allow for variation in usage rates. We borrow their 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 as a set can be represented by the integer vector $(k_1, k_2, ..., k_M)$ with $k_i \geq 0$ and $\sum_{i=1}^{M} k_i = n$. Note that multiple units of a product alternative may be selected, as well as zero units. We omit consumer and category subscripts for expositional clarity.

The normative theory of shopping and consumption 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), and 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 uncertainty about the consumer's preference for the product; 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 Canonical *n*-Pack Model

The canonical n-pack model [$CNPM$] assumes that a product from the choice set is selected on each consumption occasion. To illustrate the dynamics, let us 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 $U_1+\epsilon_{11}$ versus $U_2+\epsilon_{21}$ 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 strategic 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. Clearly, factors such as trip type (major versus fill-in, cf. Kollat and Willet, 1967) and multiple-purchase incentives 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 mathematical analysis, but would also obscure the basic insights on which our propositions are based.

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

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

Our proof is in the Appendix. 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 pre-commit 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. One can think of this component as capturing the *intrinsic utility* of the set. However, if the consumer is free to choose an alternative at the time of consumption in the presence of future preference uncertainty, then the logarithmic second component in (3.1) must also be present. This component captures the incremental value of making consumption choices from a set using *all available information at the time of consumption*. Observe that the logarithmic component does not depend on the U_i but rather 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. The two components in (3.1) capture the tension between opposing objectives: one favoring intrinsic utility, the other favoring variety. A consumer's utility-maximizing set—the optimal hedge—must balance these two objectives.

Notes: (i) 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). The same property was demonstrated by Walsh (1995) for the two-alternative case. (ii) 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. Like the logarithmic component, 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, 2017).

Our first proposition states that consumers will evaluate sets for purchase consistent with (3.1).

Proposition 1. *Consumers' choice of a set of substitutable products is consistent with the valuation function (3.1).*

The optimal consumption strategy is instructive as well. Suppose, on the t^{th} consumption occasion, the set currently has q_{it} units of alternative i remaining. Then the utility maximizing strategy is to select the alternative that maximizes $ln\,(q_{it}) + \epsilon_{it}$. Observe that the optimal consumption strategy incorporates each product's current inventory, $ln\,(q_{it})$, and random component, ϵ_{it}, but not its deterministic component, U_i. The result is that consumption choices are biased toward alternatives with higher inventory levels, implicitly preserving variety throughout the consumption sequence. This is the basis for our second proposition.

Proposition 2. (The Inventory Effect) *Let* q_{it} *represent the quantity of alternative* i *available on the* tth *consumption occasion. Then the consumer's decision rule for selecting alternative* i *is a linear function of* $ln\left(q_{it}\right)$.

The decision rule in Proposition 2 additionally implies that the probability of selecting an alternative is increasing in its inventory. It can be shown that this probability property generalizes to any error distribution (see Alptekinoglu and Semple 2017).

3.2 Generalized n-Pack Model

The generalized n-pack model $[GNPM]$ incorporates an outside option, 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 (of 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

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

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 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

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

This is an unwieldy value function, but *FNS* showed that the set that optimizes (3.2) has at least as much variety as the set optimizing (3.1). Observe also that equation (3.2) nests (3.1), fixing the value of the outside option U_0 =-∞.

FNS noted that lower consumption rates promote greater variety in the set that optimizes (3.2). A simple thought experiment provides the intuition. Imagine a consumer in the final period of their consumption horizon with r products ($r \leq n$) remaining in inventory. Multiple products may be available in the final period (and some ultimately wasted) if an attractive outside option causes a low consumption rate.[1] The possibility of waste is explicitly permitted in (3.2) because utility maximization is about matching the right product to the right occasion, not necessarily consuming every product in the set. Assuming that r products were available in the final period, utility could be maximized only if r different product alternatives were represented—maximum variety. This "end of horizon" effect, as it is known in dynamic programming, would encourage consumers with low consumption rates to include greater variety in their chosen sets. We would therefore expect consumers with low consumption rates to choose sets with more variety than sets of the same size chosen by consumers with high consumption rates, and possibly enduring greater waste. This is the focus of Proposition 3.

Proposition 3. *For a given n, the variety included in a consumer's optimal choice set is decreasing in consumption rate.*

4 Experimental Evidence

We conducted two experiments to evaluate Propositions 1 and 2.

[1]Once a set is purchased, costs are sunk, so utility maximization is an appropriate objective.

4.1 Experiment 1

Experiment 1 was conducted in two phases. In the initial phase, subjects first selected one of two available snack categories—salty snacks or candies. They were then given a list of 15 single-serve product alternatives from the selected category available in local vending machines. Subjects identified their three most preferred alternatives from the list, then assigned long-run choice probabilities to those three alternatives.[2] These long-run choice probabilities enabled us to determine subjects' intrinsic utilities for their favorite alternatives ($U_i = ln\,(p_i/p_1)$, where alternatives have been ordered by long-run choice probability and p_i is the respondent's long-run choice probability for alternative i). Subjects also provided minimal demographic information. Next, subjects chose a set of three products for future consumption. Subjects were instructed that they could choose one or more units of any of the 15 product alternatives in the category, so long as they chose a total of three. This effectively replicated the "simultaneous choice for sequential consumption" introduced by Simonson (1990). Unlike some studies of diversification bias, however, our subjects chose a set of products without forecasting (and pre-committing to) a consumption order.

In the second phase of the experiment, subjects received three products from their selected snack category, though not necessarily the set of products they had chosen in the first phase. The set was composed of two units of one product alternative and one unit of another. Based on a 2 x 2 between subjects design, subjects received:

- either their most preferred (i.e., favorite) product alternative together with their second most preferred alternative, or their most preferred (i.e., favorite) alternative together with their third most preferred alternative;

- either one or two units of their most preferred (i.e., favorite) product alternative and the remaining unit(s) of their less preferred alternative.

The first design factor manipulated the difference in preference between product alternatives

[2]Long-run consumption choice probabilities were elicited using a scenario of twice-weekly vending machine purchases over the course of a year.

in the set. All things being equal, subjects would have a higher relative preference for their favorite vis-a-vis their third most preferred alternative as compared to their second most preferred alternative. The second design factor manipulated which alternative (the more or less preferred) had a higher beginning inventory. To complete phase two of the experiment, subjects consumed the three products within five days, consuming no more than one product per day.

Proposition 1 posits a utility-based normative valuation for any set of n products. We tested this valuation using subjects' choice of a set of three products for future consumption from the first phase of Experiment 1. This choice was made from among the ten possible sets of subjects' three most preferred product alternatives.[3] Ordering alternatives by preference $a = 1, 2, 3$ with quantities of ordered alternatives (k_a) in parentheses, the possible sets are (3,0,0), (0,3,0), (0,0,3), (2,1,0), (2,0,1), (1,2,0), (0,2,1), (1,0,2), (0,1,2), and (1,1,1). We substituted subjects' intrinsic utilities U_a into (3.1) to compute their valuations V_i $i = 1, 2, \ldots, 10$ for the possible sets. These valuations V_i were used in a multinomial logit [*MNL*] choice model to determine choice probabilities for each set. For set j with valuation V_j and intercept α_j, the choice probability was specified to be $exp(\alpha_j + \beta V_j)/\sum_{i=1}^{10} exp(\alpha_i + \beta V_i)$. The intercept for (1,1,1) was set to zero for identification. The sets chosen by 81 of the 107 subjects comprised only their three most preferred product alternatives. For the other 26 subjects, we were unable to compute a valuation for the chosen set because we did not have a long-run choice probability for at least one chosen alternative. We conjecture that some subjects might have used the experiment as an opportunity to try unfamiliar products, a rationale inconsistent with the normative theory. Others may simply have included a product alternative with the same long-run choice probability as one of their three most preferred—essentially a tie. Though we were unable to evaluate choice sets for these 26 subjects, their consumption choices were included in the next section's analysis.

Table 1 shows the fit statistics for the "Normative Valuation" *MNL* choice model. The

[3] *FNS* noted that a set of n products chosen for future consumption should include at most the n most preferred product alternatives.

hit rate was 49.4%; the log-likelihood of observed choices was -100.95. For comparison, we estimated an "Intercept Only" model. The hit rate of this model was 39.5%; the log-likelihood of observed choices was -112.15. The "Normative Valuation" model was superior to the "Intercept Only" model based on both *AIC* and *BIC*, was well as the likelihood ratio test (D=22.40, p-value<0.0001). The hit rate of the "Normative Valuation" model was also materially better than the "Intercept Only" model, increasing the hit rate by 9.9% to nearly 50%. This analysis supports Proposition 1.

Table 1
Choice Set Models

Experiment	Model	Obs	Hit Rate	Log-Likelihood	AIC	BIC	Likelihood Ratio (D)	LR Test *p*-value
1	Normative Valuation	81	49.4%	-100.95	217.89	237.05	22.40	<0.0001
	Intercept Only	81	39.5%	-112.15	238.29	255.05		
2	Normative Valuation	58	48.3%	-71.78	157.55	171.97	15.71	<0.0001
	Intercept Only	58	41.4%	-79.63	171.26	183.62		
Pooled	Normative Valuation	139	50.4%	-176.29	370.57	396.98	36.93	<0.0001
	Intercept Only	139	40.3%	-194.75	405.50	428.98		

In the second phase of the experiment, subjects received two units of one product alternative and a single unit of another alternative. All 101 subjects were able to choose between the alternatives on the first consumption occasion, but only 62 were still able to choose between the alternatives on the second consumption occasion (there was no choice flexibility on the final consumption occasion). The "Strategic" model reflects Proposition 2, where a product alternative's consumption is based on the logarithm of its inventory. Ordering alternatives by preference ($a = 1, 2$) and supressing the purchase occasion subscript, the choice to consume alternative a was specified as

$$exp\left(\alpha_a + \beta ln\left(k_a\right)\right) / \left(exp\left(\alpha_1 + \beta ln\left(k_1\right)\right) + exp\left(\alpha_2 + \beta ln\left(k_2\right)\right)\right)$$

the intercept for alternative 1 was fixed at zero for identification. Fit statistics are shown in the upper panel of Table 2. The hit rate was 65.9% and the log-likelihood -101.89. For comparison, we estimated a "Myopic" model in which product alternatives are chosen based on the intrinsic preference (U_a replaces $ln\left(k_a\right)$ in the MNL choice specification) and an "Intercept Only" model.[4] For the preference-based "Myopic" model, the hit rate was 65.9% and the log-likelihood was -102.82; for the "Intercept Only" model, the hit rate was 65.9% and the log-likelihood was -105.29.[5] Given observed consumption choices, the "Strategic" model was more likely than either alternative model, although only $e^{0.93} = 2.53$ times more likely than the "Myopic" model. In addition, the "Strategic" model was superior to the "Intercept Only" model on the basis of both AIC and BIC, as well as the likelihood ratio test (D =6.79, p-value=0.0091). On the other hand, the hit rates for all three models were the same, indicating that the difference in fit was not material. Our analysis therefore provides limited support for Proposition 2.

Table 2
Consumption Choice Models

				Strategic		Myopic		Intercept Only	
Experiment	Consumption Occasion(s)	Inventory Ratio(s)	Obs	Hit Rate	Log-Likelihood	Hit Rate	Log-Likelihood	Hit Rate	Log-Likelihood
1	All	Multiple	164	65.9%	-101.89	65.9%	-102.82	65.9%	-105.29
2	All	Multiple	157	66.4%	-89.47	61.0%	-97.66	61.0%	-98.50
	1	4 to 1	61	75.4%	-32.56	62.3%	-40.64	62.3%	-40.90
	2	3 to 1	45	66.7%	-30.57	51.1%	-31.82	51.1%	-32.62
	3 and 4	2 to 1, 1 to 1	40	52.5%	-26.34	70.0%	-25.20	70.0%	-24.98

4.2 Experiment 2

The tepid support for Proposition 2 prompted us to consider why the consumption choices in Experiment 1 were not more consistent with the normative theory of shopping and consumption. We conjectured that the inventories of the two product alternatives might not have been sufficiently different to evoke strategic consumption. In fact, all consumption choices made in Experiment 1 involved inventory ratios of either 2 to 1 or 1 to 1, hardly a robust test of an inventory-based decision rule. This conjecture led us to conduct a second experiment with larger differences in inventory between product alternatives in the consumption phase.

[4] Similar "Myopic" models are commonly specified for consumption (e.g., Guo, 2010).

[5] The "Myopic" and "Intercept Only" models have the same hit rate because the estimated $\alpha_2 < \alpha_1 = 0$ and $U_2 \leq U_1$.

The first phase of Experiment 2 replicated Experiment 1 exactly, with subjects again choosing a set of $n = 3$ snack products. In the second phase of Experiment 2, subjects received a larger beginning inventory of $n = 5$ products; either: (i) four units of the subject's most preferred (i.e., favorite) product alternative and one unit of the second most preferred alternative, or (ii) one unit of the subject's most preferred (i.e., favorite) alternative and four units of the second most preferred alternative. The set of five products was then consumed over five consecutive weekly class meetings. Participants chose one product per class meeting. This longer consumption sequence enabled us to observe consumption choices with more varied inventory ratios (4 to 1, 3 to 1, 2 to 1 and 1 to 1) compared to Experiment 1, providing additional information about how the inventories of product alternatives affect strategic consumption. Recall that Experiment 1's 2x2 design distinguished between respondents' second and third most preferred product alternatives. Because we had found no significant differences in consumption choices between the second and third most preferred alternatives, Experiment 2 omitted this manipulation.

Sixty one graduate students participated in Experiment 2, completing the first phase and making at least one consumption choice during the second phase. Three subjects included a product alternative in their chosen set not among their three most preferred (possible reasons were discussed in §4.1). We were able to compute a valuation for the chosen sets of the other 58 subjects, so they were included in our analysis. Fit statistics for the "Normative Valuation" and "Intercept Only" models are reported in the second panel of Table 1. For the "Normative Valuation" model, the hit rate was 48.3% and the log-likelihood of observed choices was -71.78. For the "Intercept Only" model, the hit rate was 41.4% and the log-likelihood was -79.63. As in Experiment 1, the "Normative Valuation" model was superior to the "Intercept Only" model based on both AIC and BIC, as well as the likelihood ratio test (D=15.71, p-value<0.0001). And, while the difference in hit rates between the "Normative Valuation" and "Intercept Only" models was not as large as in Experiment 1, the 6.9% difference was nevertheless substantial. We therefore conclude that Experiment 2 provides additional support for Proposition 1.

Now we turn our attention to Experiment 2's consumption choices. Four subjects failed to complete all consumption choices; fortunately, half of the missing observations were of the minimally diagnostic final consumption choice (i.e., with a 1 to 1 inventory ratio). All 61 subjects chose between two product alternatives on the first consumption occasion (inventory ratio 4 to 1), 45 chose between two alternatives on the second consumption occasion (inventory ratio 3 to 1), 29 chose between two alternatives on the third consumption occasion (inventory ratio 2 to 1), and 11 chose between two alternatives on the fourth consumption occasion (inventory ratio 1 to 1) for a total of 146 consumption choice observations. In the second panel of Table 2, we see that the "Strategic" model's hit rate was 66.4% across all observations and the log-likelihood of observed choices was -89.47. The "Myopic" model's hit rate was 61.0% across all observations and the log-likelihood of observed choices was -97.66, while the "Intercept Only" model's hit rate was also 61.0% and the log-likelihood of observed choices was -98.50. As in Experiment 1, we observe that the "Strategic" model was more likely than either alternative model. In addition, the "Strategic" model was superior to the "Intercept Only" model based on both *AIC* and *BIC,* as well as the likelihood ratio test (D =18.05, p-value<0.0001). Finally, the hit rate of the "Strategic" model was nearly 2/3, 5.4% better than both alternative models.

We also partitioned the data by consumption occasion, which revealed an interesting pattern. On the first consumption occasion, all 61 subjects had a 4 to 1 inventory ratio. For this occasion, the "Strategic" model's hit rate was 75.4% and log-likelihood was -32.56, as compared with the "Myopic" model's 62.3% hit rate and -40.64 log-likelihood or the "Intercept Only" model's 62.3% hit rate and -40.90 log-likelihood. The large inventory disparity resulted in a far superior fit for the "Strategic" model. On the second consumption occasion, the 45 remaining subjects had a 3 to 1 inventory ratio. For this occasion, the "Strategic" model's hit rate was 66.7% and log-likelihood was -30.57, as compared to the "Myopic" model's 51.3% hit rate and -31.82 log-likelihood or the "Intercept Only" model's

51.1% hit rate and -32.62 log-likelihood. The 3 to 1 inventory disparity, while not as large as on the first consumption occasion, was nevertheless enough to result in a better hit rate and marginally better likelihood for the "Strategic" model. Pooling the third and fourth consumption occasions, subjects faced either a 2 to 1 or 1 to 1 inventory ratio, as in Experiment 1. For these final 40 consumption observations, the "Strategic" model's hit rate was only 52.5% and log-likelihood was -26.34, as compared to the "Myopic" model's 70.0% hit rate and -25.20 log-likelihood or the "Intercept Only" model's 70.0% hit rate and -24.98 log-likelihood.[6] For those consumption occasions on which subjects faced inventory ratios of 2 to 1 or 1 to 1, the "Strategic" model was marginally less likely than the two alternative models and had a lower hit rate. Focusing on the "Strategic" model's hit rate across consumption occasions, we see that it decreased from 75.4% when the inventory ratio was highest to 52.5% when the ratios were lowest. In summary, our analyses of Experiment 2's consumption choices show that, for sufficiently unbalanced inventories, the data were consistent with strategic consumption. Evaluating these results together with those from Experiment 1 suggests a boundary condition for strategic consumption—if the disparity in inventory between product alternatives is sufficiently large, consumption will be strategic; if the disparity in inventory between product alternatives is small, consumption may not be strategic.

4.3 Pooled Results

Because the first phase of Experiments 1 and 2 was exactly the same, we were able to pool the data from the two experiments. The bottom panel of Table 1 shows that the "Normative Valuation" model was superior to the "Intercept Only" model based on both *AIC* and *BIC*, as well as the likelihood ratio test (D =36.93, p-value < 0.0001). In addition, the "Normative Valuation" model's 50.4% hit rate was significantly higher than the "Intercept

[6]Note that both the "Intercept Only" model and the "Myopic" model (which nests the "Intercept Only" model) were estimated using the full dataset. Partititioning the likelihood *ex post* made it possible for the "Intercept Only" model to fit a subset of choice observations better than the more general model.

Only" model's 40.3% hit rate (p-value = 0.0458 for a test of two proportions). Finally, we ranked the ten possible sets of participants' three most preferred product alternatives based on their valuations (3.1). Figure 1 shows the choice frequency of sets ranked from 1 (highest valuation) to 10 (lowest valuation). The pattern in the figure clearly supports Proposition 1.

Figure1
Choice Sets Ranked by Valuation

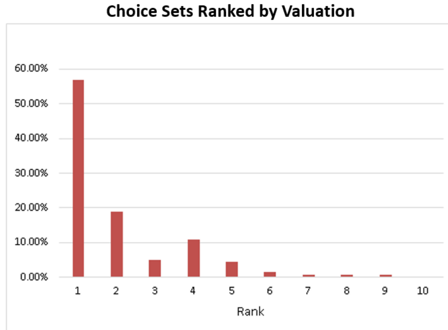

5 Panel Data Evidence

Proposition 3 states that the variety included in a consumer's optimal set of products is decreasing in consumption rate. Rather than testing this proposition by attempting to manipulate consumption rate, we chose to use observational data. Following Simonson and Winer (1992), we selected single-serve yogurt products.

Data from two major metropolitan markets (one in the midwestern US totalling 1707 households, the other in the southeastern US totalling 1031 households) were used for this analysis. The panel data captured purchases made during a 24-month period between September 2004 and September 2006. The dataset was partitioned so that the first 18 months served as a calibration period, while the final six months were used for estimation. In order to avoid intra-household preference heterogeneity, the dataset was limited to single-member households (308 and 137 households, respectively, in the two markets). We also required that households had made at least 10 yogurt purchases during the calibration period, then purchased again during the estimation period. The final dataset included

70 single-member households purchasing a total of 8,670 single-serve yogurt cups on 1,611 shopping trips during the calibration period, and purchased 2,376 cups on 443 shopping trips during the estimation period. Though this dataset is relatively small, it was purposefully constructed to avoid intra-household heterogeneity so as to provide a fair test of Proposition 3.[7]

Data from the calibration period were used for two purposes. The first was to determine panelists' long-run consumption preferences. These preferences were determined by Universal Product Code [*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 by *UPC*. 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 by averaging over the calibration period (assuming that all yogurt purchases were consumed). Initially, we conjectured that each day (morning) presented a consumption opportunity. Interestingly, we found that one panelist consumed 1.328 units/day, buying yogurt on 114 shopping trips during the calibration period (recall that each unit is a single serving). However, the other panelists consumed fewer than one unit/day.

Data from the estimation period were used to assess the relative variety of yogurt purchases. Consistent with the normative theory of shopping and consumption, we assumed that n, the number of units chosen on a given trip, was exogenous.[8] Variety was measured as the number of product alternatives m in the chosen set. Clearly, m depends on the set size n (e.g., $m \leq n$). To control for this dependency, we took advantage of the fact that the *CNPM* (no outside option) is actually a limiting case of the *GNPM* (with utility of the

[7]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.

[8]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.

outside option set to -∞) for a given set size n. To evaluate the observed m, we therefore compared it to the variety of the *CNPM*'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} for every panelist and every set size n, which we then used to determine the relative variety of observed puchases. The relative variety measure for yogurt purchases is the proportional difference between observed and optimal variety, $D = \frac{m-m^{opt}}{m^{opt}}$.[9]

Proposition 3 states that the variety included in a consumer's optimal choice set is decreasing in consumption rate, but does not specify a functional form for the relationship. We therefore began by estimating nonparametric correlations—Spearman's Rank Correlation and Kendall's Tau—as well as Pearson's R for comparison. An important characteristic of the data is that relative variety varied across a panelist's yogurt purchases, but the panelist's consumption rate did not. We will 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 average over trips, $\overline{D}_h$, was used to compute correlations at the household-level. Table 3 shows that all correlations had the expected sign: negative for $Units/Day_h$ and positive for $Days/Unit_h$. All nonparametric correlations were significant at the 0.05 level, and were 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 with proportional difference in variety was more linear for the former than the later.

[9]By computing $\frac{m-m^{opt}}{m^{opt}}$ rather than $m - m^{opt}$, we were able to avoid scaling issues.

Table 3
Correlations of Consumption Rate and Relative Variety

Correlation	$Units/Day_h$ v D_t	$Days/Unit_h$ v D_{ht}	$Units/Day_h$ v $\overline{D}_t$	$Days/Unit_h$ v $\overline{D}_t$
Spearman's Rank	-0.1707 (0.0003)	0.1707 (0.0003)	-0.3303 (0.0052)	0.3303 (0.0052)
Kendall's Tau	-0.1249 (0.0003)	0.1249 (0.0003)	-0.2107 (0.0116)	0.2107 (0.0116)
Pearson's R	-0.1282 (0.0069)	0.2239 (<.0001)	-0.1625 (0.1788)	0.2632 (0.0277)

The combination of a single household-level predictor ($Units/Day_h$ or $Days/Unit_h$) and a household-level response variable with repeated measures over time (D_{ht}) led us to apply a random coefficients model. The first-level was specified

$$D_{ht} = \theta_h + \xi_{ht} \tag{5.1}$$

and the hierarchical random intercept was specified

$$\theta_h = \delta + \gamma \cdot f\left(Units/Day_h\right) + \zeta_h, \tag{5.2}$$

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 relative variety of a purchase, D_{ht}, and the transformed consumption variable, $f\left(Units/Day_h\right)$. 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 transformation applied, the 25000 Markov Chain Monte Carlo iterations converged quickly after a short burn-in period and autocorrelation proved acceptable, resulting in a stable posterior distribution for γ.

Table 4 shows estimates for the "Random Coefficients" model specified by (5.1) and (5.2) for selected monotonic transformations of the independent variable. The table also shows estimates for two nested models models: (i) "Fixed Intercept" (i.e., $\beta_1 = \beta_2 = \ldots = \beta_H = \beta$), and (ii) "Random Coefficients($\gamma = 0$)." To compare models, we used the Deviance Information Criterion [DIC], a Bayesian analog of AIC.[10] To assess model predictions, we bootstrapped the estimation dataset, then compared actual values of D_{ht} with model predictions 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 fit the data better: $exp\,(Units/Day_h)$ and $ln\,(Days/Unit_h)$.

Table 4
Relative Variety Models for Yogurt Purchases

			Posterior Predictions		Posterior Distribution of γ		
Model	**Predictor**	***DIC***	***MAD***	***MSE***	**Mean**	**Pr(γ< 0)**	**(5% , 95%)**
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	-0.211	0.909	(-.4740,-.0470)
Random Coefficient	$\text{exp(Units/Day}_h)$	244.80	0.270	0.136	-0.100	0.870	(-.2482,.0511)
Random Coefficient	Days/Unit_h	245.55	0.258	0.127	0.015	0.016	(.0037,.0255)
Random Coefficient	$\text{ln(Days/Unit}_h)$	244.73	0.261	0.130	0.092	0.024	(.0144,.1673)

The two nested models, "Fixed Intercept" and "Random Coefficient ($\gamma = 0$)," differed greatly from one another in terms of fit—"Random Coefficient ($\gamma = 0$)" had a far lower DIC than "Fixed Intercept" (lower is better)—but were very similar in predictive accuracy. In terms of fit, the "Random Coefficient ($\gamma = 0$)" model had only a slightly higher DIC than three of the four full "Random Coefficient" models, and actually had a slightly lower DIC than the fourth. In terms of predictive accuracy, the "Random Coefficient ($\gamma = 0$)" model offered better posterior predictions than two of the four full "Random Coefficient" models, but worse posterior predictions than the other two. We conclude that unmodeled individual differences explained much more variation in D_{ht} than consumption rate did. On the other hand, the functional form of the relationship between consumption rate and variety mattered, particularly for predictive accuracy. The "Random Coefficient" model using $ln\,(Days/Unit_h)$ as the predictor had the lowest DIC, and so was preferred to the

[10] DIC is appropriate for comparison of both nested and non-nested models; see Spiegelhalter, et al. (2002) for details.

other "Random Coefficient" models based on fit. Further, this model and the "Random Coefficient" model using $Days/Unit_h$ as the predictor offered 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 in Proposition 3 is better 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 was found to be positive and significant. The mean estimate was 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 was also in the expected direction—positive for $f\left(Days/Unit_h\right)$ and negative for $f\left(Units/Day_h\right)$. The posterior estimates of γ were significant at 0.05 level only for models using $f\left(Days/Unit_h\right)$ as the predictor. Taken together, the random coefficient models and nonparametric correlations provide support for Proposition 3.

6 Concluding Remarks

This paper develops and tests three propositions from a normative theory of shopping and consumption. The first proposition links consumers' choice of a set of substitutable hedonic products to a rational valuation of the expected future utility of that set. This valuation balances the intrinsic utility of products in the set with the additional utility gained by constructing the set to allow choice on future consumption occasions, effectively hedging against future preference uncertainty. Evidence from two experiments offered support for this proposition. The second proposition predicts strategic consumption choices; specifically, that consumption choices will be linearly dependent on the natural log of inventory for available product alternatives. Evidence from the same two experiments offered qualified support for this proposition. If the inventories of available alternatives differed sufficiently, then consumption choices were well predicted by a linear model of the natural log of inventory. However, if the inventories of available alternatives were not sufficiently different, consumption choices were predicted just as well by consumers' long-run utilities as by the natural log of their inventories. Our conjecture is that strategic consumption, while normative, is only activated if there is a salient contrast between the inventories of available product alternatives. The third proposition relates the optimal variety of products in the chosen set with a consumer's usage, or consumption rate. We tested this proposition using multi-market panel data of yogurt purchases for single-person households. Though the proposition specifies no functional form for the relationship, our analysis showed that the yogurt purchases of consumers with lower usage rates included relatively more variety than the purchases of consumers with higher usage rates. Thus, our analysis supported this arguably counterintuitive proposition.

Recall that the theory tested in this paper represents an alternative to existing theories about the choice of multiple substitutable products for future consumption. Those theories explain observed behavior based on bias (i.e., *diversification bias*) or systematic variation in preferences over time (i.e., *variety-seeking*). In contrast, we presented evidence for a theory of rational hedging, in which consumers construct and preserve choices for future consumption occasions that maximize expected future utility. We would not presume to suggest that competing theories are flawed, or even that the theory we tested explains more observed variation in multi-product choice than other theories do. Rather, we would suggest that the theory we tested provides a rational baseline against which to evaluate competing theories.

Our investigation has the potential to generate a number of future research opportunities. Testing the theory of rational hedging, together with variety-seeking and diversification bias, is likely to find conditions under which different theories apply. For example, we found that smaller vs. larger differences in a consumer's inventories of available product alternatives led to different decision-rules for consumption. Exploring the boundary conditions of different choice models and decision rules could also be a fruitful avenue for future research. Our research also opens the door for future investigation of the effects of consumption rate on consumer's product choices. Studies of consumption have heretofore been limited, perhaps because consumption data are difficult to obtain. Yet consumption expectations clearly motivate consumer's product choice and category incidence decisions. This appears to be another potentially fruitful avenue for future research.

References

[1] Aydin Alptekinoglu and John H. Semple. Consuming a set of substitutable products using a general random utility framework. 2017.

[2] David R. Bell, Teck Hua Ho, and Christopher S. Tang. Determining where to shop: Fixed and variable costs of shopping. *Journal of Marketing Research*, 35(August):352–69, 1998.

[3] Moshe Ben-Akiva and Steven R. Lerman. *Discrete Choice Analysis*. The MIT Press, Cambridge, MA, 1985.

[4] Jean-Pierre Dube. Multiple discreteness and product differentiation: Demand for carbonated soft drinks. *Marketing Science*, 23(1):66–81, 2004.

[5] Edward J. Fox, Laura E. Norman, and John H. Semple. Choosing an n-pack of substitutable products. *Management Science*, 2017.

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

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

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

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

[10] Sanghak Lee and Greg M. Allenby. Modeling indivisible demand. *Marketing Science*, 33(3):364–81, 2014.

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

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

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

[14] Linda Court Salisbury and Fred M. Feinberg. Future preference uncertainty and diversification: The role of temporal stochastic inflation. *Journal of Marketing Research*, 35(2):349–59, 2008.

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

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

[17] 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

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

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

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

(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 $ln\left(k_i\right)+\epsilon_i$, is readily apparent:

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

□