The Exponomial Choice Model: A New Alternative for Assortment and Price Optimization

Aydin Alptekinoglu, John H. Semple

The Exponomial Choice Model: A New Alternative for Assortment and Price Optimization

February 20, 2015

Abstract

We propose a new discrete choice model for use in optimal pricing and assortment planning. In contrast to multinomial and nested logit (the prevailing choice models used for optimizing prices and assortments), we assume the distribution of consumer utilities is negatively skewed, an assumption motivated by conceptual arguments as well as published data. The choice probabilities in our model can be derived in closed-form as an *exponomial* (a linear function of exponential terms). The loglikelihood function is concave in model parameters, which makes empirical estimation straightforward. Our pricing and assortment planning insights differ from the literature in two important ways. First, our model allows variable markups in optimal prices that increase with expected utilities. Second, when prices are exogenous, the optimal assortment allows leapfrogging in prices, i.e., a product can be skipped in favor of a lower-priced one depending on the utility positions of neighboring products. These two plausible pricing and assortment patterns are ruled out by multinomial logit (and by nested logit within each nest). We provide structural results on optimal pricing for monopoly and oligopoly cases, and on the optimal assortments for both exogenous and endogenous prices. Finally, we show that our model generalizes to allow different variances for each choice's stochastic error term without sacrificing closed-form choice probabilities, something that is not feasible in the logit framework.

Keywords: Discrete choice theory, assortment planning, pricing, revenue management, multinomial logit, nested logit.

1. Introduction

Recent academic research in the area of optimal pricing and assortment planning has caught the attention of companies that provide decision support systems for retailers (Kök et al. 2009; Sinha et al. 2013). As a case in point, the authors were approached by a software development company building business analytics solutions for several retail chains selling packaged goods. Their clients were seeking more intelligent tools that could help store managers select the right assortments and/or prices as well as predict the associated revenue effects. Because the model needed to be applied to 1000's of retail stores with 100's of different assortments, it needed to be scalable, customizable, and tractable (in an optimization sense). The software company was already investigating several published pricing and assortment planning models.

To date, the prevailing choice models used for optimizing prices and assortments—multinomial logit (MNL) and nested logit (NL)—assume the distribution of consumers' willingness to pay (WtP) for a product is positively skewed (e.g., van Ryzin and Mahajan 1999; Talluri and van Ryzin 2004; Rusmevichientong et al. 2010; Li and Huh 2011; Davis et al. 2012; Alptekinoğlu and Grasas 2014). This stems from the error term in logit models having a Gumbel distribution, and it makes sense in many cases, particularly those where consumers have limited information regarding a product's value (e.g., wine, art, etc.). However, in situations where consumers are well informed about products and their values, one might expect the distribution of consumers' WtP to be negatively skewed simply because even an enthusiastic buyer would be put off by the thought of overpaying.

To make this concrete, imagine a person buying a new car. Most individuals know the MSRP and many more consult resources such as Kelley Blue Book or Edmunds.com to see what others are paying or to obtain the dealer's cost. A consumer who has been informed in this manner thus obtains a benchmark price for each car in their choice set. In most cases, the probability that a consumer is willing to pay 10-20% more than this benchmark price is likely much smaller than the probability that they are willing to pay 10-20% less. In fact, we would expect to find a point in the WtP distribution beyond which buyers tail off rapidly because any higher price would be regarded as overpaying or unfair. This should result in a short right-hand tail. In contrast, we would expect a proportionately longer left-hand tail simply because consumers can lower their WtP for a particular car for any number of idiosyncratic reasons (e.g., intangibles such as the car's look and feel, etc.).

All of this suggests a negatively skewed distribution of consumers' WtP. This conceptual analysis is not limited to cars. In markets where consumers are well informed about products and prices, one might expect to find negative skewness – we give an example later based on published WtP data for soybean oil. In general, the point at which the distribution of consumers' WtP drops off could be a function of many things, including product attributes (quality, etc.), environmental attributes (store location, product display, etc.), and consumer attributes (income, loyalty, etc.). Our proposed model allows all of these attributes to be included in capturing the drop-off point.

The purpose of this paper is to offer a new discrete choice model based on negatively skewed distributions of consumer utilities, and to develop its analytical consequences on pricing and assortment planning. Because a requirement of the model is that it could be customized to each store in a retail chain, it is designed to be analytically tractable, scalable, and optimization-friendly. We call the new model the *exponomial choice* (EC) model due to its use of exponomials (Duffin and Whidden 1961) to characterize the closed-form choice probabilities.

The EC model leads to new insights on pricing and assortment planning that complement previous findings from the dominant closed-form choice models used in this area, MNL and NL. These new insights suggest a stronger link between product desirability and optimal prices/assortments; hence, they may be more in step with managerial intuition. First, our model need not conform to the "constant markup" property, which holds that the profit-maximizing price for a fixed assortment (or a nest) is some constant dollar amount (markup) added to each product's unit cost (Anderson et al. 1992; Li and Huh 2011). Instead, optimal markups obey a "hockey-stick" shape; in particular, they increase in expected utilities for products that are more desirable than the outside option. Second, our model need not conform to the "contiguous price" property, which holds that the optimal assortment for exogenous prices (selected from a universal set of products or a nest) is always some contiguous set of the most expensive products (Talluri and van Ryzin 2004; Davis et al. 2012). Instead, the optimal assortment allows leapfrogging in prices; a lower-priced product can be selected over a higher-priced product due to its favorable utility position.

This paper is devoted to development of the theoretical model as well as its analytical consequences. Not only do we present results on optimal prices and assortments, but also properties that are desirable for implementation in practice, e.g., concavity of the loglikelihood function. (All proofs are provided in the Appendix.) As for its usefulness in practice, this will be addressed in a companion paper (Anonymous 2014) that includes a three-way comparison between the EC model, a generalized version of the EC model (with choice-specific variance terms discussed in §6), and the MNL model. There, it will be shown that, using the top 25 categories of household level grocery data analyzed by Briesch et al. (2013) as our test set, the EC model was superior to MNL in out-of-sample prediction for 22 of the 25 categories, and the generalized EC model was superior for 23 categories (in all 25 categories one of our two models was superior to MNL). We believe this offers empirical support for the validity of the models introduced in this paper.

2. Literature Review

The area of research most related to our work uses discrete choice modeling as the basis for analytical pricing and assortment planning models. We organize our review of this literature around the two major families of discrete choice models used in this area: logit and Hotelling models, and their generalizations. For in-depth treatment of the theoretical foundations of discrete choice models in general, the reader is referred to Manski and McFadden (1981), Ben-Akiva and Lerman (1985), Anderson et al. (1992) and Train (2009).

Logit Model: MNL and its nested variety are the most popular discrete choice models in this stream of literature. The primary reasons for this are that they offer reasonable levels of analytical tractability coupled with years of empirical research supporting their real-world relevance (McFadden 2001). van Ryzin and Mahajan (1999) derive the structure of optimal assortment for exogenous uniform prices, assuming MNL choice behavior for consumers and newsvendor costs (with lost sales) for each product. Talluri and van Ryzin (2004) address optimization of fare classes in revenue management using a general choice model that subsumes MNL as a special case. They prove the optimal fares must form an efficient subset, that is, a subset of the most expensive products where no price is leapfrogged. Cachon et al. (2005) develop an MNL-based model of consumer search to understand how it influences optimal assortment. They show that it can be optimal to offer an unprofitable product to prevent consumer search. Again building on MNL, Aydın and Hausman (2009) investigate contractual forms a manufacturer can use to induce a retailer to make supply-chain optimal assortment decisions. Rusmevichientong et al. (2010) and Rusmevichientong and Topaloğlu (2012) study a dynamic assortment problem where the firm is attempting to learn the parameters of an MNL choice model by experimenting with its assortment. Davis et al. (2013) show under MNL choice that the problem of finding the revenue-maximizing assortment subject to a set of totally unimodular constraints can be cast as a linear program. They demonstrate how this general setup applies to some specific pricing and assortment planning problems.

NL has also been used in analytical models of pricing and assortment planning, although finding a suitable nesting structure can be difficult in practice (Koppelman and Bhat 2006; Louviere and Woodworth 1983). In retail operations, for example, it helped understand the impact of category management (Cachon and Kök 2007) and consumer returns (Alptekinoğlu and Grasas 2014) on assortment selection. Li and Huh (2011) and Gallego and Wang (2012) generalize past results on optimal pricing under NL choice. Davis et al. (2012) show that the contiguous price property due to Talluri and van Ryzin (2004) applies to each nest when selecting the optimal assortment under exogenous prices. This body of work is expanding in many interesting dimensions such as generalizing the tree structure of NL (Li et al. 2013).

Hotelling Model: Another popular discrete choice model family in this literature is due to Hotelling (1929), who originally set out to explain the 'sameness' in product variety offered by competing firms. His model is variously known as the locational choice model, the address model, or simply the Hotelling model. Because the choice probabilities are generally not closed-form, it is less tractable than MNL for pricing and assortment planning purposes, and it does not lend itself to empirical inquiry as easily. Yet it has enjoyed great conceptual appeal after it was formalized by Lancaster (1966) into a theory of consumption that views products as a bundle of characteristics and consumer preferences as points or addresses in a characteristics space. Applications of the Hotelling model in business disciplines include multichannel retail competition (Balasubramanian 1998), brand loyalty (Villas-Boas 2004), and joint inventory-price optimization (Alptekinoğlu and Corbett 2010). The Hotelling model can be tied to MNL and the general body of discrete choice models in a unified manner (Anderson et al. 1992).

Our research contributes to this literature by offering a new, analytically tractable discrete choice model that has different implications for optimal prices and assortments. These implications are consistent with managerial intuition and further supported by our ongoing empirical work noted earlier.

3. The Exponomial Choice Model

We begin with the standard random utility framework where the consumers' utility for a particular choice among an assortment of choices is composed of two parts, a deterministic part and a random part. The deterministic part in our model represents what we term the choice's *ideal utility*, an upperbound on the utility that would be obtained "for the right person," and it can be a function of many attributes (product, environment, and consumer). It reflects the preferences of an enthusiastic buyer. A literal interpretation of ideal utility would be that it is the maximum utility any consumer in the population (having the same consumer attributes) derives from a particular choice. The random part of our model captures consumer heterogeneity and has one sign (negative).

The utility that a random consumer k has for choice i in our model is the linear function

$$U_k(i) = u_i - z_{ik}$$

for $i = 1, \ldots, m$, where m is the number of choices, u_i is the ideal utility for choice i, and z_{ik} are IID (independent and identically distributed) exponential random variables with mean $1/\lambda$ (in §6 we relax this condition by allowing choice-specific parameters λ_i). For example, these m choices may consist of $m - 1$ products offered by a firm plus an *outside option* (the alternative of not buying from the firm). Observe that the maximum utility that a consumer can obtain from a particular choice is it's ideal utility, but the random term means each consumer's final utilities can assume any order and thus reflect their personal preferences. The expected utility for each choice is easily obtained from its respective ideal utility by subtracting $1/\lambda$. This utility specification results in a negatively skewed distribution of consumer utilities for each choice.

In many empirical applications, modeling the ideal utility as a function of product attributes, environmental attributes, and consumer attributes would represent the model's standard form. (Indeed, the empirical analysis in our companion paper included the environmental attributes "feature" and "display.") However, models used for assortment and price optimization generally consider only product level attributes. We do the same to simplify our exposition. We also note that consumer heterogeneity in ideal utilities can be incorporated using the following specification: $U_k(i) = (u_i + d_k) - z_{ik}$ where d_k is an IID random variable. The component d_k is differenced out

Figure 1: Density of the Exponomial Utility $U_k(i) = (u_i + d_k) - z_{ik}$ with $u_i = 40$, $d_k \sim N(0, 0.5)$, and $z_{ik} \sim exp(0.5)$.

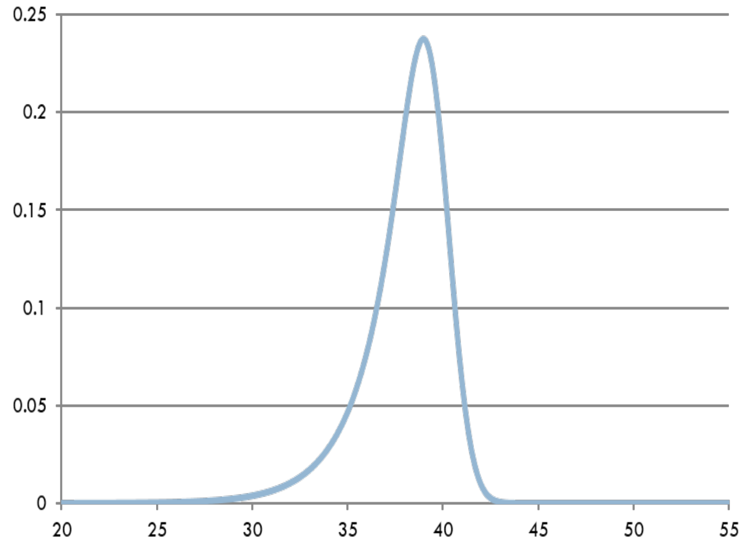

when calculating our choice probabilities, but it may help represent aggregate level consumer data better. For example, if the d_k are IID normal random variables, one might see a utility distribution like that in Figure 1.

When the choices are products to buy, the observable component of our utility specification can be operationalized as follows. Let the ideal utility for product i be $u_i = \alpha_i - \beta p_i$, where α_i captures all non-price related product attributes and can be interpreted as the *intrinsic desirability* or attractiveness of the product, p_i is the price, and β is the price sensitivity parameter. Using the approach of Besanko et al. (1998), one can then specify consumer k's maximum willingness to pay for product i as $W_k(i) = \frac{\alpha_i}{\beta} - \frac{z_{ik}}{\beta}$, where the WtP error term $\frac{z_{ik}}{\beta}$ is exponentially distributed with rate parameter $\beta\lambda$. Moreover, $W_k(i) - p_i$ is consumer k's surplus for product i, and the product that maximizes consumer k's utility is the same as the product that maximizes consumer k's surplus.

In practice, it is easier to observe WtP than utility, so we make liberal use of the connection between the two in establishing the practical relevance of our utility specification. We imagine a consumer who is well informed about various product alternatives and their market prices. Knowledge of products and prices can be obtained from a variety of sources including stated MSRPs, frequent purchases within a product category (e.g., groceries), third party information providers, or even word-of-mouth. Our model does not rely on a particular learning mechanism, it merely

Figure 2: CDF of Willingness to Pay for Genetically Modified Soybean Oil

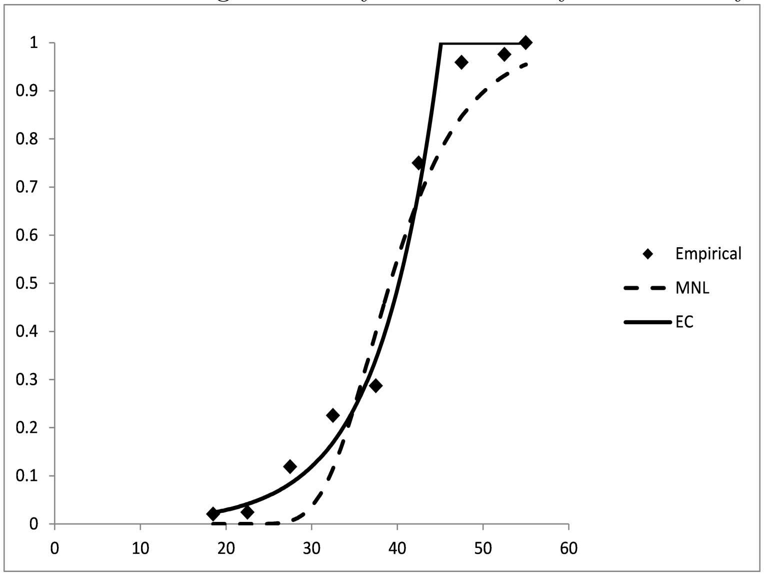

assumes that some mechanism exists. The consumer then makes an *informed* decision about their maximum willingness to pay for an item. Given extensive product/price knowledge, one would expect the distribution of their WtP to drop off rapidly above some value simply because consumers dislike overpaying. This means the WtP distribution would be negatively skewed, as would the distributions of consumer surplus and consumer utility.

Negatively skewed WtP distributions have been cited in health care (Philips et al. 2006), agriculture (Hu et al. 2006; Norwood et al. 2005), and financial security markets (Garbade and Silber 1976), among others. We could only find one paper that published data on consumers' WtP. Hu et al. (2006) collected data on consumers' WtP for genetically modified (GM) soybean oil in Nanjing, China. As noted by the authors, soybean oil is a popular product in China, and so participants were well informed about the various oil products and their market prices. The market price for GM soybean oil at the time was 40 RMB. Of the 244 participants who would buy GM soybean oil, the following WtP distribution was observed: 18.5 RMB (5), 22.5 RMB (1), 27.5 RMB (23), 32.5 RMB (26), 37.5 RMB (15), 42.5 RMB (113), 47.5 RMB (51), 52.5 RMB (4), 55 RMB and up (6). Observe that the WtP distribution has a pronounced hump slightly above the current market price with a relatively short tail to the right and a longer tail to the left. The empirical cumulative distribution function (CDF) is shown in Figure 2.

Now consider the distribution of WtP in the EC model versus its MNL analog, which is based on Gumbel distributed errors. Using the Kolmogorov distance as a measure of fit (the maximum gap between the theoretical and empirical WtP distributions, see Weber et al. 2006), we find that the EC utility framework approximates the data substantially better than the MNL model. This is even visually apparent from the tail behavior. The fitted distributions are shown in Figure 2. (The minimum Kolmogorov distance for the best EC and MNL models are 0.056 and 0.112, respectively.)

Having introduced the basic components of the EC model, we first derive the choice probabilities implied by utility maximization.

3.1 Choice Probabilities

Dropping the consumer subscript to improve readability, the probability that a consumer prefers choice i – because it offers the largest utility – is

$$\begin{aligned} Q(i) &= Prob\{u_i - z_i \geq u_j - z_j \quad \forall j,\, j \neq i\} \\ &= Prob\{z_j \geq u_j - u_i + z_i \quad \forall j,\, j \neq i\} \\ &= \int_0^\infty \prod_{j\neq i} \left[1 - F\left(u_j - u_i + z\right)\right] f(z) dz, \end{aligned} \tag{1}$$

where $f(z) = \lambda e^{-\lambda z}$ and $F(z) = 1 - e^{-\lambda z}$ for $z \geq 0$ are the probability density function (PDF) and CDF, respectively, of an exponential distribution with rate λ ($f(z) = F(z) = 0$ for $z < 0$).

The selection of an exponential distribution has considerable mathematical virtue in simplifying (1), which is an important consideration for a discrete choice model. Indeed, in the case of a linear utility model whose random part takes both positive and negative values, the analogous integral would simplify (to one of the form $\int u^k du$) provided the distribution function satisfied the property $F(x + y) = F(y)^{g(x)}$ for some function $g(x)$. It can be shown that this is the defining equation of a Gumbel distribution and the basis of the MNL model. Indeed, Yellott (1977) has demonstrated that Gumbel is the only distribution which leads to the closed-form MNL choice probability for three or more choices.

In all that follows, we assume without loss of generality that the ideal utilities are labeled in increasing order $u_1 \leq u_2 \leq \cdots \leq u_m$ (ties between ideal utilities can be broken arbitrarily).

Theorem 1. (Exponomial Choice Probabilities) *Out of m products with utilities $U(i) = u_i - z_i$*

$(u_1 \leq \cdots \leq u_m)$, where z_i follow independent exponential distributions with rate λ, the probability that the consumer chooses product $i \in \{1, \dots, m\}$, i.e., considers it as utility-maximizing, is

$$Q(i) = \frac{exp\left[-\lambda \sum_{j=i}^{m}(u_j - u_i)\right]}{m-i+1} - \sum_{k=1}^{i-1} \frac{exp\left[-\lambda \sum_{j=k}^{m}(u_j - u_k)\right]}{(m-k)(m-k+1)}. \tag{2}$$

This expression is known as an *exponomial* (a linear function of exponential terms) as first defined by Duffin and Whidden (1961). Without loss of generality, one may take $\lambda = 1$ and rescale the u's accordingly. Additionally, the choice probabilities depend on the *difference* in ideal utilities, and this means they are unaffected if translated by a common constant. This in turn means one of the ideal utilities can be set to a convenient value. We defer setting any u_i at this point, but we set $\lambda = 1$ for ease of exposition.

In the rest of this section we develop some immediate implications of the EC model.

3.2 Basic Properties

At first glance, the EC probabilities (2) appear to lack any attractive qualities, but they actually possess excellent structure. To analyze this structure in greater detail, it helps to define

$$G(i) = \frac{exp\left[-\sum_{j=i}^{m}(u_j - u_i)\right]}{m-i+1} \tag{3}$$

for $i = 1, \dots, m$. The choice probabilities in (2) can now be expressed as

$$\begin{aligned}
Q(1) &= G(1) \\
Q(2) &= G(2) - \frac{1}{m-1}G(1) \\
Q(3) &= G(3) - \frac{1}{m-2}G(2) - \frac{1}{m-1}G(1) \\
&\vdots \\
Q(m-1) &= G(m-1) - \frac{1}{2}G(m-2) - \cdots - \frac{1}{m-2}G(2) - \frac{1}{m-1}G(1) \\
Q(m) &= G(m) - G(m-1) - \frac{1}{2}G(m-2) - \cdots - \frac{1}{m-2}G(2) - \frac{1}{m-1}G(1)
\end{aligned} \tag{4}$$

Whereas it is difficult to see why the choice probabilities sum to 1 using (2), the probabilities expressed in (4) reveal this fact fairly readily: the coefficients of $G(1)$, $G(2)$, ..., $G(m-1)$ in (4) sum to zero—note there are $m-1$ terms of the form $\frac{-1}{m-1}G(1)$, $m-2$ terms of the form $\frac{-1}{m-2}G(2)$, etc. The only term that remains after summing the expressions on the right hand side of (4) is $G(m)$, which equals 1 using (3).

It is also easy to see that the choice probabilities are monotone increasing in i. Because $u_1 \le u_2 \le \cdots \le u_m$, the $G(i)$ defined in (3) must satisfy

$$(m-i+1)G(i) \quad \le \quad (m-i)G(i+1) \tag{5}$$

But then using (4),

$$Q(i+1) - Q(i) = G(i+1) - \frac{m-i+1}{m-i}G(i) \ge 0 \tag{6}$$

Hence, the higher the ideal utility of a product, the higher the choice probability associated with that product, i.e., $Q(1) \le Q(2) \le \cdots \le Q(m)$. Observe that the inequalities in both (5) and (6) are strict if and only if $u_i < u_{i+1}$.

An attractive property of EC probabilities is logconcavity.

Lemma 1. (Logconcavity) *The exponomial choice probabilities, $Q(i)$, are each a logconcave function of the ideal utilities in the region $u_1 \le u_2 \le \cdots \le u_m$.*

This result only applies to ordered ideal utilities, but it is a building block for showing that logconcavity extends to the choices' parameter space (using *unordered* utilities) provided the ideal utilities are linear functions of the unknown parameters.

In the rest of this section, we use the index j to identify and label a *product name* and not its particular ideal utility rank among available choices. In building the likelihood function this is necessary because consumers choose products and not their ranks. Suppose, for choice scenario k, consumers are offered a set of products indexed by the set S_k. Let n_{kj} represent the number of consumers who chose product j, $j \in S_k$, for choice scenario k. Assume product j has an ideal utility $u_{kj}(\gamma)$ that is a linear function of its unknown parameters, represented by the vector γ. To capture a product's rank within each choice set, we define the integer-valued rank function as (1 = lowest)

$$r_{kj}(\gamma) = rank\ of\ u_{kj}(\gamma)\ among\ all\ u_{ki}(\gamma),\ i \in S_k.$$

We maintain the notational convention that $Q(1)$ refers to the probability of the lowest utility product, $Q(2)$ refers to the probability of the second lowest utility product, etc. Thus the probability that a consumer chooses product j in choice scenario k is simply $Q(r_{kj}(\gamma))$.

Assuming consumers are exposed to a total of K choice scenarios, the loglikelihood function is

$$\mathcal{LL}(\gamma) = \sum_{k=1}^{K} \sum_{j \in S_k} n_{kj} ln\left(Q(r_{kj}(\gamma))\right). \tag{7}$$

An important property for estimation is given in the following theorem.

Theorem 2. (Concavity of Loglikelihood Function) $\mathcal{LL}(\gamma) = \sum_{k=1}^{K} \sum_{j \in S_k} n_{kj} ln\left(Q(r_{kj}(\gamma))\right)$ *is a concave function of the unknown parameters* γ*.*

Theorem 2 is the primary reason estimating the unknown parameters is both fast and easy on real data, something that we experienced first-hand by applications to 25 different categories of consumer level grocery data (Anonymous 2014).

3.3 Patterns of Cannibalization

Inherent in any discrete choice setting, new choices cannibalize existing demand. In this subsection we show that cannibalization behaves differently in EC than MNL. In the MNL model, the IIA (independence of irrelevant alternatives) property holds that the ratio of two choice probabilities is constant regardless of the remaining alternatives in the choice set. This implies demand cannibalization is proportional across existing products. The IIA property does not hold for the EC model. A simple numerical example should make this clear.

Numerical Example. Consider an initial choice set consisting of three products with ideal utilities $u_1 = 1.1$, $u_2 = 1.2$, $u_3 = 1.3$, plus an outside option (no buy) with ideal utility $u_0 = 1.0$ (and zero expected utility, given $\lambda = 1$). Then consider an expanded choice set that includes a fourth product with ideal utility $u_4 = 1.4$. The choice probabilities for both scenarios are given in Table 1 below.

Table 1: Choice probabilities for three and four product examples.

	Outside	Product 1	Product 2	Product 3	Product 4
Ideal Utilities	1	1.1	1.2	1.3	1.4
Probabilities	0.137	0.201	0.283	0.378	NA
Probabilities	0.074	0.119	0.183	0.265	0.36

The introduction of product 4 reduces the choice probability of product 1 by a relative fraction of (.201-.119)/.201 = .408, for product 2 by (.283-.183)/.283 = .353, and for product 3 by (.378-.265)/.378 = .299. In other words, from an assortment planning perspective, introducing a product that is more attractive than all the alternatives in a choice set appears to have a disproportionately larger impact (in percentage terms) on the less attractive products. This turns out to be provable.

Theorem 3. (Addition of a Superior Choice) *Consider ideal utilities* $u_1 \leq u_2 \leq \cdots \leq u_I$. *Let the choice probability for item* i *be denoted by* $Q^I(i)$. *Suppose a new item with ideal utility* u_{I+1} *is introduced to the choice set so that* $u_1 \leq u_2 \leq \cdots \leq u_I \leq u_{I+1}$. *Let the new choice probabilities for the expanded set be denoted by* $Q^{I+1}(i)$. *Then*

$$\frac{Q^{I+1}(i)}{Q^I(i)} \leq \frac{Q^{I+1}(i+1)}{Q^I(i+1)} \quad i = 1, 2, \ldots, I-1.$$

Now consider the relative change in choice probabilities as in the numerical example above. Theorem 3 implies $\frac{Q^I(i)-Q^{I+1}(i)}{Q^I(i)} \geq \frac{Q^I(i+1)-Q^{I+1}(i+1)}{Q^I(i+1)}$ for $i = 1, 2, \ldots, I-1$, that is, when a superior product is added to the choice set, the percent reduction in choice probability is greater for products with lower ideal utility. Therefore, unlike MNL, the EC model implies that new products can cannibalize existing demand at non-proportional rates.

A related result is obtained by studying the elasticities of the choice probabilities. It is not hard to show using (3) and (4) that the elasticities are

$$E_k^i \equiv \frac{\partial Q(i)}{\partial u_k} \frac{u_k}{Q(i)} = \begin{cases} -u_k & i < k \\ u_k \cdot \left[\frac{(m-k+1)G(k)}{Q(k)} - 1\right] & i = k. \\ -u_k \cdot \frac{Q(k)}{Q(i)} & i > k \end{cases} \tag{8}$$

Observe that the monotonicity of choice probabilities implies that $E_k^i < E_k^{i'} \leq 0$ for any three choices $i < k < i'$ with $u_i < u_k$ or $u_k < u_{i'}$. This means that when the ideal utility for a mid-tier product is increased (as might occur if its price were reduced on a promotion), it has a proportionately larger effect on products of lower ideal utility. To the extent that higher product quality translates into higher ideal utility, this result is fully consistent with the empirical findings of Blattberg and Wisniewski (1989), who showed that price reductions in mid-tier quality products stole market share disproportionately from lower-tier products.

4. Pricing under Exponomial Choice

When choices are products to buy, the consumer utility is typically a function of many factors (quality, aesthetics, etc.) including price. We assume that the ideal utility of product j is separable and linear in price p_j, i.e.,

$$u_j(p_j) = \alpha_j - \beta p_j \qquad j = 1, ..., m. \tag{9}$$

The intercept α_j captures all non-price related factors and measures the *intrinsic desirability* of the product, i.e., the ideal utility of product j if $p_j = 0$. The coefficient β ($\beta > 0$) captures the price sensitivity of consumers. It is important to note that the labeling of products in (9) does not imply any ordering of their ideal utilities as it does in Theorem 1; this is because the ordering of products via their ideal utilities involves prices. It remains to be seen how the ideal utilities of these products are ordered once their prices are decided, which is the focus of our structural results in this section.

The objective is to maximize expected revenue, but profit maximization—with marginal cost c_j per unit of product j—can be easily accommodated. Rewriting the ideal utility equation in (9) as $u_j(p_j) = (\alpha_j - \beta c_j) - \beta(p_j - c_j)$, it is easy to see that the profit maximization model would have profit margin $(p_j - c_j)$ in place of price, and $(\alpha_j - \beta c_j)$ would become the measure of intrinsic desirability (= the ideal utility of product j assuming $p_j = c_j$). Hence, we set the marginal cost of all products to zero without loss of generality.

4.1 Optimal Pricing for a Fixed Assortment

In revenue management as well as assortment planning, it is common to optimize the expected revenue of a given assortment. In the setting we consider next, a fixed assortment is offered to the market and a manager coördinates prices to optimize the expected revenues from the assortment.

Suppose for the moment that we knew how the products' ideal utilities (including the outside option) were ordered at optimality. We could then relabel products so that $u_1 \leq u_2 \leq \cdots \leq u_m$. Additionally suppose the outside option were in the n^{th} position ($1 \leq n \leq m$). The remaining $m-1$ options are then products offered by a single firm that wants to price them optimally. The firm's expected revenue per choice could then be expressed as

$$R(p_1, \ldots, p_m) \quad = \quad \sum_{i=1}^{m} p_i Q(i). \tag{10}$$

Observe that we must impose the constraint $p_n = 0$ (price of the outside option) to make the revenue expression valid. This does not sacrifice any generality; because the price of the outside option is fixed, its original price component can be rolled into the intercept and its price reset to 0.

We show next that we essentially know the optimal ordering of final ideal utilities of products based on their intercepts; only the position of the outside option is unknown. The result rests on a simple swapping argument.

Theorem 4. (Intrinsic Desirability of Products and Their Optimal Prices) *Suppose all products have been relabeled so that $u_1 \leq u_2 \leq \cdots \leq u_m$ at optimal prices, and the outside option is in position n. Then it cannot happen in the revenue maximizing solution that $u_i(p_i) > u_{i'}(p_{i'})$ and $\alpha_i < \alpha_{i'}$ for two arbitrary products i and i' ($i \neq n$, $i' \neq n$).*

The importance of the preceding theorem is that it implies intrinsically less desirable products (as measured by their α_i) cannot be assigned higher ideal utilities (and therefore higher probabilities) than intrinsically more desirable products once prices are optimized. This makes the search for optimal prices considerably easier. Once the position of the outside option is known—and there are only m possibilities—the remaining positions are taken by products in order of increasing α_i. This means the particular choice probability formula (2) to apply to each product is known as well. To solve the firm's pricing problem, one may assume each possible position for the outside option and solve m optimization problems. We refer to these as *conditional price optimization* problems because they are conditioned on the position of the outside option.

Each conditional optimization can be transformed into a separable concave programming problem, which means standard optimization techniques can be used for its solution. We now outline the key steps of this transformation.

Conditional Price Optimization. The ideal utility of the outside option is assumed to be in position n $(1 \leq n \leq m)$ with $p_n = 0$ at optimality. The firm's $m-1$ products are assigned to the remaining positions in order of increasing intrinsic desirability. For the purposes of this conditional optimization, products and prices are henceforth labeled according to their position in this list.

By Theorem 4 and our assumption regarding the position of the outside option at optimality, the optimal prices must satisfy $u_1(p_1) \leq u_2(p_2) \leq \cdots \leq u_m(p_m)$. The structure of the probabilities in (4) implies the expected revenue defined in (10) can be expressed as

$$\begin{aligned} R(p_1,\ldots,p_m) \quad &= \quad G(1)\left[p_1 - \frac{1}{m-1}\sum_{j=2}^{m} p_j\right] + G(2)\left[p_2 - \frac{1}{m-2}\sum_{j=3}^{m} p_j\right] \\ &\quad + \cdots + G(m-1)\left[p_{m-1} - p_m\right] + G(m)\left[p_m\right]. \end{aligned} \tag{11}$$

We transform the price variables into new "y" variables using the linear transformation $y = Ap$, where $p^T = (p_1, p_2, \ldots, p_m)$, $y^T = (y_1, y_2, \ldots, y_m)$, and the matrix A is upper triangular with elements $A_{ij} = \frac{-1}{m-i}$ for $j > i$, $A_{ij} = 1$ for $j = i$, and $A_{ij} = 0$ for $j < i$ (the rows of A are the coefficients of prices in the bracketed terms of equation (11)). We then define parameters

$$\begin{aligned} \delta_i \quad &= \quad \sum_{j=i+1}^{m} (\alpha_j - \alpha_i) \quad i = 1, 2, \ldots, m-1 \\ \delta_m \quad &= \quad 0. \end{aligned}$$

The y's are then transformed to equivalent $w's$ using the nonlinear transformation

$$\begin{aligned} w_i \quad &= \quad exp\left[-\delta_i - (m-i)\beta y_i\right] \quad i = 1, 2, \ldots, m-1 \\ w_m \quad &= \quad z_m. \end{aligned}$$

Table 2: Solutions to the five conditional price optimization problems.

n	w_1	w_2	w_3	w_4	w_5	$p_{(1)}$	$p_{(2)}$	$p_{(3)}$	$p_{(4)}$	$p_{(5)}$	Exp. Rev.
1 (Lowest)	0.63	0.63	0.63	0.79	1.65	0	1	1.1	1.39	1.65	1.237
2	0.2	0.67	0.67	0.79	1.68	1.3	0	1.1	1.42	1.68	1.253
3	0.2	0.3	0.72	0.79	1.72	1.39	1.39	0	1.45	1.72	1.268
4	0.2	0.3	0.76	0.76	1.72	1.41	1.41	1.5	0	1.72	1.266
5 (Highest)	0.2	0.3	1	1	0	1.5	1.5	1.5	2	0	1.202

The equivalent conditional optimization to be solved is

$$Max \quad R(w_1, \ldots, w_m) = w_m + \textstyle\sum_{i=1}^{m-1} \frac{w_i}{(m-i+1)} \cdot \frac{[ln(w_i)+\delta_i]}{-(m-i)\beta} \tag{12}$$

$$s.t. \quad -\frac{ln(w_n)+\delta_n}{(m-n)\beta} - \frac{ln(w_{n+1})+\delta_{n+1}}{(m-n)(m-n-1)\beta} - \frac{ln(w_{n+2})+\delta_{n+2}}{(m-n-1)(m-n-2)\beta} - \cdots - \frac{ln(w_{m-1})+\delta_{m-1}}{2\beta} + w_m \leq 0 \tag{13}$$

$$w_1 \leq w_2 \leq \cdots \leq w_{m-1} \leq 1 \tag{14}$$

The objective function (12) is concave and separable. The nonlinear but separable and convex constraint in (13) ensures that the outside option has price $p_n = 0$. The linear inequality constraints in (14) ensure the optimal utilities satisfy the ordering of the u_i specified at the outset. To determine the optimal revenue, one solves m different versions of the problem formulated in (12)-(14), each one adjusted to account for the conditional position of the outside option. This optimization has good structure and could even be solved using a piecewise linear approximation and thus linear programming.

Numerical Example. Consider a set of four products with ideal utilities $u_1 = 9-p_1$, $u_2 = 9.1-p_2$, $u_3 = 9.5 - p_3$, $u_4 = 10 - p_4$, and outside option $u_0 = 8$. There are five different conditional price optimization problems to solve. The solutions to each of these as well as the corresponding prices and revenues are given in Table 2. The w's are those that solve (12)-(14), whereas the prices $p_{(1)}, \ldots, p_{(5)}$ correspond to the prices of the *ranked alternatives*; for example, when the outside option is constrained to be the least desirable product, its price is $p_{(1)} = 0$, and the remaining prices $p_{(2)} = 1.00$, $p_{(3)} = 1.10$, $p_{(4)} = 1.39$, $p_{(5)} = 1.65$ reflect the prices for products 1, 2, 3, and 4 respectively. With the exception of the outside option (whose price is constrained to be 0), products are always ranked in ascending order of intrinsic desirability, i.e., by their intercepts α_i. In this example, the global optimal prices are those where the outside option is ranked third from the bottom (total revenue of 1.268) and the corresponding prices are 1.39 (product 1), 1.39 (product 2), 1.45 (product 3), and 1.72 (product 4). It is easy to check that the ideal utilities for the five alternatives are 7.61 (product 1), 7.71 (product 2), 8.00 (outside option), 8.05 (product 3), and 8.28 (product 4). The associated choice probabilities are .040, .065, .205, .240, .450. Note that 20.5% of potential buyers do not buy (they choose the outside option).

A closer examination of the results suggests that those products whose ideal utilities are below that of the outside option have equal prices. For example, the final optimal prices are $1.39 for both products 1 and 2. This is not a coincidence, as the following theorem explains. Without loss of generality, we assume the price of the outside option has been reset to 0 (i.e., its price term has been rolled into its intercept) and the products are indexed and relabeled according to their optimal ideal utilities $u_1 \leq u_2 \leq \cdots \leq u_m$. Recall that we set the marginal cost of all products to zero without loss of generality. Prices can thus be viewed as equivalent to profit margins.

Theorem 5. (Monotonicity of Optimal Prices) *Assume product i's ideal utility is linear in prices ($u_i = \alpha_i - \beta p_i$) and the α_i are distinct (to avoid ties). Then (**a**) all products whose optimal ideal utilities are less than the outside option have identical optimal prices; (**b**) all products whose optimal ideal utilities are greater than the outside option have optimal prices that increase monotonically. That is, $p_1^* = \cdots = p_{n-1}^* < p_{n+1}^* < \cdots < p_m^*$, where n is the ideal utility position of the outside option.*

The theorem is relatively easy to visualize. The optimal prices follow a "hockey stick" pattern, where the outside option divides the constant price of the lower utility products (the "shaft") from the increasing prices of the higher utility products (the "blade").

This result yields an appealing managerial insight. It prescribes profit margins or markups that must increase in expected utilities for products that are more desirable than the outside option, but a constant markup for all other products. Thus, offering higher quality—one way to increase the utility of a product—means higher profit margins only if the product is ultimately more desirable than the outside option. In contrast, the MNL model cannot accommodate increasing margins; it prescribes a constant markup for all products in the assortment (Anderson et al. 1992).

Suppose all products including the outside option are relabeled in order of their optimal ideal utilities (after price optimization is complete). Then an interesting consequence of the proof of

Theorem 5 is that for any product whose optimal ideal utility is less than the outside option, its (constant) price satisfies the equation

$$p_i = \overline{p}_{i+1} + \frac{1}{(m-i)\beta} \tag{15}$$

where p_i is the price of the product having the i-th lowest utility (i.e., u_i), $\overline{p}_{i+1}$ is the average price of all products $\{i+1, i+2, \ldots, m\}$ (i.e., the set of products having higher ideal utilities than product i), and β is the price sensitivity parameter. In the special case where the outside option has the highest optimal utility (and so has index m), we have $p_m = 0$, $p_{m-1} = +\frac{1}{\beta}$, and so $p_i = +\frac{1}{\beta}$ for $i = 1, 2, \ldots, m-2$. When marginal costs c_i are included (recall intercepts are then adjusted and the p_i represent margins), the optimal sticker price is $p_i^* = c_i + \frac{1}{\beta}$. Observe that this special case of EC echoes the constant markup property of MNL, except the EC formula is closed-form and represents a form of "pass through" pricing based solely on the consumers' price sensitivity parameter.

4.2 Price Competition Among Single-Product Firms

In contrast to a manager coördinating all prices for the assortment (§4.1), imagine $m-1$ firms each offering a single product and setting its own price independently through oligopolistic market competition. Although the prices ultimately determined would differ from those selected by a monopolist, we show next that the products' optimal ideal utilities would still be ordered according to each product's intrinsic desirability (akin to Theorem 4). This means there is at least agreement on how products should be ranked in monopoly and oligopoly solutions.

Assume we know which product occupies position i in the list of ordered ideal utilities in the Nash equilibrium, and assume the ideal utility for the product is a linear function of price ($u_i = \alpha_i - \beta p_i$ with $\beta > 0$). The unconstrained optimal revenue $R_i(u_i)$ obtained by the firm owning product i, expressed as a function of the ideal utility u_i, is

$$R_i(u_i) = \frac{\alpha_i - u_i}{\beta} \cdot Q(i) \tag{16}$$

where $u_i \leq \alpha_i$. Differentiation of equation (16) with respect to u_i reveals the optimal ideal utility satisfies the equation

$$(\alpha_i - u_i) = \frac{Q(i)}{\frac{\partial Q(i)}{\partial u_i}}.$$

The right hand side ratio $Q(i)/\frac{\partial Q(i)}{\partial u_i}$ has three important properties summarized below.

Property 1. *(**a**) Monotonicity and Positivity: $Q(i)/\frac{\partial Q(i)}{\partial u_i}$ is monotone nondecreasing in u_i and strictly positive. (**b**) Agreement on Boundary: $Q(i)/\frac{\partial Q(i)}{\partial u_i} = Q(i-1)/\frac{\partial Q(i-1)}{\partial u_{i-1}}$ when $u_i = u_{i-1}$, i.e., $Q(i)/\frac{\partial Q(i)}{\partial u_i} \mid_{u_i=u_{i-1}} = Q(i-1)/\frac{\partial Q(i-1)}{\partial u_{i-1}}$. (**c**) Independence from Higher-Ranked Ideal Utilities: The ratio $Q(i)/\frac{\partial Q(i)}{\partial u_i}$ is independent of the ideal utilities $u_{i+1}, \ldots, u_m$.*

These three simple properties allow us to establish that the following construction process results in a Nash equilibrium.

Construction of a Nash Equilibrium. Suppose each product is offered by a separate firm and the products have been ordered (relabeled) according to their intrinsic desirability as $\alpha'_1, \alpha'_2, ...\alpha'_{m-1}$; the primes ($'$) signify an initial and temporary ordering because the outside option has yet to be included. Let the outside option have ideal utility α_0. Beginning with the least desirable product, the one having intercept α'_1, solve

$$(\alpha'_1 - u_1) = \frac{Q(1)}{\frac{\partial Q(1)}{\partial u_1}} = \frac{1}{(m-1)}$$

The optimal solution is $u_1 = \alpha'_1 - \frac{1}{m-1}$. If $u_1 \leq \alpha_0$, set $\alpha_1 = \alpha'_1$, $u_1^* = \alpha_1 - \frac{1}{m-1}$. If $u_1 > \alpha_0$, set $\alpha_1 = \alpha_0$, $u_1^* = \alpha_0$ and relabel the remaining temporary intercepts using $\alpha'_{j-1} \to \alpha'_j$ for $j = 2, ..., m$. The general step is to find u_i that solves

$$(\alpha'_i - u_i) = \left[\frac{Q(i)}{\frac{\partial Q(i)}{\partial u_i}} \right]_{u_k = u_k^* \; k=1,2,...,i-1} \tag{17}$$

Observe that the previous optimal ideal utilities for less desirable products ($k = 1, 2, ..., i-1$) are substituted and fixed in this equation. That this equation always has a solution with $u_i \geq u_{i-1}^*$

requires some inductive reasoning. Observe

$$\begin{aligned} \alpha_i' - u_{i-1}^* \geq \alpha_{i-1}' - u_{i-1}^* \quad &\geq \quad \left[\frac{Q(i-1)}{\frac{\partial Q(i-1)}{\partial u_{i-1}}} \right]_{u_k = u_k^* \; k=1,2,\ldots,i-1} \\ &= \quad \left[\frac{Q(i)}{\frac{\partial Q(i)}{\partial u_i}} \right]_{u_i = u_{i-1}^*; \, u_k = u_k^* \; k=1,2,\ldots,i-1} \end{aligned} \tag{18}$$

The first inequality is because the α_i' are nondecreasing; the second inequality is because u_{i-1}^* either solves its version of equation (17) or equals α_0 (in which case we have ">" for the second $\geq$ in 18); the last equality is due to Property 1(b). However, the right hand side of (17) is positive nondecreasing by Property 1(a), and the left hand side of (17) equals 0 when $u_i = \alpha_i'$; thus (17) has a solution u_i on the interval $[u_{i-1}^*, \alpha_i']$. If this solution satisfies $u_i \leq \alpha_0$, then set $u_i^* = u_i$ and $\alpha_i = \alpha_i'$. If $u_i > \alpha_0$, then set $u_i^* = \alpha_0$ and relabel the remaining temporary intercepts using $\alpha_{j-1}' \to \alpha_j'$ for $j = i+1, ..., m$.

That this produces a Nash equilibrium may require some additional thought. We have shown the construction process naturally orders the optimal ideal utilities $u_1^* \leq u_2^* \leq \cdots \leq u_m^*$. Moreover, we claim no firm has an incentive to deviate from their optimal ideal utility. This is because each firm selected its own optimal ideal utility with full knowledge of all products having lesser ideal utilities, and the selections by all firms having products with higher ideal utilities have no impact on their decision (see Property 1(c)). Thus the construction does indeed produce a Nash equilibrium in the ideal utilities. Optimal prices can be backed out by inverting the linear utility-price equation for each product.

Numerical Example. Consider the previous four-product example with ideal utilities $u_1 = 9 - p_1$, $u_2 = 9.1 - p_2$, $u_3 = 9.5 - p_3$, $u_4 = 10 - p_4$, and outside option $\alpha_0 = 8$. Here, $\alpha_1' = 9$, $\alpha_2' = 9.1$, $\alpha_3' = 9.5$, $\alpha_4' = 10$. Solving for $\alpha_1' - u_1 = \frac{1}{m-1}$ yields $u_1 = 9 - \frac{1}{4} = 8.75$. Because $8.75 > \alpha_0 = 8$, we set $u_1^* = 8$, $\alpha_1 = 8$ and relabel the remaining intercepts $\alpha_2' = 9$, $\alpha_3' = 9.1$, $\alpha_4' = 9.5$, $\alpha_5' = 10$. Fixing $u_1^* = 8$ and solving $\alpha_2' - u_2 = Q(2) / \frac{\partial Q(2)}{\partial u_2}$ yields $u_2 = 8.673$, and so we set $u_2^* = 8.673$, $\alpha_2 = 9$. Continuing in this fashion we obtain the results in Table 3. The total expected revenue is .6286, approximately half of what was achieved by a monopolist setting prices. As expected, competition brings down prices, and only 0.58% of consumers fail to buy a product. Recall the corresponding figure was 20.5% in the monopoly case.

Table 3: Nash equilibrium in prices.

Choice	u*	Price	Probability	Exp. Rev.
No Buy	8	---	0.0058	---
1	8.673	0.327	0.1062	0.0348
2	8.743	0.357	0.1399	0.05
3	8.966	0.534	0.2891	0.1545
4	9.152	0.848	0.4589	0.3893

5. Assortment Planning under Exponomial Choice

In this section we treat two canonical assortment planning problems with exogenous and endogenous prices.

5.1 Optimal Assortment Given Exogenous Prices

Here we assume the assortment is endogenous but prices are exogenous. This occurs, for example, when a manufacturer forces its retailers to follow an MSRP.

The IIA property of the MNL model implies that, when a new product is introduced to the choice set, the choice probabilities for all prior products are reduced in equal proportion. This is enough to imply that the optimal assortment consists of some contiguous subset of the most expensive products (Talluri and van Ryzin 2004). One can construct an algorithmic proof of this result that highlights the role of the IIA property (details available upon request).

In contrast, the EC model does not exhibit the IIA property and so optimizing the assortment is not as easy. However, this is not without some interesting consequences. The optimal assortment in the EC model does not need to be a contiguous set of the most expensive choices. In other words, rank-ordering products by price from the highest to the lowest, the EC-optimal assortment may leapfrog some products with higher prices in favor of some products with lower prices. For example, Table 4 shows 10 products plus an outside option (choice 11) ranked in order of descending prices. The optimal assortment $\{1, 2, 4, 6, 7\}$ skips products 3 and 5 even though they have higher prices than products 6 and 7. The intuitive reason is that the skipped products—which have high ideal utilities—would significantly cannibalize the higher priced choices.

Finding an optimal assortment 100% of the time proved to be a daunting task given the combi-

Table 4: Optimal assortment vs. heuristic solution (exogenous prices)

	u = α−p	α	p	Optimal	Heuristic
1	3.51	8.73	5.22	×	O
2	-5.91	-0.70	5.20	×	O
3	4.07	9.06	4.98		O
4	2.25	6.58	4.34	×	
5	7.06	10.96	3.90		
6	-1.24	2.33	3.57	×	O
7	-5.20	-1.84	3.35	×	
8	1.35	3.77	2.42		
9	0.18	1.28	1.10		
10	-1.50	-0.56	0.94		
11	1	1	0	NA	NA

natorial complexity of the problem. However, we now describe a simple and very fast heuristic that solves the problem with near certainty. The heuristic is based on the idea of backward elimination: starting with the complete list of products, eliminate one product at a time until expected revenue stops improving.

Backward Elimination Heuristic for Assortment Optimization under Exogenous Prices:

1. Initialize the assortment to include all products and the outside option, $S = \{1, 2, ..., m\}$. Let index n denote the outside option.

2. Identify the product in $S \setminus \{n\}$ that would improve the expected revenue the most if removed, $i = argmax\{R(S \setminus \{i\}) - R(S) \,|\, i \in S \setminus \{n\}\}$, where $R(S_0)$ is the expected revenue defined in (10) for any choice set S_0.

3. If no product improves the expected revenue when removed, then stop. Otherwise, remove product i from S, i.e., $S \leftarrow S \setminus \{i\}$, and repeat step 2.

We report two numerical experiments to help understand the skipping behavior and the heuristic's performance. Both experiments involved solving $500,000$ randomly generated problems with 10 products plus an outside option ($m = 11$) using complete enumeration versus backward elimination.

Experiment 1 (E1) is a conservative test of skipping; ideal utilities and prices were randomly drawn but monotone, which tends to reduce the incidence of skips. We generated these parameters by taking independent draws from uniform and normal distributions as follows: $u_1 = X_1$ and

$u_i = u_{i-1} + X_i$ for $i = 2, ..., 10$, where $X_i \sim U(0, 1)$ and $U(a, b)$ denotes the uniform distribution with support $[a, b]$; $p_1 = Y_1$ and $p_i = p_{i-1} + Y_i$ for $i = 2, ..., 10$, where $Y_i \sim U(0, 1)$; $u_{11} \sim N(5, 2)$, where $N(\mu, \sigma)$ denotes the normal distribution with mean μ and standard deviation σ, and $p_{11} = 0$ for the outside option; and $\alpha_i = u_i + p_i$ for $i = 1, ..., 11$.

Experiment 2 (E2) is a conservative test of the heuristic; intrinsic desirability parameters (α_i) and prices (p_i) were independently drawn with no particular ordering, which creates large variability in choices and tends to influence the performance of the heuristic adversely. We generated these parameters as follows: $\alpha_i \sim U(-4, 12)$ and $p_i \sim U(0, 6)$ for $i = 1, ..., 10$; $\alpha_{11} = 1$ and $p_{11} = 0$ for the outside option; and $u_i = \alpha_i - p_i$ for $i = 1, ..., 11$.

We observed one or more skipped products in the optimal solution for 7.74% and 29.36% of E1 and E2 problem instances, respectively. Skips, defined as any product not included in the optimal assortment that contained a lower-priced product, tended to happen when the skipped product was sufficiently attractive that it would cannibalize too much demand from higher-priced products. Within the (sub)sample where skipping occurred, the modal number of skips was 1 and the maximum was 7 in both experiments. In many cases, the optimal solution had multiple runs of skipped products. Hence, skips had a substantial presence in optimal assortments, even for highly structured problem instances with ordered ideal utilities and prices (E1).

Out of all randomly generated E1 and E2 problem instances, the heuristic solution matched the optimal solution 99.87% and 99.97% of the time (645 and 158 mismatches, respectively). In those rare cases where there was a mismatch, the optimality gap was typically quite small, averaging 0.05% (E1) and 1.71% (E2) weighted by optimal revenues. Only 27 of the mismatches in E2 had optimality gaps in excess of 10%, and those were all low-revenue problems (average optimal revenue of 0.98 vs. 4.33 for all problems) where even small misses meant large percentage gaps. Two-thirds (105) of the mismatches in E2 had gaps below 3%. As for the example in Table 4—one of the mismatches from E2—the heuristic produced a notably different assortment, but missed the optimal revenue of 5.055 by only 0.215%. In sum: E2 produced less frequent but worse mismatches than E1; yet, even with the extreme choice variability in E2, the heuristic performed remarkably well; it produced the optimal solution with near-perfect certainty and accuracy when ideal utilities and prices were ordered (E1).

To explain why such a heuristic should work well in practice, we consider its application to a class of problems with good structure–those where the choice model is MNL. The idea here is similar to that of testing a nonlinear maximization technique on negative definite quadratic forms. If it works well on a class of problems with good structure, it bodes well for its performance in general. This is the motivation for the following result.

Theorem 6. *Assume prices are exogenous and the choice model is MNL. Backward elimination produces an optimal assortment in a finite number of steps.*

A related problem is to find the optimal assortment having no more than C different products. Rusmevichientong et al. (2010) and Farias et al. (2011) developed optimal algorithms based on a series of product additions and exchanges. In contrast to their forward-searching approach, our algorithm is based on backward elimination and involves no exchanges—once a product is removed it never returns. Theorem 6 suggests backward elimination may also be a valuable idea.

Surprisingly, the case where both the assortment and prices are endogenous can be solved optimally without a heuristic. This case is treated in the next subsection.

5.2 Optimal Assortment and Prices

In order to analyze the case where the assortment and prices are both endogenous, we prove an important structural result that reduces the number of possible assortments to $m-1$. Once again the notion of intrinsic desirability is front and center.

Theorem 7. (Structure of the Optimal Assortment under Endogenous Prices) *The optimal assortment is composed of products with highest intrinsic desirability, i.e., labeling the products such that* $\alpha_1 \le \alpha_2 \le \cdots \le \alpha_m$, $S^* = \{m-k, \ldots, m\} \setminus \{n\}$ *for some* $k \in \{0, 1, \ldots, m-1\}$, *where* n *is the position of the outside option* ($1 \le n \le m$).

The importance of the preceding theorem is that we can now optimize both the assortment and prices simultaneously. Theorem 7 indicates only $m-1$ possible assortments. The optimal prices for each of these $m-1$ assortments can be computed using the pricing techniques described in §4.1. The optimal revenues can then be computed and compared so that both the optimal assortment and the optimal prices are obtained jointly. Note that the structure shown in Theorem 7 is not necessarily optimal when prices are given exogenously (see Table 4 for an example, where the two most intrinsically desirable products, 3 and 5, are not in the optimal assortment).

6. Concluding Remarks

We propose a new, tractable discrete choice model for use in pricing and assortment planning. The model is based on the premise that consumers would not choose to overpay if they were well-informed about products and prices, thus creating a negatively skewed willingness-to-pay distribution. The model has closed-form choice probabilities that are linear functions of exponential terms. It implies a concave loglikelihood function and disproportional demand cannibalization between products (i.e., it violates the IIA property).

6.1 Summary of Pricing and Assortment Planning Insights

We offer two significant managerial insights based on the new model, which contrast with the prevailing choice models (MNL and NL) currently used in price and assortment optimization. First, the EC model shows that optimal markups (profit margins) increase in expected utilities for all products that are more desirable than the outside option. Second, when deciding optimal assortments given prices set by management, the EC model shows that skipping or leapfrogging higher-priced products in favor of lower-priced ones – due to their favorable utility position – can be optimal. The managerial appeal of these insights lies in product desirability (and the resulting utility) playing a major role in optimal pricing and assortment planning decisions. By contrast, logit-based models imply constant margins and no skipping, entirely independent of the products' desirability levels.

6.2 Extensions and Future Research Directions

Avenues for future research include (i) applying the EC model in areas that use discrete choice models as analytical building blocks and (ii) pursuing various generalizations of the EC model. With respect to (i), it would be interesting to see, for example, what the EC model may bring to bear on choice-based revenue management (e.g., Chaneton and Vulcano 2011) and pricing (e.g., van Ryzin 2013), inventory management under price-dependent demand (e.g., Aydın and Porteus 2008), product line design (e.g., Alptekinoğlu and Corbett 2010) and multiproduct competition (e.g., Cachon et al. 2008). Because optimal pricing has good structure in the EC model (and even a simple closed-form expression under certain conditions), it could potentially simplify more complex problems if used as the embedded choice model.

With respect to (ii), the baseline version of the EC model generalizes in many ways. For example, independent but non-identical random terms—requiring a rate parameter λ_i associated with each choice i—would still yield closed-form choice probabilities, which we state without proof.

Theorem 8. (Generalized Exponomial Choice Probabilities) *Out of m products with utilities $U(i) = u_i - z_i$ ($u_1 \leq \cdots \leq u_m$), where z_i follow independent exponential distributions with rate λ_i, the probability that the consumer chooses product $i \in \{1, \ldots, m\}$, i.e., considers it as utility-maximizing, is*

$$Q(i) = \frac{\lambda_i}{L_i} exp\left[-\sum_{j=i}^{m} \lambda_j (u_j - u_i)\right] - \sum_{k=1}^{i-1} \frac{\lambda_i \lambda_k}{L_k L_{k+1}} exp\left[-\sum_{j=k}^{m} \lambda_j (u_j - u_k)\right] \tag{19}$$

where $L_i \equiv \sum_{j=i}^{m} \lambda_j$ for $i = 1, \ldots, m$.

This generalization allows error terms with different variances, which is not analytically tractable in the logit framework. It may also enable richer choice settings; as a case in point, we know the choice probabilities are no longer necessarily ascending in ideal utilities (instead, it can be shown that the ratio $Q(i)/\lambda_i$ is monotone increasing in i), but this could be viewed as a positive feature in some settings.

The EC model can be nested using a fairly simple formula for the expected value of maximum utility out of a nest, which we omit, but the model may no longer be consistent with utility maximization (although it is consistent with a two-stage sequential choice framework, one described by Anderson et al. (1992) for NL). A mixed version is certainly possible using existing simulation techniques, but there is also the possibility of closed-form expressions. This is because the choice probabilities in the mixed version involve an integral with respect to the distribution of unknown parameters, and the EC model appears to be more integral-friendly than the existing choice models.

Despite promising empirical results on grocery data (Anonymous 2014), a full understanding of the EC model's strengths and weaknesses in empirical applications requires much more work. Fortunately, there is a large and well established literature to draw from on choice estimation (Train 2009). Structural results we present in this paper, especially the concavity of the loglikelihood function (Theorem 2), pave the way for the EC model to be considered as an empirical building block in much the way that MNL has (e.g., Guadagni and Little 1983). Indeed, the vast array of results and machinery now available in the choice modeling literature should serve as a valuable guide for future parallel investigations of the EC model.

References

Alptekinoğlu, A. and C. J. Corbett (2010). Leadtime - variety tradeoff in product differentiation. *Manufacturing Service Oper. Management 12*(4), 569–582.

Alptekinoğlu, A. and A. Grasas (2014). When to carry eccentric products? Optimal retail assortment under consumer returns. *Production Oper. Management 23*(5), 877–892.

Anderson, S. P., A. de Palma, and J. F. Thisse (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, MA.

Aydın, G. and W. H. Hausman (2009). The role of slotting fees in the coordination of assortment decisions. *Production Oper. Management 18*(6), 635–652.

Aydın, G. and E. L. Porteus (2008). Joint inventory and pricing decisions for an assortment. *Oper. Res. 56*(5), 1247–1255.

Balasubramanian, S. (1998). Mail versus mall: A strategic analysis of competition between direct marketers and conventional retailers. *Marketing Sci. 17*(3), 181–195.

Ben-Akiva, M. and S. R. Lerman (1985). *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press, Cambridge, MA.

Besanko, D., S. Gupta, and D. Jain (1998). Logit demand estimation under competitive pricing behavior: An equilibrium framework. *Management Sci. 44*(11), 1533–1547.

Blattberg, R. C. and K. J. Wisniewski (1989). Price-induced patterns of competition. *Marketing Sci. 8*(4), 291–309.

Briesch, R. A., W. R. Dillon, and E. J. Fox (2013). Category positioning and store choice: The role of destination categories. *Marketing Sci. 32*(3), 488–509.

Cachon, G. P. and A. G. Kök (2007). Category management and coordination in retail assortment planning in the presence of basket shopping consumers. *Management Sci. 53*(6), 934–951.

Cachon, G. P., C. Terwiesch, and Y. Xu (2005). Retail assortment planning in the presence of consumer search. *Manufacturing Service Oper. Management 7*(4), 330–346.

Cachon, G. P., C. Terwiesch, and Y. Xu (2008). On the effects of consumer search and firm entry in a multiproduct competitive market. *Marketing Sci. 27*(3), 461–473.

Chaneton, J. M. and G. Vulcano (2011). Computing bid prices for revenue management under customer choice behavior. *Manufacturing Service Oper. Management 13*(4), 452–470.

Davis, J. M., G. Gallego, and H. Topaloğlu (2012). Assortment optimization under variants of the nested logit model. Working paper, Cornell University, Ithaca, NY.

Davis, J. M., G. Gallego, and H. Topaloğlu (2013). Assortment planning under the multinomial logit model with totally unimodular constraint structures. Working paper, Cornell University, Ithaca, NY.

Duffin, R. J. and P. Whidden (1961). An exponomial extrapolator. *Journal of Mathematical Analysis and Applications 3*(3), 526–536.

Farias, V. F., S. Jagabathula, and D. Shah (2011). Assortment optimization under general choice. Working paper.

Gallego, G. and R. Wang (2012). Multi-product price optimization and competition under the nested logit model with product-differentiated price sensitivities. Working paper, Columbia University, New York, NY.

Garbade, K. D. and W. L. Silber (1976). Price dispersion in the government securities market. *Journal of Political Economy 84*(4), 721–740.

Guadagni, P. M. and J. D. C. Little (1983). A logit model of brand choice calibrated on scanner data. *Marketing Sci. 2*(3), 203–238.

Hotelling, H. (1929). Stability in competition. *The Economic Journal 39*(153), 41–57.

Hu, W., F. Zhong, and Y. Ding (2006). Actual media reports on gm foods and chinese consumers' willingness to pay for gm soybean oil. *Journal of Agricultural and Resource Economics 31*(2), 376–390.

Kök, A. G., M. L. Fisher, and R. Vaidyanathan (2009). Assortment planning: Review of literature and industry practice. In N. Agrawal and S. A. Smith (Eds.), *Retail Supply Chain Management*, pp. 99–153. New York, NY: Springer.

Koppelman, F. S. and C. Bhat (2006). A self instructing course in mode choice modeling: Multinomial and nested logit models. Technical Report, U.S. Department of Transportation, Federal Transit Administration.

Lancaster, K. J. (1966). A new approach to consumer theory. *Journal of Political Economy 74*(2), 132–157.

Li, G., P. Rusmevichientong, and H. Topaloğlu (2013). The d-level nested logit model: Assortment and price optimization problems. Working paper, University of Southern California, Los Angeles, CA.

Li, H. and W. T. Huh (2011). Pricing multiple products with the multinomial logit and nested logit models: Concavity and implications. *Manufacturing Service Oper. Management 13*(4), 549–563.

Louviere, J. J. and G. Woodworth (1983). Design and analysis of simulated consumer choice or allocation experiments: An approach based on aggregate data. *Journal of Marketing Research 20*(4), 350–367.

Manski, C. and D. McFadden (1981). *Structural Analysis of Discrete Data*. MIT Press, Cambridge, MA.

McFadden, D. (2001). Economic choices. *American Economic Review 91*(3), 351–378.

Norwood, F. B., R. L. Luter, and R. E. Massey (2005). Asymmetric willingness-to-pay distributions for livestock manure. *Journal of Agricultural and Resource Economics 30*(3), 431–448.

Philips, Z., D. K. Whynes, and M. Avis (2006). Testing the construct validity of willingness to pay valuations using objective information about risk and health benefit. *Health Econ. 15*, 195–204.

Rusmevichientong, P., Z.-J. M. Shen, and D. B. Shmoys (2010). Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Oper. Res. 58*(6), 1666–1680.

Rusmevichientong, P. and H. Topaloğlu (2012). Robust assortment optimization in revenue management under the multinomial logit choice model. *Oper. Res. 60*(4), 865–882.

Sinha, A., A. Sahgal, and S. K. Mathur (2013). Category optimizer: A dynamic-assortment, new-product-introduction, mix-optimization, and demand-planning system. *Marketing Sci. 32*(2), 221–228.

Talluri, K. and G. van Ryzin (2004). Revenue management under a general discrete choice model of consumer behavior. *Management Sci. 50*(1), 15–33.

Train, K. E. (2009). *Discrete Choice Methods with Simulation* (2 ed.). Cambridge University Press, New York, NY.

van Ryzin, G. (2013). Models of demand. In O. Ozer and R. Phillips (Eds.), *The Oxford Handbook of Pricing Management*, pp. 340–380 (Chapter 18). Oxford, UK: Oxford University Press.

van Ryzin, G. and S. Mahajan (1999). On the relationship between inventory costs and variety benefits in retail assortments. *Management Sci. 45*(11), 1496–1509.

Villas-Boas, J. M. (2004). Consumer learning, brand loyalty, and competition. *Marketing Sci. 23*(1), 134–145.

Weber, M. D., L. M. Leemis, and R. K. Kincaid (2006). Minimum kolmogorov-smirnov test statistic parameter estimates. *Journal of Statistical Computation and Simulation 76*(3), 195–206.

Yellott, J. I. (1977). The relationship between Luce's choice axiom, Thurstone's theory of comparative judgment, and the double exponential distribution. *Journal of Mathematical Psychology 15*(2), 109–144.