NORTHWESTERN UNIVERSITY

The Dynamics of Demand in Seasonal Goods Industries: An Empirical Analysis

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Marketing

By

Gonca Pinar Soysal

EVANSTON, ILLINOIS

June 2008

ABSTRACT

The Dynamics of Demand in Seasonal Goods Industries:

An Empirical Analysis

Gonca Pinar Soysal

This dissertation develops and estimates a dynamic model of consumer choice behavior in markets for seasonal (short lifecycle) goods where products have a finite selling season, consumer valuations change over time and availability is limited. In these markets, retailers often use dynamic markdown policies in which an initial retail price is announced at the beginning of the season and the price is subsequently marked down as the season progresses. Strategic consumers face a tradeoff between purchasing early in the season when prices are higher but goods are available and purchasing later when prices are lower but the stock-out risk is higher.

My structural model incorporates two features essential to modeling the demand for seasonal goods: change in consumer valuations over a finite season and limited

availability. In this model, heterogeneous consumers have expectations about future prices and availability levels and strategically time their purchases. I estimate the model using aggregate sales and inventory data from a fashion goods retailer.

The results indicate that ignoring the change in consumer valuations over the season or consumers' expectations about future availability can lead to biased demand estimates. I find that demand is very responsive to price changes in the earlier periods but that responsiveness decreases significantly throughout the season. I also find that strategic consumers delay their purchases to take advantage of markdowns and that these strategic delays hurt the retailer's revenues. Retailer revenues facing strategic consumers are 18% lower than they would have been facing myopic consumers. Limited availability on the other hand reduces the extent of strategic delays by motivating consumers to purchase earlier. I find that, the impact of strategic delays on retailer revenues would have been as high as 37% if there were no stock-out risk.

By means of three counterfactual experiments, I show that the highest retailer profits are achieved by offering early and small markdowns. On the other hand, given current markdown percentages, the retailer can improve profits by delaying the markdowns or carrying less stock. Facing later markdowns, less price sensitive consumers accelerate their purchases and buy at a higher price. When the retailer limits the initial stock, however, increased stock-out risk in the later periods forces the customers to buy earlier at higher prices. As long as the reduction in availability is not great, the profit gain from earlier sales can overcome the loss due to the reduction in overall sales.

Acknowledgements

I would like to express my deepest gratitude to my advisor Lakshman Krishnamurthi for his guidance and encouragement through the whole process. I would like to thank my dissertation committee members Eric Anderson, Karsten Hansen and Raquel Bernal for their insightful comments and suggestions.

I am grateful to all the great people who have taught me, been a role model to me and supported me during the doctoral program. I learned a lot from my friends, who I hope will be lifelong colleagues. Federico Rossi, Jong-hee Song, Lei Wang, Manish Tripathi, Sandeep Connor, Kanishka Misra and Huseyin Karaca are only a few of all the special people that made this journey a memorable one.

I dedicate this dissertation to my husband, Metehan and my daughter, Nazli. Metehan provided tremendous support, encouraged me even through tough times and always believed in me. Nazli is the joy of my life. I hope I can inspire her to be always curious about life and have a passion for learning.

Metehan ve Nazli’ ya

Contents

List of Tables 10

List of Figures 11

1. Introduction 12

2. Related Literature 23

3. Model 33

3.1 Overview 33
3.2 The Utility Specification 37
3.3 Time Varying Valuations 39
3.4 Availability 40
Translating Aggregate Inventory Data into a Measure of Retail Distribution 42
3.5 Expectations and Evolution of States 48
3.6 Consumer's Decision Rule and Dynamic Optimization Problem 52

4. Estimation 55

4.1 Overview 55
4.2 Calculation of the Purchase Probabilities 55
4.3 Consumer Heterogeneity and Market Shares 57
4.4 Estimation Strategy 60
4.5 Instruments 63
4.6 Market Size 65
4.7 Identification 66

4.8 Key Modeling Assumptions 67
5. The Data 72
6. Empirical Results 76
6.1 Parameter Estimates 76
6.2 Model Comparison 80
6.3 Price Elasticities of Demand 82
6. 4 Availability Elasticities of Demand 84
7. Counterfactuals 86
7.1 Uniform Single Markdown Policy 87
7.2 Change in the Timing of Markdowns 89
7.3 Change in Availability 90
7.4 Impact of Strategic Consumer Behavior on Retailer Revenue 91
8. Discussion 94
REFERENCES 99
APPENDIX 105
Algorithm Used to Translate Aggregate Inventory Levels to Availability 105
TABLES and FIGURES 107

List of Tables

Table 1: Own & Cross Price Elasticities - Homogeneous Multinomial Logit Model 107

Table 2: Total Revenue and Quantity Sold at Different Relative Price Points.......................... 108

Table 3: Availability Expectations Process Parameter Estimates... 109

Table 4: Demand Estimates for a 2 Segment Specification.. 110

Table 5: Demand Estimates for the First Restricted Model.. 111

Table 6: Demand Estimates for the Second Restricted Model 112

Table 7: Model Comparison 113

Table 8: Price Elasticities Averaged Across SKUs by Period.. 114

Table 9: Availability Elasticities Averaged Across SKUs by Period....................................... 115

Table 10: Revenues from a Uniform Single Markdown Policy 116

Table 11: Effects of a Change in the Timing of Markdowns 117

Table 12: Effects of a Change in the Initial Stock.. 118

List of Figures

Figure 1: Aggregate Inventory to Availability Mapping 119

Figure 2: Histogram of Season Length across SKUs.. 120

Figure 3: Sales and Prices for a Sample SKU.. 121

Figure 4: Change in Base Valuations over Time for a Sample SKU.. 122

Figure 5: Evolution of the Segment Sizes over Time 123

Figure 6: Evolution of the Sales Shares over Time 124

Figure 7: Simulated Sales for a Sample SKU 125

Figure 8: Average Price Elasticities over Time 126

Figure 9: Impact of Strategic Consumer Behav. & Limited Availability on Retailer Revenue 127

1. Introduction

This study specifies and estimates a dynamic model of consumer choice behavior in markets for seasonal (short lifecycle, perishable) goods where products have a finite selling season, consumer valuations change over time and availability is limited. The empirical context for the analysis is the fashion apparel market.

Some examples of seasonal goods are fashion apparel, holiday merchandise and concert and airline tickets. Seasonal goods exhibit unique demand characteristics when compared to everyday staples like consumer packaged goods or durable goods. First, there is a well defined, finite selling horizon; goods are introduced into the market, sold over a (usually short) season and are discontinued. Second, consumer valuations change over the season and result in intertemporal variation in demand. For example, consumers in the market for a swimsuit prefer to buy the suit as soon as the summer starts so that they can get the most use out of the product. Third, goods are mostly unique and hedonic and it is difficult to forecast demand for a specific product.

These demand characteristics coupled with supply side limitations create challenges for the seasonal goods retailer in pricing and inventory management. First, demand uncertainty is high due to the uniqueness of the products and time varying valuations of the customers. Second, replenishment lead times are relatively long in some seasonal goods industries (e.g., fashion apparel), compared to the length of the selling season. This limits the retailer's opportunity to replenish the inventory during the season. Third, the end of season salvage value is very low. So, the seasonal goods retailer faces the challenge of maximizing his profits by ordering a fixed amount of inventory before the selling season and selling this inventory over a finite horizon.

After setting an initial stock level, seasonal goods retailers often resort to intertemporal (dynamic) pricing policies and prices for seasonal goods exhibit substantial variation within the season. In the fashion goods industry for example, it is common practice to employ markdown pricing. Every new product line is introduced at a "retail price" and the price is marked down a number of times until the inventory is cleared or the selling season ends. In the rest of the dissertation, I will discuss the fashion apparel market but my methodology and results apply to any market where the selling horizon is well defined and availability is limited.

Intertemporal pricing can help the retailer in two ways. First, it enables the retailer to segment the market and take advantage of differences in consumers' time preferences and price sensitivities (intertemporal price discrimination). While some consumers may prefer to buy a product early in the season and pay a higher price, others may choose to wait and purchase later in the season at a lower price. Second, it helps the retailer to manage demand uncertainty. If the retailer has overstocked in the initial period for example, he can reduce prices later in the season to boost demand and clear his shelves to make room for the next season's products. Price discrimination and demand uncertainty are the two main reasons behind the intertemporal variation in prices in the seasonal goods industry.

Recent years have witnessed a significant increase in the percentage of retail products sold on sale and in the magnitude of markdowns. The dollar value of total markdowns (on all merchandise sold in department stores), as a percentage of total sales, has increased from 5.2 percent in 1955 to 16.1 percent in 1984 (Pashigian, 1988). Surveys conducted by the National Retail Federation (1998) indicate that in a sample of department and specialty stores, mark-downs as a percentage of sales have risen from 6% in 1967 to 20% by 1998 for department stores, and from 10% to 28% for specialty stores. The same survey also reports that more than 72% of all fashion products sold in 1998 were sold at a discount. Both an explosion in the variety of products offered by the retailers and diversification in consumer tastes are believed to be driving factors behind this increase (Fisher et al., 1994, Pashigian, 1988, Pashigian and Bowen, 1991).

Advances in information technology and marketing research have increased the ability of retailers to collect and analyze consumer data, making it easier for them to employ complex pricing and inventory management strategies. Consumers have been getting increasingly more sophisticated as well. In making a purchase, a strategic consumer weighs the benefits of purchasing today against the benefits of waiting and purchasing in the future at a lower price. As a recent Wall Street Journal (2002) article reports, it is now possible for the consumers to “crack the retailer’s pricing code” and “not pay retail.” Articles in the popular press report that we are witnessing an intensifying cat-and-mouse game between retailers who hope to charge full price for everything and consumers who wait for a sale.

The game between the retailers and the consumers has another interesting dimension in the fashion goods market (and in all markets where availability is limited) in contrast to durable goods markets. As supply is limited, a consumer cannot wait for a sale without taking into account the stock-out risk. So, consumers need to trade off decreasing prices against the possibility of not being able to find the product later in the season. The retailer also faces a tradeoff. Having a lot of stock might increase his chances of meeting the demand but, on the other hand, limiting the stock might motivate strategic consumers to buy earlier at higher prices. For example, Zara, a large Spanish producer and retailer of fashion goods, is well known for its success in implementing a deliberate limited stock strategy. Zara limits the number of clothing products in each store and offers the products over a very short time period to create a sense of urgency among consumers. As a consequence, Zara sells a much higher percentage of the products at retail price compared to their competitors (Ghemawat and Nueno, 2003).

This study develops a structural model of the dynamic decision process on the consumer (demand) side and uses the estimates from this model to investigate retailer pricing and inventory policies. I model strategic consumer behavior where consumers form expectations about future prices and availability and take stock-out risk into account when timing their purchases. I am not aware of prior empirical work that models consumers' expectations about product availability. I also allow for consumers' valuations to vary over the finite season. Previous work on the role of price expectations in consumer choice behavior assumes that consumers have the option of buying the product again next period with certainty if they choose to delay their purchases and that consumers' valuations stay constant over time. Ignoring consumers' availability expectations and time varying valuations would result in biased demand estimates in the seasonal goods context where availability is limited and valuations change over time. The current model also allows for heterogeneity in price and markdown sensitivities, seasonality parameters and time preferences.

This demand model enables a seasonal goods retailer to decompose the effects of different factors that contribute to change in demand over time. Understanding the separate effects of time variation in valuations, increasing stock-out risk and changing market size and composition on demand and market responsiveness is important for the retailer as each of these factors has different implications for the retailer's pricing and inventory strategy. Suppose, for example, the retailer finds out that sales start out strong but slow down significantly in the middle of the season. The retailer would take different actions depending on the reason behind the decline in sales. If the decline is because valuations for both high and low valuation consumers drop very quickly over time, his long-run strategy might be to offer small, early markdowns to boost demand while the consumers are still interested in the product. If on the other hand, the valuations do not decrease significantly but sales slow down because all the high valuation consumers buy and exit the market, the remaining market is composed of only low valuation consumers. In such a case, his long-run strategy might be to offer later, deeper markdowns. This way, he would reduce the high valuation consumers' incentive to wait for the sale and capture the demand from low valuation consumers by lowering the price after the high valuation consumers exit the market.

The structural approach allows me to obtain behavioral predictions that are invariant to the effects of policy changes and makes it possible to simulate various pricing/stocking policies and study their profit implications. A seasonal goods retailer can use my model in jointly determining optimal initial inventory levels, retail prices and the magnitude and depth of markdowns.

My focus is on permanent markdowns as opposed to temporary promotions (e.g., Labor Day sale) where prices are reduced for a limited time and then go back up again. In my model, I assume that there is a separate market for each specific product. Every period (week), each consumer in the market for a specific product decides to buy the product or wait until the next period. Consumers are strategic and heterogeneous with respect to their response parameters (e.g., price sensitivities and time preferences). Strategic consumers have expectations regarding the likelihood of future states like prices and availability and choose to buy the product and exit the market if the expected discounted sum of utilities from buying in that period exceeds that from waiting. One challenge faced when working with aggregate data is that actual stock-outs are not observable unless data on store-level inventory and individual level store visits are available. In this study I offer a simulation based methodology that enables one to translate aggregate inventory data to an availability measure that reflects stock-out risk.

The data used to estimate my model comes from a specialty apparel retailer that sells its own private label fashions. Aggregate weekly sales, on-hand-inventory and cost data are available for 105 SKUs from the women's coats category, for a period of two years. Each SKU is introduced, sold over a finite season and is discontinued. The season length varies within a range of 11 to 31 weeks and the median season length is 19 weeks. There is significant variation in sales and prices across SKUs. Some SKUs sell as few as 600 units while others sell over 25,000 units. Each SKU is marked down at least once during its sales horizon and the median number of markdowns is 3.

To preview my results, I note that the model produces a good fit to the data. My estimates indicate that the market is composed of two distinct consumer segments. The first segment consists of low price sensitivity consumers who account for 79% of the total market. The second segment consists of high price sensitivity consumers who start purchasing late in the season and account for the majority of the end-of-season sales. Base valuations of both segments are time sensitive. Base valuations of the low price sensitivity consumers decrease over time at a faster rate compared to that of the high price sensitivity consumers except for a brief period early in the season. The estimates also show that consumers get extra utility from purchasing at markdown prices.

I compare my model to two alternative models. The first model does not take consumers' availability expectations into account and assumes that consumers expect to find the product in stock with probability 1 every period until the end of the season. The second model assumes consumer base valuations do not change over time within the season. These models produce unreasonable demand estimates (e.g., the price sensitivity estimate is positive in the first model for the high valuation segment) and the DM (Distance Measure Statistic) test suggested by Newey and West (1987) shows that my model outperforms both of these models.

Price elasticities for both segments decrease over time. Demand is very responsive to price changes in the early periods and pricing decisions during these periods are very critical since early markdowns can significantly increase sales but might affect total revenue very negatively if not timed optimally. The market is not very responsive to price changes at the end of the season. I also calculate availability elasticities of demand and show that, in the earlier periods, they are larger than the price elasticities in magnitude. This demonstrates the importance of availability as a strategic variable for the retailer.

The demand estimates indicate that strategic consumers delay their purchases to take advantage of markdowns and that these strategic delays hurt the retailer's revenues. In order to quantify the impact of strategic consumer behavior on retailer revenues, I keep everything else constant, simulate sales and calculate resulting retailer revenues under the assumption that consumers are myopic. I show that retailer revenues facing strategic consumers are 18% lower than they would have been facing myopic consumers. Limited availability, on the other hand, helps to reduce the extent of strategic delays and their impact on retailer's revenues, as increasing stock-out risk over time motivates the consumers to purchase earlier. In order to quantify the extent to which limited availability helps to dampen the impact of strategic consumer behavior on retailer revenues, I keep everything else constant and simulate sales and calculate resulting retailer revenues under the assumption that consumers are strategic but do not face stock-out risk, i.e. they expect the product to be available with probability 1 every period until the end of the season. I show that the impact of strategic delays on retailer revenues would have been much larger, 37%, if there were no stock-out risk and products were available throughout the season.

Given the estimated demand model, I run three counterfactual experiments to investigate different elements of a retailer's markdown and inventory policy: timing and depth of markdowns and level of total stock offered (availability). The first experiment examines the properties of a uniform single markdown policy to gain insights into the retailer's tradeoff between the timing and depth of markdowns. Assuming the retailer sets a certain markdown percentage and timing across all products, I simulate sales and calculate resulting revenues for different combinations of markdown time and depth. Results indicate that early and deep markdowns result in the lowest revenues. Early and small markdowns have the most favorable revenue outcomes. Given a set timing for markdowns, revenue improves by reducing the depth of the markdown but given a specific markdown depth, revenue first improves by delaying the markdowns but starts getting worse after a certain period.

The second experiment examines how keeping the availability and markdown percentages fixed but slightly changing the timing of markdowns impacts the retailer's performance. I find that delaying the first markdowns by one period (compared to the current situation) results in a 4% increase in the retailer's revenue.

The third experiment examines how changing the retailer's stocking policy impacts the retailer's performance. I simulate sales for both segments under 5, 10, 15, 20 and 25% reductions in the initial stock offered. I find that a slight decrease in the initial stock offered can improve retailer's profits, even though reducing availability has a negative effect on the total quantity sold. For example, when the retailer reduces initial stock levels by 5%, total profits increase by 8%.

The rest of the dissertation is organized as follows. Section 2 introduces the related literature. Section 3 presents the model and Section 4 outlines the estimation strategy. Section 5 introduces the data used in the empirical application. Section 6 presents demand estimates. Section 7 presents the counterfactual experiments and discusses pricing and inventory management implications corresponding to the demand estimates. Section 8 concludes with a discussion of the results and future directions for research.

2. Related Literature

This study is closely related to three main streams of literature. The first is the economics literature on intertemporal price discrimination. The second is the operations literature on revenue management. And the third is the recent marketing literature on structural models of strategic consumer behavior. I will briefly discuss each research stream highlighting ideas and articles closely related to my study.

Interest in intertemporal demand considerations in the economics literature first arose in the area of durable goods monopoly pricing. In his seminal paper, Coase (1972) argued that consumer rationality eventually eliminates a durable goods monopolist's market power as consumers foresee price reductions in the future and refuse to pay the monopoly price. Stokey (1981) formalized this result with a rational expectations model of pricing, considering the existence of a perfectly competitive second hand market.

Stokey (1979) and Landsberger and Meilijson (1985) were the first to study a monopolist's intertemporal price discrimination problem. Assuming consumers are strategic and know the future path of prices with certainty, Stokey (1979) showed that differences in consumers' rates of time preference can be a possible reason for declining prices but variation in tastes by itself does not make price discrimination profitable. If all consumers' reservation prices fall at the same proportionate rate, the retailer does not find it profitable to cut prices over time, but if consumers with high valuations are more impatient, it is profitable for the retailer to exploit this impatience by setting a high initial price and lowering the price over time. She also noted that some ignored factors like limited capacity and imperfect insight can lead to declining prices. Landsberger and Meilijson (1985) showed that it might be profitable to price discriminate if the consumers have higher discount rates than the seller even if consumers are homogeneous in time preferences. When the discount rate is large, high valuation consumers choose not to wait. Both of these papers assume full consumer rationality (consumers know with certainty the entire future price policy of the firm and act strategically), but limit the strategies on the supply side to full-commitment pricing policies where the supplier announces the entire price path at the beginning of the season. Full-commitment strategies are usually not sub-game perfect as the retailer might have an incentive to deviate once the initial period passes.

Besanko and Winston (1990) also assume rational (strategic) consumers but characterize a finite horizon, sub-game perfect Nash equilibrium pricing policy. They compare the resulting optimal policy to that of a monopolist facing myopic consumers and show that at any state, prices are lower with rational consumers. This implies that the first-period price with myopic consumers will be higher than the first period price with rational consumers but as the sales paths over time will be different, the prices with myopic consumers might fall below those with rational consumers in future periods. Another interesting finding is that profit loss from assuming that the consumers are myopic could be rather significant for a retailer if, in fact, consumers are strategic. The seller would/might set prices too high in the initial periods, s and/so no rational consumers would buy in these periods. Through a numerical example, Besanko and Winston show that a retailer would more than double his profits following the equilibrium pricing policy if the consumers are in fact rational. In sum, the above theoretical studies on intertemporal price discrimination are very useful for investigating demand side considerations. However, they do not address supply side considerations such as limited supply.

A specific case of intertemporal pricing policies, clearance sales, have also received attention in the economics literature. Lazear (1986) offers a theory of clearance sales in the following context: a risk-neutral retailer will sell a line of dresses that he has already purchased over time but he is uncertain about what the consumers will pay for individual dresses. The retailer benefits from the ability to sell the goods over time in two ways. First, if he cannot sell a specific dress in the first period, he gets a second chance in the following period. Second, the outcome of the first period provides him with additional information about demand where the nature of the information depends on market characteristics, number and attributes of the buyers. In Lazear’s model, the retailer prices a single dress in this specific line and all buyers have the same valuation (reservation price) for the good where the valuation is drawn from a known distribution. With a two-period model, Lazear shows that the level of the initial price and the price drops are positively related to the number of consumers and the uncertainty about the valuations. Also, high variability in valuations results in higher average prices and more goods left unsold. He extends his results to T periods and shows that, as the time horizon increases, the initial price increases, the price drops by smaller amounts and the final price drops to zero. Pashigian (1988) extends Lazear's model to allow for industry equilibrium and provides some empirical evidence from sales offered by department stores. He offers the growing importance of "fashion" (variety) as an explanation for changes in markups and markdowns over time and between merchandise groups. Pashigian and Bowen (1991) provide further empirical evidence for the clearance pricing theory. They consider three hypotheses to explain the observed pricing practices: demand uncertainty, price discrimination and peak-load pricing. Of the three, they conclude that peak-load pricing is the least useful for explaining observed pricing patterns. Because many of the observed pricing patterns can be explained by both uncertainty and the price discrimination, they find it very difficult to distinguish between these two hypotheses.

The second area of interest is the vast operations literature on revenue management that studies the dynamic pricing problem of selling a fixed inventory (capacity) over a short selling season. Interest in this field has stemmed from the use of revenue management (i.e. yield management) by most airlines as well as by many hotels, car rental companies and similar industries. Here, I will discuss studies closely related to my work. For an extensive review, the reader is referred to Elmaghraby and Keskinocak (2002).

Gallego and van Ryzin (1994), Bitran and Mondschein (1997) and Bitran et al. (1998) all study analytical dynamic pricing models of selling a fixed inventory over a fixed selling horizon. Gallego and van Ryzin (1994) model the demand as a Poisson process with an arrival rate that is a one-to-one function of price. They determine the optimal price path as a function of the stock level and the length of the horizon. Bitran and Mondschein (1997) also use a Poisson demand model but in their model the arrival rate depends on time instead of price and purchase rate depends on consumers' reservation prices. They compare the profits from using a more realistic, periodic pricing review policy to those from using a continuous policy, and show that the loss is small when the appropriate number of reviews is chosen. Their model also restricts the number of price changes within a season to a fixed number and shows that profits from using a periodic pricing review policy are very close to those from a continuous policy. Another interesting finding from this study is that uncertainty in demand for new products leads to higher prices, larger discounts and more unsold inventory. Bitran et al. (1998) generalize the models discussed above to a multiple store setting where prices and inventories are coordinated considering reallocations after an initial distribution. These studies derive two structural properties of the optimal policy which state that at any given time the optimal price decreases in the number of items left, and for any given number of items the optimal price decreases over time. Zhao and Zheng (2000) allow demand to vary over time and show that the second property holds if the willingness of a customer to pay a premium for the product does not increase over time. This condition would hold for fashion goods but not for travel services. These papers provide a very extensive treatment of the supply side considerations but they do not incorporate important demand side considerations like consumer heterogeneity and strategic behavior. They also focus principally on the pricing problem and assume the initial inventory level is exogenously determined.

Smith and Achabal (1998) and Mantrala and Rao (2001) consider the joint dynamic pricing and initial inventory level determination problem of the retailer. Smith and Achabal (1998) evaluate a deterministic, continuous demand model where demand is allowed to be a function of the current inventory level as well as price, time and seasonal changes. They argue that sales slow down when the inventory falls below a critical level (“the fixture fill rate”) due to limited shelf space and assortment. They show that initial prices increase and markdowns get deeper as sensitivity of the sales rate to the inventory on-hand increases. Mantrala and Rao (2001) present a decision support system that can be used by retailers to jointly address the problems of optimal markdown pricing and determining the optimal initial inventory levels.

Studies discussed under the revenue management research stream provide extensive analytical models and incorporate many important considerations on the supply side. On the other hand, these studies do not incorporate important demand side considerations like consumer heterogeneity and strategic behavior and the area lacks empirical research with realistic demand models. To my knowledge, Heching et al. (2002) is the only recent empirical study in this area. They estimate a simple demand model using data from a specialty retailer of women's apparel and obtain estimates of revenues under various pricing policies. The analysis suggests that the firm would have increased its revenue if it had had smaller mark-downs earlier in the season.

The third literature stream consists of the recent economics and marketing literature on structural models of strategic consumer behavior. A large number of articles have proposed dynamic models of consumer decision making where there is uncertainty about product quality, future prices, promotions or product introductions. One of the pioneering studies, Erdem and Keane (1996), presents a structural dynamic choice model where "forward looking" consumers are uncertain about attributes of a set of brands, and learn about these brands (update their beliefs about brand attributes) through advertising exposure and usage experience. The model is tested with scanner panel data for laundry detergent purchases. They find that, although statistically the forward-looking model fits the data better, forward-looking and myopic models produce similar parameter estimates and policy implications. Gönül and Srinivasan (1996) consider a model where there is uncertainty about future promotions. They model the impact of consumers' expectations of coupon availability in future periods on their current purchase decisions. Estimating their model with scanner panel data for disposable diaper purchases, they find results consistent with the notion that consumers hold beliefs about future coupons and that these beliefs affect their purchase decisions.

A number of articles have studied dynamic models of consumer decision making where there is uncertainty about future prices. Erdem, Imai and Keane (2003) and Hendel and Nevo (2005) study demand models for frequently purchased goods that are storable and are subject to stochastic price fluctuations. Estimating their model using scanner panel data for ketchup purchases Erdem, Imai and Keane (2003) show that price expectations have important effects on demand elasticities, and long-run cross price elasticities (allowing for the effect of price-cut on future expected prices) are more than twice as large as the short-run cross-price elasticities (holding expectations fixed). Hendel and Nevo (2005) estimate their model using scanner panel data for laundry detergents and report similar findings. They show that static demand estimates overestimate own price elasticities, underestimate cross-price elasticities and overestimate the substitution to the no-purchase option. All these models are constructed for frequently purchased consumer goods and are estimated using scanner panel data. Luan (2005) is another related study and models forward looking consumers' consumption decisions about sequentially released products in the motion picture context.

High-tech durables markets are similar to seasonal goods markets in that prices exhibit a declining pattern over the lifecycle of a product creating an incentive for consumers to delay purchases and repeat purchases are rare. My model is similar to discrete choice models of durable goods adoption developed in the high-tech durables context. Melnikov (2000) models strategic consumers' adoption behavior using data from the computer printer market but does not allow for consumer heterogeneity. Song and Chintagunta (2003) analyze the impact of price expectations on the diffusion patterns of new high-tech products using aggregate data accounting for consumer heterogeneity, but do not allow for econometric errors in the demand function and do not account for price endogeneity. Erdem et al. (2005) investigate how consumers learn about and choose between two different brands of personal computers. Their model requires the use of individual level data and does not allow for availability considerations. Nair (2004) empirically estimates a dynamic structural pricing model in the video game industry using aggregate data in an infinite horizon setting allowing for consumer heterogeneity and econometric error terms but does not allow for availability considerations or time variant valuations.

This study presents the first empirical model that investigates how consumers' availability expectations and the rate of change in valuations over time affect demand, and how the retailer can use limited availability together with markdowns as a tool against strategic consumers to induce sales at higher prices. This study lies on the interface of the recent economics and marketing literature on structural dynamic discrete choice demand models and the vast operations literature on revenue management.

3. Model

3.1 Overview

This study offers a dynamic structural model of demand in markets for seasonal goods where consumers are strategic and heterogeneous. My modeling approach is closer to the adoption models used in the literature for durable goods (e.g., Song and Chintagunta, 2003; Melnikov, 2000; Nair, 2006) than to the demand models used for consumer packaged goods. This is because just as in the durable goods markets, repeat purchase is not a significant source of sales in the seasonal goods industry. Unlike the durable goods adoption models, however, my model allows for change in consumer valuations over a finite horizon and for consumers' consideration of stock-out risk.

In this model, each item is treated as a separate item. An item refers to a stock keeping unit (SKU). Different colors of the same style are treated as different items. Each period, a consumer in the market for a specific SKU decides whether to buy the item and exit the market, or wait until the next period and make the decision again. Consumers are strategic and choose to buy the product if the expected discounted sum of utilities from buying in that period exceeds that from waiting. Consumers also have the option of not buying the item at the end of the season. When calculating expected future utilities, consumers take into account their expectations about future states of the world (e.g., prices, stock-out risk).

My model captures three important characteristics of seasonal goods demand. First, consumers’ responsiveness to prices and other marketing variables changes through the season as a result of the change in consumer valuations and the increase in stock-out risk over the season. Second, because the product is a durable, consumers who purchase the product exit the market and the potential market for a specific item shrinks through the season. Third, the composition of remaining consumers in the market changes through the season as long as there is heterogeneity in the consumer population. As an example, assume that consumers have different price sensitivities and time preferences (i.e., some consumers are impatient as their valuation drops over time at a faster rate as compared to more patient consumers). If consumers face declining prices, less price sensitive and/or impatient consumers, everything else equal, will purchase the product in the earlier periods and exit the market. Thus, the proportion of more price sensitive, more patient consumers in the remaining market will increase over time.

Capturing the effect of variation in consumers’ valuations over the finite season as well as consumers’ consideration of future stock-out risk is important. An empirical regularity in the data is that, except for a brief period early in the season, sales for a specific SKU decline over time at a given price. Sales increase in the periods where prices are marked down but start decreasing immediately after the markdown period. Decreasing consumer valuations and increasing stock-out risk both reduce a consumer's incentive to wait and contribute to the decrease in sales over time. Estimates from a demand model would be biased if one did not account for change in consumer valuations and stock-out risk. Ignoring these effects would result in underestimation of the price sensitivity parameter (since sales and prices are both higher in earlier periods and lower in later periods) and/or overestimation of the markdown sensitivity parameter (since we observe higher sales immediately after a markdown and lower sales in the later periods) and might produce counter-intuitive parameter estimates (e.g., positive price sensitivities).

It is also important to account for strategic consumer behavior and consumer heterogeneity. As discussed above, a number of studies that investigated consumer behavior in the consumer packaged goods (CPG) industry have found that consumers form expectations about product quality (Erdem and Keane, 1996), coupon availability (Gönül and Srinivasan, 1996) and future prices (Erdem, Imai and Keane, 2003). These studies have shown that strategic models fit the data better than myopic models. In the seasonal goods markets, consumers have even higher incentives to behave strategically in timing their purchases as they face significant reduction in prices over a short season and are also subject to stock-out risk. The retailer needs to account for strategic behavior as it affects the shape of the aggregate sales (adoption) curve and induces price dynamics in the market (Song and Chintagunta, 2003). Facing strategic consumers, the retailer needs to take intertemporal substitution into account since a price reduction will influence sales in other periods. Besanko and Winston (1990) showed that the reduction in profit from assuming that the consumers are myopic, when in fact consumers are behaving in a strategic manner, could be rather significant for a retailer. A retailer should also understand and account for heterogeneity in the consumer population. Allowing for heterogeneity provides a flexible pattern for the aggregate sales curve (Song and Chintagunta, 2003). Understanding consumer heterogeneity enables the retailer to take advantage of differences in the population by adjusting prices dynamically through the sales season.

I incorporate consumer heterogeneity through an aggregate analog to the latent class models used with household purchase data (Kamakura and Russel, 1989). I assume that each consumer belongs to one of a finite number of segments and each segment is characterized by its own time preference, price and markdown sensitivity and seasonality parameters.

Finally, my structural approach allows me to obtain behavioral predictions that are invariant to the effects of policy changes and to simulate various pricing/stocking policies and study their profit implications. The interested reader is referred to Chintagunta et al. (2005) for a discussion that compares structural and reduced form modeling approaches.

3.2 The Utility Specification

A consumer ***i***'s conditional indirect utility from purchasing product ***j*** in period ***t*** is defined as:

$$U_{ijt} = \alpha_{ij}(t) + \beta_{ip}\, p_{jt} + \beta_{im}\, d_{jt} + \beta_{is}\, s_t + \xi_{jt} + \varepsilon_{ijt} \qquad \textbf{(1)}$$

Where p_{jt} is the price of product ***j*** in period ***t***, d_{jt} is the markdown dummy and s_t is a seasonal dummy. $\alpha_{ij}(t)$ is consumer ***i***' s time varying base valuation for product ***j***, β_{ip} is sensitivity to price, β_{im} is markdown sensitivity and β_{is} is the seasonality parameter. ξ_{jt} is a product and time specific demand shock and ε_{ijt} is a mean-zero stochastic term.

U_{ijt} is defined as a onetime utility the consumer gets from purchasing the product and includes not only the instantaneous (current period) utility but also the discounted sum of all future utilities the consumer will get from owning this product.

The markdown dummy, d_{jt}, is set to 1 if the product has been (permanently) marked down and is set to 0 if the product is still sold at retail (full) price. This variable is included to capture the "mere markdown" effect, i.e. the possibility that consumers might get extra utility from purchasing on sale. The seasonal dummy, s_t, is included to capture the possibility that utility from a product might be higher (lower) during peak (low) seasonal periods. A close examination of the seasonality patterns in the data reveals a strong demand peak in the six week holiday shopping period that starts after Thanksgiving and ends after Christmas. A regression of overall sales on relative prices and a set of dummies for all possible seasonal periods (e.g., Mother's Day, Labor Day, January-February slow shopping period, holiday shopping period) reveals that the holiday shopping period is the only period that has a significant effect on the overall demand. In my application, s_t is set to 1 for the holiday shopping period and 0 for all other periods.

Note that ξ_{jt} is a product and time specific demand shock and ε_{ijt} is a mean-zero stochastic term. ξ_{jt} controls for any additional product and time specific factors consumers observe that influence the purchase decision but that the econometrician does not observe. In the fashion apparel context, ξ_{jt} corresponds to demand shifters such as a specific SKU appearing in an advertisement or a T.V. show in a specific week. ξ_{jt} also serves as the econometric error term in the estimation of demand. The instantaneous utility from not buying product j in period t is normalized to ε_{ij0t}:

$$U_{ij0t} = \varepsilon_{ij0t}$$

3.3 Time Varying Valuations

I assume consumers' base valuations for a specific product change over the season as a function of time and model changing valuations using the following quadratic form:

$$\alpha_{ij}(t) = a_j + \beta_{it1}\, t + \beta_{it2}\, t^2 \tag{2}$$

As discussed earlier, consumers' valuations are time sensitive in the seasonal goods markets. In the context of my empirical application, winter coats for example, this is for two main reasons. The first reason is the time sensitivity of fashion and the second is the limited seasonal usefulness of the good. For example, if one gets a winter coat early in the season, one is able to wear it when it is in fashion and get more use from it when the weather is still cold. If on the other hand one purchases it at the end of the season, one can use it next year, but one needs to store it and the coat might not be as fashionable (or might not fit) next year. The story is of course different for airline tickets where consumers prefer buying the good later in the season as it gives them a chance to resolve any uncertainty they might have about the trip before the travel. The quadratic specification fits the data well and gives us enough flexibility. I tried alternative specifications and additional terms (logarithmic and cubic) but they did not improve the fit significantly, whereas removing the squared term made the fit much worse.

Notice that I am estimating a separate fixed effect, a_j, for each SKU. This controls for the significant variation in total sales across SKUs. The fixed effect corresponds to the population mean valuation of the (unobserved by the econometrician) product characteristics. The final form of the utility function is as follows:

$$U_{ijt}(S_t) = a_j + \beta_{it1}\,t + \beta_{it2}\,t^2 + \beta_{ip}\,p_{jt} + \beta_{im}\,d_{jt} + \beta_{is}\,s_t + \xi_{jt} + \varepsilon_{ijt} \qquad \textbf{(3)}$$

3.4 Availability

In a limited stock environment, a consumer is likely to face a stock-out in any period. Dana (2001) defines availability as the likelihood a consumer is served. Bruno and Vilcassim (2006) operationalize availability, in a consumer packaged goods context and a multi-store environment, as the probability of finding the product in a store in a given shopping trip. For the purposes of this study I resort to a similar definition and define availability of a specific item in a time period as the probability that a consumer visiting a store in that period finds the item in stock.

It is well accepted that consumers directly value high availability (service rates) and that availability impacts demand. However, accounting for the effect of availability on demand is a challenge. One reason is that availability is often not directly observable. To account for the effect of availability on individual consumer purchase decisions in a multi-store retail environment, one would need real-time data on individual consumer store visits, purchases and real-time inventory data at the store and SKU level. Practitioners and researchers typically have access only to market-level data where sales and inventory information is aggregated across time and/or stores. The data set I use in my empirical application comes from a multi-store retailer. In my data set, sales and inventory information is aggregated across stores and total sales and opening inventory level for each (active) SKU is reported for 104 weeks.

In the absence of detailed real-time data, different measures of *"retail distribution"* have been used as a proxy for availability (Bruno and Vilcassim 2006). Retail distribution is an aggregate measure and is defined, in its simplest form, as *"the number of outlets carrying a product (has the product in stock) as a percentage of total outlets."* Data on retail distribution is available and widely used in the CPG industry. In other industries though, distribution data are typically not available and decision makers need to use more readily available data on aggregate inventory. Although aggregate inventory does not directly reveal retail availability, it is an observable measure of availability. In a multi-store environment, keeping the mechanism that distributes inventory to individual stores fixed, higher levels of aggregate inventory correspond to higher levels of availability. In this study, I propose a method that can be used to translate aggregate inventory data into a measure of retail distribution to serve as a proxy for availability. This method is explained in detail in the next section.

Translating Aggregate Inventory Data into a Measure of Retail Distribution

The method used to translate aggregate inventory data (aggregated over stores) into a retail distribution measure in a multi-store retail environment consists of three main steps. In the first step, given an aggregate inventory level we determine all possible ways inventory can be distributed across stores, i.e., form all feasible distribution vectors. In the second step, we calculate retail distribution corresponding to each feasible distribution vector. In the third step, we calculate expected retail distribution by aggregating retail distribution over all feasible distribution vectors. Below, I explain each step in greater detail and provide a simple example.

<u>Step 1: Distribution</u>

Given an aggregate inventory level and multiple stores, we start by enumerating all possible combinations of store level inventory distributions (distribution vectors) that would be consistent with the aggregate level of inventory. Given 2 stores (A and B) and 2 units of inventory, for example, there are 3 possible distribution vectors. If we define the distribution vector as (i_A, i_B) where i_A represents the inventory level in store A and i_B represents the inventory level in store B, the three possible vectors are d_1=(1,1), d_2=(2,0) and d_3=(0,2). In other words, each store can have 1 unit of inventory; Store A can have 2 units of inventory where Store B has no inventory or vice versa.

More generally, consider at any time t, ***N*** units of inventory distributed to S stores. Define i_s as the inventory level in store s and the vector $d_N = (i_1, i_2, ..., i_S)$ as the store level inventory vector. d_N satisfies two conditions. First, stores can have only positive inventory, i.e., $i_s \geq 0$ s=1,…,S, and, second, store level inventories should sum up to the aggregate inventory level, i.e., $\sum_{s=1}^{S} i_s = N$.

Step 2: Calculation of Retail Distribution

Next, we calculate the corresponding level of retail distribution for each one of the possible vectors. In this example, retail distribution is 1 for d_1 as both stores have the item in stock, and 0.5 for d_2 and d_3 since for these vectors only 1 out of the 2 stores has the item in stock.

More generally, we define $I_s \in \{0,1\}$ as the indicator of the event "item is in stock at store ***s***." Retail distribution, $A(d_N)$, corresponding to the inventory vector d_N is calculated as the number of stores with positive stock as a percentage of total number of stores; i.e., $A(d_N) = \sum_{s=1}^{S} I_s / S$.

Step 3: Aggregation

Finally, we integrate over the probability distribution of store level inventory vectors to calculate the expected value of retail distribution. This last step requires making appropriate assumptions about the distribution of inventory to individual stores and the consumer visit probabilities to individual stores. If, for example, we assume that each unit of inventory is equally likely to end up in any store, distribution d_1 would be twice as likely as d_2 and d_3. If we also assume that consumers are equally likely to visit any store, the resulting retail distribution can be calculated. Define $Av(N)$ as the expected value of aggregate retail distribution (availability) corresponding to an aggregate inventory level of ***N***. Since the aggregate inventory level is 2, we calculate the corresponding expected retail distribution as $Av(2)$ =1/2*1+1/4*0.5+1/4*0.5=0.75. This means that a consumer visiting a random store will find the item in stock with probability 0.75.

More generally, the expected value of aggregate retail distribution (availability) corresponding to an aggregate inventory level of ***N*** is given by $Av(N) = \sum_{d_N} \Pr(d_N) * A(d_N)$ where $\sum_{d_N}$ represents the sum over all feasible distribution vectors d_N and $\Pr(d_N)$ represents the probability of observing the distribution vector d_N.

If we assume that each unit of inventory is independently distributed across the stores and p_s represents the probability of an individual unit of inventory falling in the sth

store, then the distribution of the vector d_N is a multinomial distribution. The probability of exactly i_s units falling into store s is given by:

$$\Pr(d_N) = \Pr(i_1, i_2, ..., i_S) = \begin{cases} \dfrac{n!}{i_1! i_2! ... i_S!} p_1^{\,i1} p_2^{\,i2} ..... p_S^{\,S} & when \ \sum_{s=1}^{S} i_s = N \\ 0 & otherwise \end{cases}$$

When the number of stores and possible inventory levels are small, it is not difficult to enumerate all possible store level inventory vectors, calculate the corresponding distribution levels and multinomial probabilities corresponding to individual inventory vectors and calculate the expected retail distribution. However, this procedure becomes computationally very intensive with many stores and high levels of inventory. This is because the number of all possible inventory vectors grows very rapidly with the number of stores and with the aggregate inventory level and multinomial probabilities get hard to compute.

With over 400 stores and aggregate inventory levels higher than 9000 units for some SKUs and time periods, I need to resort to simulation methods to devise the function that maps aggregate inventory levels to retail distribution (availability). For S stores and an aggregate inventory level ***N***, I simulate the distribution of inventory to stores by making the assumption that each unit of inventory is equally likely to be in any one of the S stores. After repeating this K times, I aggregate over the resulting distribution of distribution vectors to calculate availability (the probability that a customer visiting any one of the stores at random will find the item in stock). A similar approach has been taken in a recent study by Bruno and Vilcassim (2006) in structural demand estimation for the chocolate confectionary industry in the UK. The algorithm is detailed in the Appendix.

In this model, I assume that each item is equally likely to be in any store in order to devise the function that translates aggregate inventory levels to availability. In addition, the retail distribution measure weights distribution across stores equally implying equal consumer visit probability to any store. In my discussions with the retailer I learned that although at the beginning of the season inventory is allocated according to store level sales forecasts, inventory relocation between stores is very costly and very infrequent in later periods. As resulting inventory levels are a function of many stochastic factors affecting store level sales, it is reasonable to assume that later in the season, each remaining item is equally likely to be in any store and use the simple retail distribution measure as a proxy to availability.

If the retailer redistributes inventory to ensure higher levels of inventory in more important (higher-traffic) stores, a more appropriate approach would be to assume different inventory distribution probabilities across stores and use a weighted distribution measure. “All-commodity volume weighted (ACV) distribution” is one such measure widely used in the CPG industry and weights the stores based on their total sales.

As long as the retailer’s re-distribution policy is based on store visit probabilities (weights), the two approaches should produce very similar results. In order to validate the robustness of my inventory measure to the assumptions about distribution and visit probabilities, I re-calculated my availability measure under the assumption that 10% of the stores are high-traffic stores (I assumed that consumers are twice as likely to visit these stores) and the retailer pays special attention to keep stock levels high in these stores (I assumed that each unit of inventory is twice as likely to fall into one of these high-traffic stores). The correlation between the original availability vector and the new availability vector is 0.9985 and the mean of the absolute difference between the two vectors is 0.0121. My availability measure should provide a reasonable approximation to actual availability as long as the inventory distribution is based on consumers’ store visit probabilities.

Before estimating my model, I calculate and tabulate the retail distribution (availability) measure, $Av(N)$, for every observed aggregate inventory level $\boldsymbol{N}$ in my data set. Figure 1 presents the resulting mapping from $\boldsymbol{N}$ to $Av(N)$ constructed using the actual number of stores the retailer has. At the estimation stage, I translate observed aggregate inventory levels into availability using this table. For example, if I observe that

100 units of item ***j*** is available across the retailer's stores in period t, I set the availability of item ***j*** in time ***t***, λ_{jt}, to $Av(100)$.

3.5 Expectations and Evolution of States

I define the state vector, S_t, as the vector of all variables that influence a consumer's purchase decision at time t. A consumer in the market for product j at time t faces the state vector S_t that includes price p_{jt}, markdown status d_{jt}, availability λ_{jt}, seasonality s_{jt}, time and product specific demand shock ξ_{jt} and unobservable error terms.

I assume that the unobservable error terms ($\varepsilon_{ijt}, \varepsilon_{ij0t}$) evolve over time independently from the other state variables. Partitioning S_t into X_t and $\varepsilon_t = (\varepsilon_{ijt}, \varepsilon_{ij0t})$ where X_t represents all state variables except the unobservable error terms, the transition probabilities have the following form:

$$P(S_{t+1} \mid S_t) = P(X_{t+1}, \varepsilon_{t+1} \mid X_t, \varepsilon_t) = P(X_{t+1} \mid X_t) * P(\varepsilon_{t+1} \mid \varepsilon_t)$$

This is the well known "conditional independence" assumption widely used in the literature (Rust, 1994). I further assume that the unobservable error terms are i.i.d. extreme value distributed.

Consistent with the majority of studies in the dynamic choice models literature (Song and Chintagunta (2003), Erdem, Imai and Keane (2003)) I assume that consumers have rational expectations about the future values of state variables. Rational expectations assumptions have been questioned by Manski (2004) because these assumptions may be intrinsically implausible in many contexts and because data on expectations enables one to achieve identification under weaker assumptions. It would be ideal to collect data on individual consumers' expectations and incorporate this information into the dynamic choice model. Erdem et al. (2005) relax the rational expectations assumption using survey data on self-reported consumer price expectations. I do not, however, have access to such rich data on consumer expectations. In the absence of such data, I use the rational expectations assumption to provide a reasonable approximation to consumer expectations.

It is well accepted that consumers rely on past experience, advertising and other signals to predict future states of the world like future prices and availability. In modeling consumers' availability expectations, I assume that a consumer observes the current availability at the store and computes expected future availability relying on past experience. In order to predict future availability, I assume a consumer uses current period availability, time in the season and current price relative to retail price (relative price). In order to approximate the process a consumer uses to compute expected future availability using his past experience, I estimate a linear model that links current availability to past availability, relative price and time in the season using my data set. I specify the following linear model for availability, time and price:

$$\lambda_{j,t+1} = \beta\lambda_{jt} + \gamma(p_{jt}/p_{j0}) + \eta\, time_{j,t} + \theta(p_{jt}/p_{j0})\, time_{jt} + e_{jt} \qquad \textbf{(4)}$$

$$e_{jt} \sim N(0, \sigma_e^2)$$

where λ_{jt} is the observed availability of item j at time t, p_{jt} is the price of item j at time t and p_{j0} is the retail (starting) price for item j and $time_{jt}$ is the number of periods since the beginning of the season for item j at time t. Note that the aggregate weekly inventory information I have for each SKU is translated into availability information for the same SKU and week using the method described in Section 3.4. The availability process parameters are estimated in a first stage using price and availability data and reported in the results section. In the demand estimation stage, I assume that consumers know and use these parameters to form their estimates of availability for future periods. A similar strategy has been employed to model consumers' price expectations in the CPG context by Erdem, Imai and Keane (2003). Note that the specified process is adaptive. In other words, consumers are assumed to observe actual realizations of prices and availability every period and update their availability expectations for future periods.

For tractability and computational ease I assume that consumers can correctly predict future prices and time and product specific demand shocks. In the absence of data on actual consumer expectations, I believe that actual prices serve as a good approximation to consumers' expectations on future prices. Compared to availability information, price information is more readily available, easier to observe and remember. For the retailer in question, prices are uniform across stores and across different distribution channels. It is also easier to predict future prices as markdowns are strategically set by the retailer depending on a product's sales performance in the season whereas availability is an outcome variable that is a function of many stochastic factors. I believe after some experience in buying from the company, most consumers are likely to have a good understanding as to whether a certain product will be a "hot item" or not and when in the season and how big the markdowns will be for that specific product. The researcher, on the other hand, does not observe the "soft" product characteristics related to the appeal of a particular item like its design and fashionability; so these effects are captured through a fixed effect at the demand estimation stage. An alternative approach would be to model price and availability expectations by specifying a joint process. In order to have confidence in my model's predictions of how consumers' price and availability expectations affect choice dynamics, this joint process should be realistic. The disadvantage of not observing the soft product characteristics on the researcher side makes it hard to specify a process that realistically captures the joint evolution of prices and availability. Since the focus of this study is on modeling consumers' availability expectations and how these expectations affect their purchase decisions, I decided to approximate consumers' expectations on prices by their true realized values and specify a process that models consumers' availability expectations.

3.6 Consumer's Decision Rule and Dynamic Optimization Problem

Recall that $U_{ijt}(S_t)$ represents the utility consumer ***i*** gets from purchasing item ***j*** in period ***t*** when the state of the world is S_t and $U_{ij0t}(S_t)$ is the instantaneous utility from the "*no-purchase*" option under the same conditions. Then, the value of buying product ***j*** at time ***t*** for a strategic consumer ***i*** is given by:

$$V_{ijt}(S_t) = U_{ijt}(S_t)$$

The value of the *"no-purchase*" option (waiting) at time ***t*** for a strategic consumer ***i*** is the value from delaying the purchase. The value of the *"no-purchase"* option in period ***t*** is modeled as the sum of (a) the discounted expected value that a consumer can get at time ***t+1*** and (b) the instantaneous utility the consumer can get from the "no-purchase" option. With the "no- purchase option" the consumer gets to choose again next period between purchasing and waiting. Therefore, his expected next period value is the maximum of the value from choosing to wait and the value from choosing to buy. One important point to note here is that the consumer will make this choice only if the product is available next period. If the product is not available, he will get zero utility. So, a strategic consumer ***i*** calculates the expected value from waiting in period t, taking expected availability into account as follows:

$$V_{ij0t}(S_t) = E\left[\lambda_{j,t+1} \max\{V_{ij,t+1}(S_{t+1}), V_{ij0,t+1}(S_{t+1})\} + (1-\lambda_{j,t+1}) \times 0 \mid S_t\right] + \varepsilon_{ij0t}$$

Where $\lambda_{j,t+1}$ is availability in period *t+1*. The expression simplifies as follows:

$$V_{ij0t}(S_t) = E\left[\lambda_{j,t+1} \max\{V_{ij,t+1}(S_{t+1}), V_{ij0,t+1}(S_{t+1})\} \mid S_t\right] + \varepsilon_{ij0t} \qquad \textbf{(5)}$$

The individual consumer's decision rule is such that consumer ***i*** buys item ***j*** and exits the market in period ***t*** only if his value from buying in period ***t*** exceeds his value from waiting and he had chosen to wait in all previous periods:

$$V_{ijt} \geq V_{ij0t} \;\text{ and }\; V_{ij\tau} < V_{ij0\tau} \text{ for all } \tau < t$$

On the other hand, the consumer does not buy product j in period t and stays in the market if his value from waiting in period t exceeds that from buying and he had chosen to wait in all previous periods:

$$V_{ijt} < V_{ij0t} \text{ and } V_{ij\tau} < V_{ij0\tau} \text{ for all } \tau < t$$

4. Estimation

4.1 Overview

This section describes the estimation of the model parameters. I have described the consumer's decision process in the model section. Under distributional assumptions about the stochastic term, ε_{ijt}, I compute the (unconditional) probability that consumer ***i*** buys product ***j*** in period ***t***. Next, I compute the market share by aggregating these probabilities across heterogeneous consumers for each product and time period. And finally, I use the GMM estimation strategy suggested by Berry (1994) to estimate the model parameters. This estimation strategy allows efficient estimation of a large number of parameters as well as dealing with price endogeneity.

I will first discuss the computation of unconditional purchase probabilities.

4.2 Calculation of the Purchase Probabilities

Recall the specification of the value function for the purchase and no-purchase options respectively;

$$V_{ijt} = \alpha_{ij}(t) + \beta_{ip}\, p_{jt} + \beta_{im}\, d_{jt} + \beta_{is}\, s_t + \xi_{jt} + \varepsilon_{ijt} \qquad \textbf{(6)}$$

$$V_{ij0t}(S_t) = E\left[\lambda_{j,t+1} \max\{V_{ij,t+1}(S_{t+1}), V_{ij0,t+1}(S_{t+1})\} \mid S_t\right] + \varepsilon_{ij0t} \qquad \textbf{(7)}$$

The expectation in **(7)** is taken with respect to the distribution of future variables unknown to the consumer conditional on the current information.

Remember that I have assumed that the unobservable error terms $(\varepsilon_{ijt}, \varepsilon_{ij0t})$ evolve independently from the other state variables. I further assume that the unobservable error terms are i.i.d. extreme value distributed.

I define W_{ijt} and W_{ij0t} as the observable (by the retailer) part of the value functions for the purchase and no-purchase options respectively.

$$V_{ij0t} = W_{ij0t} + \varepsilon_{ij0t} \qquad \text{and} \qquad V_{ijt} = W_{ijt} + \varepsilon_{ijt}$$

I can write down W_{ijt} and W_{ij0t} as a function of state variables as follows:

$$W_{ijt}(S_t) = \alpha_{ij}(t) + \beta_{ip}\, p_{jt} + \beta_{im}\, d_{jt} + \beta_{is}\, s_t + \xi_{jt} \tag{8}$$

$$W_{ij0t}(S_t) = E\left[\lambda_{j,t+1} \max\{V_{ij,t+1}(S_{t+1}), V_{ij0,t+1}(S_{t+1})\} \mid S_t\right] \tag{9}$$

Note that in equation **(9)**, following Rust (1987), calculation of the expectation with respect to the distribution of future variables unknown to the consumer can be simplified. The integration with respect to the extreme value errors can be done analytically and $W_{ij0t}(S_t)$ can be expressed by the following equation:

$$W_{ij0t}(S_t) = \int \lambda_{j,t+1} \ln\{\exp[W_{ij,t+1}(S_{t+1})] + \exp[W_{ij0,t+1}(S_{t+1})]\}]dF(S_{t+1} \mid S_t) \qquad \textbf{(10)}$$

Equations **(8)** and **(10)** define W_{ijt} and W_{ij0t}, as a function of state variables, respectively. Since I have assumed that the unobserved error terms in equations **(6)** and **(7)** follow an i.i.d. extreme value distribution, the individual level unconditional purchase probabilities, have the following logit form:

$$P_{ijt} = \frac{\exp(W_{ijt})}{\exp(W_{ij0t}) + \exp(W_{ijt})} \qquad \textbf{(11)}$$

The aggregate purchase probability (market share) for each product j and period t is calculated by integrating P_{ijt} over the consumer heterogeneity distribution. Before specifying the aggregate purchase probabilities, I will discuss my approach in modeling consumer heterogeneity and the evolution of heterogeneity.

4.3 Consumer Heterogeneity and Market Shares

I model consumer heterogeneity using a random coefficients approach. I use a discrete approximation to the parameter distribution and my method is an aggregate analog of the latent-class models widely used for individual level data (Kamakura and Russel 1989). I assume there are K segments in the population and consumers in segment *k* (*k*=1,…,K) share the common parameters β_k where β_k is a vector consisting of the two time preference parameters $(\beta_{kt1}, \beta_{kt2})$, the price sensitivity parameter (β_{kp}), the markdown sensitivity parameter (β_{km}) and the seasonality parameter (β_{ks}); that is, $\beta_k = (\beta_{kt1}, \beta_{kt2}, \beta_{kp}, \beta_{km}, \beta_{ks})$. These parameters are common across all products. The initial size of segment *k*, i.e., proportion of consumers who belong to segment k in the potential market, is represented by π_k and $\sum_{k=1}^{K} \pi_k = 1$. As segments are allowed to be heterogeneous in their time preferences, price and deal sensitivities, they would exhibit different adoption patterns and segment sizes would change over time within the season. Segments with lower price sensitivities and less patient segments would adopt earlier and the proportion of consumers belonging to these segments in the remaining market would fall over time.

Let M_{j0} be equal to the market size for product j, i.e., the number of potential consumers that are in the market for product j. Define N_{jkt} to be the number of remaining consumers from segment k in the market for product j at any period *t*. N_{jkt} is determined by the total market size M_{j0}, segment size π_k and the proportion of consumers from segment *k* who have not bought the item at any period before *t*. Then, the evolution of N_{kjt} in the market is given by:

$$N_{kjt} = N_{kj,t-1}(1-P_{kj,t-1}) \quad \text{or} \quad N_{kjt} = M_{j0}\,\pi_k \prod_{l=1}^{t-1}(1-P_{kjl})$$

If I define θ_{kjt} as the size of segment k in the market for product j at time period t, θ_{kjt} can be calculated as follows:

$$\theta_{kjt} = \frac{N_{kjt}}{\sum_{m=1}^{K} N_{mjt}} = \frac{\pi_k \prod_{l=1}^{t-1}(1-P_{kjl})}{\sum_{m=1}^{K} \pi_m \prod_{l=1}^{t-1}(1-P_{mjl})}$$

Aggregating over the heterogeneity distribution, market share for product j at time t, MS_{jt} can be calculated as follows:

$$MS_{jt} = \sum_{k=1}^{K} \theta_{kjt}\, P_{kjt} = \sum_{k=1}^{K} \theta_{kjt}\, \frac{\exp(W_{kjt})}{\exp(W_{kj0t} + W_{kjt})} \qquad \textbf{(12)}$$

Now that I have explained the calculation of market shares, I will next discuss the strategy employed to estimate the model parameters.

4.4 Estimation Strategy

The model parameters to be estimated consist of product fixed effects represented by the vector *a*, $a=(a_1,......,a_J)$, segment specific parameters represented by the vector β, $\beta=(\beta_1,....,\beta_K)$ where $\beta_k=(\beta_{kt1},\beta_{kt2},\beta_{kp},\beta_{km},\beta_{ks})$ and the initial segment sizes $\pi=(\pi_2,...,\pi_K)$. I resort to the GMM estimation strategy suggested by Berry (1994) for two main reasons. First, the method is computationally efficient when estimating a large number of nonlinear parameters. Second, it allows me to account for the potential endogeneity between product and time specific demand shocks and prices.

Define $X_{jt}=(t,t^2,p_{jt},d_{jt},s_t)$ as the set of covariates. Let $\delta_{jt}=a_j+X_{jt}\beta_1+\xi_{jt}$ denote segment 1's mean utility for product j at time t. Also, let $\overline{\beta}_k=(\beta_k-\beta_1)$ denote segment ***k***' s parameter difference relative to segment 1 for k=2,…,K. Using this notation, I can now rewrite the share equation **(12)** as:

$$MS_{jt}=\theta_{1jt}\frac{\exp(\delta_{jt})}{\exp(W_{1j0t})+\exp(\delta_{jt})}+\sum_{k=2}^{K}\theta_{kjt}\frac{\exp(X_{jt}\overline{\beta}_k+\delta_{jt})}{\exp(W_{kj0t})+\exp(X_{jt}\overline{\beta}_k+\delta_{jt})} \qquad \textbf{(13)}$$

W_{kj0t}, the observable part of the value from waiting for consumer ***k***, product ***j*** and time period ***t***, is a function of ($\delta_{j,t+1},\ldots,\delta_{jT}$), observed covariates and model parameters. Next, I will discuss how one can compute W_{kj0t} starting from period T and working backwards for t=T-1,T-2,…,1 and k=2,…,K. Note that $\overline{\beta}_k = 0$ for segment 1. Remember that using equation **(10)** I can express W_{kj0t} as:

$$W_{kj0,t}(S_t) = \int \lambda_{j,t+1} \ln\{\exp[W_{kj,t+1}(S_{t+1})] + \exp[W_{kj0,t+1}(S_{t+1})]\}]dF(S_{t+1} \mid S_t)$$

Since, conditional on the current information, consumers are uncertain only about the distribution of future availability, integration will be performed over the distribution of future availability given current period availability. I can re-write equation **(10)** as:

$$W_{kj0t} = \int \lambda_{j,t+1} \ln\{\exp[\delta_{j,t+1} + X_{j,t+1}\overline{\beta}_k] + \exp[W_{kj0,t+1}(\lambda_{j,t+1},..,\lambda_{j,T}]\}dF(\lambda_{j,t+1},..,\lambda_{j,T} \mid \lambda_t) \quad \textbf{(14)}$$

The value from waiting is calculated by simulated integration of **(14).** Remember that Equation **(4)** specifies the linear process consumers use to form their expectations about future values of availability given current availability, time and relative price. For each time period t and product j, I draw N random vectors consisting of $(e_{j,t}, e_{j,t+1},....,e_{j,T-1})$. Error terms in these random vectors are i.i.d. normal with mean zero and variance σ_e^2. As discussed in Section 3.5, σ_e^2 is estimated from data together with other coefficients in Equation **(4)** in the first stage. For each random vector, there is a corresponding future availability vector $(\lambda_{j,t+1}, \lambda_{j,t+2},....,\lambda_{j,T})$. W_{kj0t} is calculated by averaging the value of the integrand over N random availability vectors as can be seen in equation **(15).** Value from waiting in any period ***t*** conditional on an availability vector is calculated starting from period T and working backwards. Value from waiting in the last period is normalized to zero, $W_{kj0T} = 0$.

$$W_{kj0t} = \frac{1}{N}\sum_{l=1}^{N}\lambda_{j,t+1}(l)\ \ln\{\exp[\delta_{j,t+1} + X_{jt}\,\overline{\beta}_k] + \exp[W_{kj0,t+1}(\lambda_{j,t+1}(l),..,\lambda_{j,T}(l))]\} \tag{15}$$

Now that I have completed the discussion of how I compute the market shares for each product and each period, next I will discuss the estimation algorithm.

The estimation algorithm can be summarized in three main steps. In the first step, given a value of the unknown parameters, segment 1 mean valuations, δ_{jt}, that equate the observed market shares to the computed market shares are computed by inverting **(13)**. Inversion is made using the contraction-mapping algorithm suggested by Berry, Levinsohn and Pakes (1995) since the function cannot be inverted analytically.

In the second step, the implied demand shocks, ξ_{jt}'s, are calculated using the equality $\delta_{jt} = a_j + X_{jt}\beta_1 + \xi_{jt}$ and interacted with the instruments to form the GMM objective function. The GMM objective function is derived from the sample analog of the moment conditions. Define Z as the instruments matrix and ξ as the stacked vector of all

ξ_{jt}'s. Then the GMM objective function equals $(\xi' Z) W (Z' \xi)$ where W is the appropriate weighting matrix.

In the third step, the parameter values which minimize the objective function are found through a search over all possible parameter values. At the true parameter values Θ^*, the population moment should equal to zero.

My estimation strategy is similar to those used in the recent literature for estimating aggregate discrete choice demand models of differentiated goods. I have included dynamics and allowed for a discrete heterogeneity distribution. Interested readers are referred to Berry, Levinsohn and Pakes (1995), Nevo (2000) and Nevo (2001) for a more detailed discussion of the estimation strategy.

4.5 Instruments

As discussed earlier, one potential strength of the estimation strategy I use is its ability to account for the potential endogeneity between product and time specific demand shocks and prices. Concern for endogeneity arises from the fact that the fashion goods retailer observes the product and time specific demand shocks (ξ_{jt}'s) and takes these demand shocks into consideration in his pricing decisions. As a result, prices and demand shocks could be correlated.

Since I am estimating a fixed effect for each product, these fixed effects capture any product characteristics that do not vary over time periods as well as product specific means of unobserved components. Therefore, I do not need to account for the correlation between prices and the product specific mean of the unobserved product characteristics. The ξ_{jt}'s capture time period specific deviations from the observed mean and I need a set of exogenous instrumental variables to account for the potential endogeneity problem due to the correlation between these deviations and the prices.

I use average men's coats prices from the same period as instruments for women's coats prices. Men's coats prices would be correlated with women's coats prices due to common cost components but would not be correlated with any time specific shock to a specific women's coat SKU, like this particular SKU appearing on a magazine ad for example. I interact men's coats prices with product fixed-effects to make the instruments product specific. In order to assess the strength of these instruments I run a regression with data pooled across all products and time periods. The instruments explain 56% of the variation in women's coats prices ($R^2 = 0.56$, $F(105, 1826) = 22$). The first stage regression of prices on the entire instrument matrix produces an R^2 of 0.92 ($F(212, 1719) = 106$).

4.6 Market Size

In order to estimate demand, I need to have a measure of the initial (potential) market size for each SKU. Knowledge of the initial market size allows me to calculate the observed market share of purchasers (and non-purchasers) using sales data every time period. There are two important considerations in defining the initial market size. First, one should allow for a nonzero share of the outside good (consumers who choose not to purchase at all at the end of the season). Second, one should check the sensitivity of the results to the market definition (Nevo 2000). A retailer would have information on the initial market potential for his products, but the researcher needs to infer the market potential from the data. Previous studies define the market size by choosing a variable to which the market size is proportional, and choosing the proportionality factor. Nevo (2000) calculates the size of the market for ready-to-eat cereal to be one serving of cereal per capita per day. Bresnahan et al. (1997) define the potential market to be the total number of office-based employees in estimating demand for computers. In my application the retailer places the initial order before the season and cannot re-order due to long lead times. Therefore, the size of the order the retailer places for a specific SKU gives us a good idea regarding the retailer’s guess of the potential market size.

I chose the order size as the variable to which the market potential is proportional to and selected a proportionality factor of 1.25 to allow for a nonzero share of the outside good. My demand estimates are not sensitive to the proportionality factor.

4.7 Identification

My identification strategy follows that of BLP. The reader is referred to Nevo (2000) for a detailed discussion of the identification of the random coefficients multinomial logit model in a static setting. Gowrisankaran and Rysman (2007) discuss the identification in a dynamic setting. Identification of my model closely follows that of Nair (2007) which further discusses identification of the binomial logit model in a dynamic setting.

To summarize, what helps me identify the SKU level fixed effects is the variation in mean level of sales across different coats. Price, markdown and seasonality parameters are identified from the within coat variation in these characteristics over the coat's season. The change in market share of product j associated with a change in a characteristic of j (e.g. price) over time identifies the (first segment's) parameter associated with that characteristic (e.g., the price coefficient). Heterogeneity is identified from the deviations from the standard logit implied elasticities (own elasticities in my case since the model is a binary logit model; own and cross price elasticities in the multinomial case). Without heterogeneity, my model will be a standard binomial logit model and standard logit implies own elasticities proportional to the outside good's market share. Consider a 1% change in price from period t to period t+1, holding everything else constant. If my consumer population was homogenous, i.e., I had only one segment, this change would have resulted in the change in market share from period t to period t+1 that is proportional to the outside good's market share in period t. But instead, if I observe a change in market share larger than the change implied by logit elasticity, this would imply the existence of a second consumer segment that is more price sensitive than the first segment and the extent of the deviation from the implied elasticity helps to identify the extent of the difference in two segments' price sensitivities. See Nair (2007) for a similar example on how rate of change in market shares helps identify the relative sizes of the segments. A number of studies have also provided simulation based evidence on identification of heterogeneity from aggregate data with logit demand models in static and dynamic settings (Chintagunta 1999, Song and Chintagunta 2003).

4.8 Key Modeling Assumptions

In this section, I will discuss two key modeling assumptions and address potential concerns related to these assumptions. The first assumption concerns cross-demand effects from products within the same category and second one relates to the impact of competition from retailers selling similar products on retailer's demand.

This study treats each SKU as a separate market and does not consider demand effects across different coats. This assumption is motivated by the characteristics of "fashion" categories. I expect substitution effects to be small in the women's coats category and in all categories on the "fashion" end of the spectrum. In contrast to categories closer to the "staples" end of the spectrum (e.g., men's white dress shirts), fashion categories are associated with greater use of colors, prints and unique designs (Pashigian 1988) and different products within these categories often serve unique tastes. The retailer in my application offers only a few models of women's coats each with a few color options at the same time and models and colors are fairly unique and different from each other. In order to investigate the magnitude of the substitution effects in the data, I estimate a homogenous aggregate multinomial logit demand model that allows for demand effects across coats that are sold at the same time. Explanatory variables include SKU level fixed effects as well as price. Table 1 presents own and cross price elasticities averaged across time periods for a subset of products. The estimated cross-price elasticities are very small (order of magnitude of 10^{-2}) even for different colors of the same model. Products 1 through 4 for example are "almond", "blue opal", "black" and "fez" colors of the same model coat. Own price elasticity averaged across products and time periods on the other hand is -5.36. The estimates suggest that consumers do not perceive different coat SKUs offered for sale at the same time by the retailer to be close substitutes. As the within-category substitution effects are small, my binary logit specification that controls for the substitution between the purchase of the focal SKU and the outside good (i.e. delaying the purchase) should do a good job at estimating demand parameters in this category. Substitutability effects, however, are important in “staple” categories with lower fashion elements (e.g., men’s dress shirts or khaki pants) where alternative products are more substitutable. Similarly, capturing cross-category complementarity effects (e.g., men’s dress shirts and ties) would be important for a seasonal goods retailer when making cross-category pricing decisions. Accounting for these effects on the other hand significantly increases the computational load in a dynamic SKU level model.

My model does also not explicitly account for the effect of competitors’ (other coat sellers’) within-season pricing decisions on demand. One big limitation I face is data availability. While company level data on sales and prices is more readily available, data from multiple competitors is hard to come across not only for academicians but also for practitioners. However, within the context of this study, I expect competitive effects to be small. While retail prices and temporary promotions (e.g., Mother’s Day Sale) are set before the season starts, taking competitors’ strategy into account, decisions regarding the timing and depth of markdowns are operational decisions and competitive reactions to markdown prices within the season are not very likely due to the nature of the decision making process. Markdown decisions are made at the SKU level, dynamically during the short season, taking into account inventory and time left until the end of the season (Bitran and Mondschein, 1997). It is for example very common for different colors of the same model to be marked down at different points in time during the season. Since different products are discounted at different points in time and at different amounts, permanent markdowns are rarely advertised (Smith and Achabal, 1998, Bitran and Mondschein, 1997). Competitive reactions seem very unlikely for unadvertised permanent markdowns (Smith and Achabal, 1998) taken at the SKU level, during the short season especially for fashion products that are fairly unique.

My use of average men's coats prices from the same period as instruments for women's coats prices helps me further address any concerns regarding the endogeneity of prices due to un-captured competitive reactions from retailers with similar products. Let's consider the case where a retailer observes a competitor discounting a similar coat and marks the price of a specific coat SKU down accordingly. One would be concerned about endogeneity of price since the impact of competitor's action on demand would be captured by the unobserved demand shocks and as a consequence prices would be correlated with the unobservable demand shocks. But since the products are marked down at the SKU level, my instrument, average men's coats price from the same period, is very unlikely to respond to the competitor's price cut on a particular women's SKU. As a result, my instruments should be uncorrelated with the unobservable demand shocks due to competitive actions and I don't need to be concerned about endogeneity due to competitive reactions.

5. The Data

The data used in the analysis comes from a specialty retailer that sells its own private label fashion apparel. The retailer has multiple stores throughout the US and the aggregate sales data records weekly sales and starting inventory levels as well as unit acquisition costs at the SKU level. The data cover a two year period (104 weeks), including the years 2003 and 2004.

I estimate the model using data from the "*Women's Coats*" category. I believe this category is suitable for the purposes of this study as the SKU's in this category are sold typically over a season, it is a high involvement category, consumer purchase frequency is low, and repeat purchases from the same consumer especially for the same SKU are unlikely.

I excluded data for about 30 SKUs for which I do not observe the whole sales cycle and ended up with 105 SKUs from this category. I observe each SKU through its season (lifecycle). SKUs were introduced and discontinued at different times during the 2 year observation period. So, different products have different seasons and different season lengths. Figure 2 displays the histogram of season lengths for the 105 SKUs in my sample. The season length varies from 11 to 31 weeks with a median season length of 19 weeks. Each observation corresponds to an SKU-week combination and I have a total of 1932 observations.

One important observation is that there is significant variation across SKUs in total sales and retail price. Total sales range from around 600 units to around 25,000 units and the median is around 6000 units. Retail price ranges from $100 to $350 and the median retail price is $200. In my sample, all SKUs face at least 1 markdown, the maximum number of markdowns is 5, and the median number of markdowns faced by an SKU during its lifecycle is 3. Average first markdown is 38% of the (initial) retail price and the average second markdown is an additional 21% of the retail price. Table 2 summarizes total revenue and quantity sold at relative price points (price as a percentage of the retail price). I can see that for the products in my sample, only 43% of the quantity sold and 57% of the revenue from sales is from sales at full (retail) price. 45% of the quantity sold and 36% of the revenue corresponds to sales that took place when the price relative to retail price was in the range 40 to 80%.

Looking at the prices and sales over time one can easily see important patterns. Figure 3 plots unit sales and prices over time for a sample SKU. Prices are rescaled to protect the identity of the retailer. The SKU is offered for sale at the stores in period 1 at a retail price of 100. Sales start to pick-up over a few periods and then fall down quickly until the first markdown. The first markdown occurs in period 13 and is around 40% and sales make a significant jump at that period. The increase in sales is 500% compared to the period just before the markdown. Following the first markdown, sales fall down even more quickly until the second markdown in period 18. The second markdown is around an extra 20% of the retail price and this time sales increase by 250%. After the second markdown, sales decrease quickly for a few periods and then die slowly. I also observe that sales do not respond much to price changes in the later periods although the response is quite significant in the earlier periods. This is because of reduced valuations, shrinkage in the size of the potential market and limited availability.

I observe significant response to price changes in the early periods but sales drop very quickly over time at a given price. The drop in sales over time at a given price can be the result of decreasing valuations over time (i.e., consumers prefer to purchase earlier than later at a given price) and/or shrinkage in the market size. The spikes on the other hand could be the result of dramatic promotion response, strategic consumer waiting or capturing different segments of customers. Let's concentrate on the period before the first markdown for example. There might be at least two different explanations behind the sales decrease before the first markdown. One explanation might be that strategic consumers are familiar with the retailer's discounting pattern and are delaying their purchases to take advantage of lower prices. This is similar to the pre-promotion dip documented in the marketing literature in the CPG context. Another explanation might be that the retailer faces different consumer segments with different levels of price sensitivities. This first drop in sales might mean that the retailer has only a small segment of low price sensitivity consumers and this segment is saturated early in the season and the retailer needs to lower prices to capture demand from more price sensitive consumers. This emphasizes the importance of a model like mine, since understanding the reasons behind these patterns of observed demand is very important for the retailer's policy. So, the retailer wants to know: Do consumers strategically wait for markdowns? To what extent does strategic waiting explain the observed demand accumulation? And can limiting availability help to dampen the effect of strategic waiting by creating urgency in consumers?

6. Empirical Results

6.1 Parameter Estimates

a. Availability Expectations Process Parameter Estimates

Table 3 reports the OLS estimates of the parameters of the availability expectations process specified in equation **(4)**. The estimates indicate that availability in the next period is closely related to availability in the current period. The relative price parameter is significant and has a positive sign, meaning next period availability is high when the relative price is high, and vice versa, i.e., a markdown in the current period would decrease availability in the next period. Time in season (number of periods since the beginning of the season) is significant and has a negative sign, indicating that controlling for price, availability falls over time. The interaction between the time in season and the relative price is also significant and has a negative sign indicating that later markdowns have a smaller impact on availability compared to earlier markdowns.

b. Demand Model Parameter Estimates

Table 4 reports the GMM estimates of demand model parameters for a two segment specification. Following Besanko et al. (2003), I determine the number of segments by adding segments until one of the segment size parameter estimates is not statistically different from zero. The estimates for the three segment specification yield an insignificant segment size parameter for the third segment. So, the data identifies two segments. As discussed in the estimation section, I estimate the demand parameters for the first segment and deviations of the second segment's parameters from those of the first segment. Segment 2 parameters reported in Table 4 are calculated using these estimates and the standard errors are adjusted accordingly. I also estimate 105 product fixed effects which are not reported here. Product fixed effect estimates lie in the range (-2.6842, 0.4913). My demand estimates reflect a larger, less price sensitive segment, Segment 1, and a smaller, more price sensitive segment, Segment 2. Price sensitivity parameters for both segments have the expected negative sign. Segment 1 corresponds to 79% of the total potential market at the beginning of the season and the estimated price sensitivity parameter for this segment is -0.008. The estimated price sensitivity parameter for Segment 2 is -0.036. Mere markdown effect (markdown sensitivity) is positive and significant for both segment and markdown sensitivity parameter for Segment 1 is slightly larger than that for Segment 2 (0.892 vs. 0.748). Seasonality parameter for Segment 1 is positive (1.198) and significant. This indicates that the fashion sensitive segment gets extra utility from purchases in the 6 week holiday shopping period. Seasonality parameter for Segment 2 on the other hand is negative (-9.429) but not significant.

Time preference parameters for both segments are significant. Base valuations for Segment 1 slightly increase early in the season and decrease rapidly afterwards. Base valuations for Segment 2 on the other hand, decrease, though at a slower rate, and exhibit a slight increase at the end of the season. The time preference parameter estimates for both segments are consistent with intuition. Segment 1 represents a fashion sensitive segment. One would expect fashion sensitive consumers to value the latest fashions and be willing to pay higher prices in order to use the products when they are in fashion, or at the appropriate time in the season. Slight increase in the valuations earlier in the season can be attributed to the fact that apparel items for the upcoming season are usually introduced a few months earlier than the time they are intended to be used. Winter coats for example start to appear at the retail stores starting in August and are used no earlier than November in most regions of the country. Segment 2 on the other hand represents a less fashion sensitive, bargain hunter segment, whose members have a lower willingness to pay for the same items compared to the fashion sensitive segment. Bargain hunters are willing to wait for lower prices later in the season and are known for shopping the clearance sales at the end of the season. Figure 4 represents how estimated base valuations change over time for both segments for a sample SKU. One can see that valuations of the fashion sensitive segment fall rapidly in the season so that after Period

23, base valuations of the bargain hunting segment exceed those of the fashion sensitive segment.

We have seen that the fashion sensitive segment accounts for a significant portion of the potential market at the beginning of the season. As consumers in this segment represent a relatively less price sensitive and less patient group, they make purchases and exit the market early in the season. As a result, relative sizes of the two segments in the remaining potential market change dramatically over time. Averaging across all SKUs, size of the fashion sensitive segment reduces from 79% in Period 1 to 18.5% in Period 25. Figure 5 illustrates the evolution of the sizes of both segments in the remaining market over time (averaged across all SKUs) for 25 periods.

As can be seen in Figure 6, the fashion sensitive segment accounts for almost all sales in the early periods but their share of sales falls down to 52% by period 25 and all the way down to 1.7% by period 30.

In order to understand the purchase behavior of the two segments, I investigate the simulated sales for a sample SKU. Figure 7 represents the sales simulated using the demand estimates from my model, for both segments across time for a sample SKU. Fashion sensitive segment (Segment 1) consumers start purchasing early in the season and some consumers of this segment take advantage of the early markdowns. The bargain hunter segment (Segment 2) consumers on the other hand, start purchasing later in the season and their purchases account for a significant portion of the end-of-season sales. Simulated sales for the two segments show that although Segment 2 customers account for majority of the sales at the end of the period and are important in clearing the shelves of the retailer for the next season, they do not start buying until very late in the season.

6.2 Model Comparison

In order to demonstrate the importance of accounting for consumers' time varying valuations and availability expectations I present estimates from two restricted models and compare these models to my benchmark model. The first model does not account for consumers' availability expectations and the second model does not account for time-varying valuations. My benchmark model on the other hand, takes both of these considerations into account.

In the first model, consumers have expectations about future prices and take these expectations into account when making a purchase decision but they do not take the stock-out risk into account. This model is a restricted version of my original model where $\hat{\lambda}_{j,t+1}$ in equation **(9)** is set to 1 for all products and time periods. Table 5 presents the parameter estimates from this restricted model. This model produces a positive price coefficient for the first segment and underestimates the price sensitivities for both segments. This is because the model ignores the availability effect and attributes consumers' incentive to accelerate purchases (to avoid stock-outs) to lower price sensitivity.

In the second model, consumers have different base valuations for different products but these valuations stay constant throughout the season. This model also is a restricted version of my original model where β_{it1} and β_{it2} in **(2)** are set to zero so that $\alpha_{ij}(t) = a_j$ for all periods. Table 6 presents the parameter estimates from this restricted model. This model overestimates the price sensitivities for both segments and the markdown sensitivity for the second segment. This is because this model ignores the fashion effect and attributes consumers' incentive to accelerate their purchases (to get the most use of the product when it is "in fashion") to higher markdown sensitivity.

In order to compare the restricted models to the benchmark model, I use the DM (Distance Metric) statistic of Newey and West (1987) which is the GMM counterpart of the likelihood ratio test. Table 7 reports the GMM objective function values and DM test results for the three models. The DM test produces test statistics of 5.4 and 38.1 for the no availability and no time sensitivity models respectively, both of which are significant. This indicates that the benchmark model that takes consumers' availability expectations and time sensitivity of valuations into account should be preferred to the two alternative models and is most consistent with the data.

6.3 Price Elasticities of Demand

Table 8 summarizes price elasticities of demand simulated from my model for the first 15 periods. The elasticities were obtained by first simulating the predicted sales using the observed prices, discounting the prices in a specific period by a small amount for each product, simulating the predicted sales once again for the new prices and computing the change in sales in each period relative to the initial values. One important point to note is that in order to isolate the price affects, I hold consumers' availability expectations constant at the levels that correspond to the expectations at observed prices throughout the simulations and exclude the mere markdown effect since the price decrease in my simulations is temporary. Price elasticity presented in the i^{th} row and j^{th} column presents the percent change in sales in period j due to a 1% change in price in period i. A price discount in any period t has three effects on the demand. First, the price discount increases the purchase probabilities in the current period for all customers, thus increasing the demand in the current period. Second, as some consumers choose to accelerate their purchases and purchase in the current period instead of waiting and purchasing in a later period, it decreases the demand in all subsequent periods. Another factor that contributes to the decrease in demand in future periods is the change in the composition of customers. Segment 1 consumers are more likely to accelerate their purchases in the earlier periods as a result of a price decrease and as a result, a smaller proportion of less price sensitive consumers remain in the market. Third, as consumers can foresee a future price discount in my model, some consumers might delay their purchases and wait for the price discount, thus the discount might reduce the demand in periods before the price decrease. In line with these effects, diagonal elements of Table 8 representing the current period elasticities are positive and off-diagonal elements representing the cross period (intertemporal) elasticities are negative. Intertemporal price elasticities for periods closer to the discount period are larger compared to those for periods that are further.

Next, I present the change in average current period elasticities over time. Figure 8 illustrates the changes in the overall current period price elasticity as well as current period price elasticities for both segments throughout the season (averaged across all SKUs). Current period price elasticities for both segments decrease over time. Price elasticity is a function of the current period valuations of the purchase and non-purchase (delay) options and the level of price. The decrease we observe in current period elasticities over time is due to two effects. First, since the prices decrease over the season, a 1% price decrease affects purchase probabilities more significantly in the earlier periods when the prices are higher. Second, the value of the non-purchase option is higher in the earlier periods and this makes consumers facing a price decrease now more likely to accelerate their purchases and buy in the discount period rather than to delay their purchases. Both of these effects contribute to the decrease in current period price elasticities over time. Throughout the season, Segment 1 consumers (fashion sensitive segment) exhibit lower price elasticity compared to Segment 2 (bargain hunter segment) consumers. Since the relative size of the fashion sensitive segment is very large, overall price elasticity is very close to the price elasticity of Segment 1 in the early periods. The overall price elasticity starts increasing later in the season, at about period 20 as the relative size of Segment 2 increases.

Price elasticities of demand indicate that demand becomes less and less responsive to price changes throughout the season. This clearly indicates that correct timing of the markdowns is very critical for the retailer. Earlier markdowns can significantly accelerate demand as demand is very sensitive to price changes in the earlier periods. On the other hand, a higher than necessary early markdown can have important profit implications as markdowns are permanent.

6. 4 Availability Elasticities of Demand

Next, I investigate the effect of consumers' availability expectations on demand. Table 9 summarizes availability elasticities of demand simulated from my model for the first 15 periods. The elasticities were computed by first simulating the predicted sales using the observed prices, decreasing the availability in a specific period by a small amount for each product, simulating the predicted sales once again for the new availability and computing the change in sales in each period relative to the initial values. The elasticity presented in the ith row and jth column represents the percent change in sales in period j due to a 1% change in availability in period i. A decrease in availability in period t does not affect the availability expectations in earlier periods and thus does not affect the sales in earlier periods. This explains the zeros on the lower right half of the availability elasticities matrix. Current period elasticities (diagonal entries) on the other hand are positive since a consumer that observes a decrease in availability expects future period availabilities to be lower as well. This reduces the value from the delay option and accelerates purchases. Since I am simulating a temporary change in availability, availability levels in all future periods are unchanged. As a result, demand in later periods decreases only due to shrinkage in the remaining potential market and change in the consumer mix (as discussed in Section 6.2). An important point to note here is that availability elasticities of demand are rather substantial and much larger than the price elasticities in the earlier periods. This shows the importance of joint consideration of prices and availability levels in retailer's decisions.

7. Counterfactuals

An important strength of a structural demand model is that we can forecast how consumer behavior will change in response to fundamental changes in pricing and inventory management policy. In this section I investigate effects of three such policy changes through policy experiments and conduct a fourth experiment to quantify the impact of strategic consumer behavior on retailer revenues. The first experiment investigates the retailer's tradeoff between the timing and depth of markdowns through a uniform single markdown policy. This experiment shows that the highest retailer profits are achieved by small and early markdowns. The second experiment is aimed at studying the effects of a change in the timing of markdowns while keeping the current markdown percentages constant. This experiment shows that as long as the consumer expectations adjust accordingly, the retailer can improve his revenue and profits by a slight change in the timing of the markdowns, more specifically by offering later markdowns in my case. The third experiment aims to study the effects of a permanent change in the level of the initial stock offered thus affecting the availability throughout the season. In this experiment I show that, counter to intuition, the retailer can actually improve his revenue and profits by supplying less. The fourth experiment aims to quantify the impact of strategic consumer behavior on retailer revenues and determine the extent to which limited availability helps to dampen this impact.

7.1 Uniform Single Markdown Policy

In this experiment I investigate the retailer's tradeoff between the timing and depth of markdowns when setting a uniform single markdown across all products. I focus on a single markdown as the first markdown is the prevalent decision and 82% of all sales at markdown prices take place at the first markdown price for the retailer in my application. In this policy, the seller changes the price only once during the season and sets the same percentage markdown for all products. In order to investigate the sales and revenue impacts of different uniform single markdown policies, I keep the retail (initial) prices and initial inventories fixed and vary the timing and depth of the markdown. I vary the depth of the markdown between 5% and 50% in 5% increments and vary the timing of the markdown between periods 1 and 30 in 1 period increments. For each depth-timing combination, I simulate segment level sales, calculate overall sales and resulting inventory levels and calculate resulting total revenue. Note that the total revenue is the relevant performance metric here since the entire inventory is purchased at the beginning of the season and salvage value is zero. Table 10 summarizes revenue outcomes of each time-depth combination. Due to confidentiality concerns, results are rescaled so that the maximum table entry corresponds to 100. After calculating and normalizing the total revenues, I divide the table entries into 4 regions where Region 1 corresponds to the top (fourth) quartile of all table entries, Region 4 corresponds to the bottom (first) quartile and so on. Each region is represented by a different color. Results indicate that early and deep markdowns (Region 4) result in the lowest revenues. Since the market consists of mostly high valuation, impatient customers, marking down the prices too early and too deep does not have a big impact on sales but the revenue loss due to sales at lower prices is significant. In this region, given a set timing for markdowns, revenues improve by reducing the depth of the markdown and given a specific markdown depth, revenues improve by delaying the markdowns. Late markdowns (Region 3) should be preferred to early and deep markdowns but these markdowns also do not have favorable revenue outcomes. After around Period 23, I do not see much variation in revenue outcomes. This is due to the reduced price responsiveness of the market as a result of reduced valuations. In this region, earlier markdowns are slightly preferred to later markdowns but the depth of the markdown does not affect the revenue outcome. Early and small markdowns (Region 1) have the most favorable revenue outcomes. In this region, given a set timing for markdowns, revenues improve by reducing the depth of the markdown but given a specific markdown depth, revenues first improve by delaying the markdowns (except for the 5% markdown) but start getting worse after a certain period. Since the market is composed of mostly high valuation customers, results indicate the retailer should focus on smaller markdowns but these markdowns should be offered in the earlier periods when market is still responsive to price changes. Smaller markdowns result in higher revenues but if the business conditions and rules (i.e., competitive forces, companywide policies, etc.) indicate a certain markdown percentage, markdown time should be set carefully since marking down too early as well as marking down too late can result in lower revenues.

7.2 Change in the Timing of Markdowns

In this experiment, I investigate whether the retailer can improve his revenues by keeping all other elements of his current pricing and inventory policy the same but slightly changing the timing of markdowns. Timing of markdowns is an important piece of the retailer's pricing and inventory management policy. I study the effects of a change in the timing of markdowns through two scenarios. In the first scenario, I hold everything else constant and simulate the effects of advancing the first markdowns for all products by one period. I set the price for the period before the first markdown to the first markdown price and also allow for consumers' availability expectations to adjust in accordance with the change in sales. The second scenario is similar to the first, except that I delay the first markdowns by one period. Table 11 reports the resulting change in sales and retailer revenue relative to the original policy. Given the current markdown depths, delaying the first markdowns by one period results in a significant increase in the retailer's revenue. Since the bargain hunting segment does not start purchasing until much later in the season, offering later markdowns forces a higher percentage of the fashion sensitive segment customers to purchase earlier at retail price and thus increases retailer's revenue and profits.

7.3 Change in Availability

In this section, I investigate the results of a change in the inventory policy. I have seen that availability elasticities are quite large and the initial stock ordered by the retailer at the beginning of the season is an important part of the retailer's pricing and stocking strategy. I simulate sales for both segments under 5 to 25% reduction in the initial stock offered, varying the reduction in 5% increments. In doing the simulations, I hold the pricing strategy constant and allow for consumers' availability expectations to adjust in accordance with the change in the initial period availability and the changes in sales in all periods. Table 12 reports the resulting change in sales, retailer revenue and retailer profits relative to the original policy. Note that the retailer profit is the relevant measure here since the policy change involves a reduction in the initial stock ordered and changes the cost of acquisition.

Results show that although reducing availability has a negative effect on the total quantity sold, a 5% decrease in the initial stock offered can improve retailer’s profits. A slight decrease in the availability forces the fashion sensitive segment to buy earlier at higher prices and the profit gain from earlier sales overcomes the loss due to the reduction in overall sales. Reducing availability further on the other hand, results in lower profits (compared to a 5% decrease) and a reduction of more than 25 % would result in lower profits compared to the current situation.

7.4 Impact of Strategic Consumer Behavior on Retailer Revenue

In this section, I aim to quantify the impact of strategic consumer behavior on retailer profits and measure the extent to which limited availability helps to dampen this effect.

We have seen that although the fashion sensitive segment accounts for a large portion of the potential market, almost half of the sales take place at markdown prices. Under the current pricing and inventory policy, some fashion sensitive consumers strategically delay their purchases to take advantage of lower price. But how significant are these strategic purchase delays and what is the impact of strategic consumer behavior on retailer revenue?

In order to quantify this effect, I keep the retailer's current price schedule and initial stock levels for all the products, use the demand estimates from my model and simulate sales and resulting revenues assuming consumers are myopic. Myopic consumers maximize immediate utility and purchase when purchase utility exceeds utility from the outside option (which is normalized to 0). Under myopic consumer behavior, we observe earlier sales at higher prices and the resulting total revenue when compared to retailer's current revenue helps me quantify the effect of strategic consumer behavior on retailer' s revenues. As can be seen in Figure 9, under strategic consumer behavior retailer revenue is about 17.8% lower than it would have been under myopic consumer behavior.

Another interesting question is whether limited product availability is helpful in dampening the effect of strategic consumer behavior on retailer revenue. As I have discussed before, limited product availability within the season creates stock-out risk that is increasing over time and reduces the strategic consumers' incentive to delay their purchases and wait for lower prices. In order to quantify the extent to which limited availability (stock-out risk) dampens the impact of strategic behavior on retailer revenue, I once again keep the retailer's initial stock levels and pricing schedule fixed, use the demand estimates and simulate sales under strategic consumer behavior but assume that consumers do not discount future utilities due to stock-out risk. As there is no future stock-out risk, consumers are more likely to wait for lower prices and we observe further delays in purchases and retailer revenues are lower than the current revenues. As summarized in Figure 9, this analysis shows that if the availability was not limited, strategic consumer behavior would have resulted in a 36.5% reduction in retailer revenue and stock-out risk considerably helps to dampen the effect of strategic behavior on retailer revenue.

8. Discussion

In this study, I estimate a dynamic structural model of consumer choice behavior in a market for seasonal goods. My model accounts for two features essential to modeling the demand for seasonal goods: time sensitive consumer valuations and consumers' consideration of stock-out risk. In my model, heterogeneous consumers have expectations about future prices and availability levels and strategically time their purchases. The results indicate that ignoring the time variation in valuations or consumers' consideration of stock-out risk can have strong implications on the demand estimates. I find that a model that ignores the stock-out risk underestimates price sensitivities for both segments whereas a model that ignores the time variation in valuations overestimates price sensitivities for both segments and overestimates markdown sensitivity for the bargain hunter segment.

My analysis shows that the retailer in the empirical application faces a large, less price sensitive segment and a much smaller, more price sensitive segment. Estimates for the time sensitivity parameters indicate that less price sensitive consumers are less patient than the more price sensitive consumers, i.e., their valuations decrease over time very quickly. Although the more price sensitive segment is essential in clearing up the excess inventory at the end of the season, their share of total sales and revenue is quite small.

Calculated price elasticities suggest that demand is very responsive to price changes in the earlier periods but the responsiveness decreases significantly through the end of the season. This finding highlights the big impact of price changes in the earlier periods on the retailer's sales and revenues. Calculated availability elasticities suggest that demand is very sensitive to changes in availability levels in the earlier periods and that availability elasticities are much larger than the price elasticities early in the season.

Through three counterfactual experiments I show that the highest retailer profits are achieved by offering early and small markdowns while early and deep markdowns are very detrimental to retailer profits. Given the current markdown percentages on the other hand, the retailer can improve his profits by delaying the markdowns or carrying less stock. Facing later markdowns, some less price sensitive consumers accelerate their purchases and buy at retail price rather than waiting for the lower price. When the retailer limits the initial stock, a slight decrease in availability increases the stock-out risk in the later periods and forces the customers to buy earlier at higher prices. As long as the reduction in availability is not large, profit gain from earlier sales can overcome the loss due to the reduction in overall sales. I also show that the fact that strategic consumers delay their purchases to take advantage of lower prices results in a 17.8% reduction in retailer's revenue. However, stock-out risk later in the season motivates consumers to purchase earlier at higher prices and I show that if the consumers had not taken stock-out risk into account when timing their purchases, strategic delays would have been more pronounced and the loss in revenue due to strategic behavior would have been twice as large (36.5%).

This study contributes to the current literature on both methodological and substantive grounds. With regard to the methodological contribution, I develop an estimable structural model of strategic consumer choice in the presence of stock-out risk. I also provide a method to estimate the model parameters utilizing aggregate (company level) inventory data. With regard to the substantive contribution, I demonstrate that limited availability and time sensitive valuations can affect the aggregate sales curve, and show that my model can effectively explain interesting regularities in the data like big sales spikes at the markdown periods and rapid decrease in sales over time at a given price. My demand model enables the retailer to understand the different factors resulting in change in demand over time: time varying valuations, reduced availability over time, shrinking potential market and changing consumer mix over time. Accounting for each of these factors separately gives the retailer the opportunity to set optimal initial stock levels and dynamically set optimal prices over the course of the season for different products.

In this study I use three counterfactual experiments to investigate the performance implications of changes in the retailer’s pricing and inventory policy rather than solving the retailer’s optimization problem. One possible extension of this study is formulating and solving the retailer's optimal dynamic pricing and initial inventory level determination problem using the demand model formulated and estimated in this study. A study that investigates the retailer's optimal pricing and initial inventory level problem should account for the retailer's initial demand uncertainty and incorporate retailer demand learning through the season.

Allowing for within and cross category demand effects between products is another important extension of this study. I expect substitution effects to be small in the category I study; women's coats. This category has a high fashion element and products offered to the market at the same time are fairly unique and serve different tastes. In order to analyze the cross price effects I estimate a homogeneous aggregate multinomial logit demand model using sales and prices for the entire set of SKUs. The cross price elasticity estimates from this model are very small (order of magnitude of 10^{-2}) suggesting that within category substitution effects are small. These effects, however are important for a seasonal goods retailer in pricing products from complementary categories (e.g., shirts and ties) or substitute products in categories with lower fashion elements (e.g., men's dress shirts). On the other hand, accounting for these effects at the SKU level in the fashion apparel context brings a computational challenge due to the large number of SKUs simultaneously offered for sale.

I have taken the initial steps in developing a realistic demand model for seasonal goods products accounting for limited availability and time varying valuations as well as strategic consumer behavior and consumer heterogeneity. Future research can benefit from richer data on consumer expectations and availability. I hope that this study encourages further research that resolves some of the open issues raised here.

REFERENCES

[1] Berry, S. (1994). "Estimating Discrete-Choice Models of Product Differentiation," Rand Journal of Economics 25, 242-262.

[2] Berry, S., Levinsohn, J. and Pakes, A. (1995). "Automobile Prices in Market Equilibrium," Econometrica 60(4), 841-890.

[3] Besanko, D., Dube, J.P. and Gupta, S. (2003). "Competitive Price Discrimination Strategies in a Vertical Channel Using Aggregate Retail Data," Management Science 49(9), 1121-1138.

[4] Besanko, D. and Winston, W. (1990). "Optimal Price Skimming by a Monopolist Facing Rational Consumers," Management Science 36(5), 555-567.

[5] Bitran, G.R., Caldentey, R. and Mondschein, S.V. (1998). "Coordinating Clearance Markdown Sales of Seasonal Products in Retail Chains," Operations Research, 46, 609-624.

[6] Bitran, G.R., and Mondschein, S.V. (1997). "Periodic Pricing of Seasonal Products in Retailing," Management Science, 43, 64-79.

[7] Bresnahan, T., Stern, S., and Trajtenberg, M. (1997). "Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the Late 1980's," RAND Journal of Economics, 28, S17–S44.

[8] Bruno, H. and Vilcassim, N.J. (2006). "Structural Estimation with Varying Product Availability," Working Paper, London Business School.

[9] Chintagunta, P., Erdem, T., Rossi, P. and Wedel, M. (2005). ‘‘Structural Modeling in Marketing: Review and Assessment,’’ Marketing Science, 25(6), 604-616.

[10] Coase, R. (1972). “Durability and Monopoly,” Journal of Law and Economics 15, 143-149.

[11] Dana, James. (2001). “Competition in Price and Availability when Availability is Unobservable,” Rand Journal of Economics 32(3), 497-513.

[12] Desiraju, R. and Shugan, S.M. (1999). “Strategic Service Pricing and Yield Management,” Journal of Marketing, 63 (January), 44-56.

[13] Elmaghraby, W. and Keskinocak, P. (2003). “Dynamic Pricing in the Presence of Inventory Considerations: Research Overview, Current Practices and Future Directions,” Management Science, 49(10), 1287-1309.

[14] Erdem, T., Imai, S. and Keane, M.P. (2003). ‘‘Brand and Quantity Choice Dynamics under Price Uncertainty,” Quantitative Marketing and Economics, 1, 5–64.

[15] Erdem, T. and Keane, M.P. (1994). ‘‘Decision-Making under Uncertainty: Capturing Choice Dynamics in Turbulent Consumer Goods Markets’’, Marketing Science, 16(1), 1–21.

[16] Erdem, T., Keane, M.P., Öncü, T.S. and Strebel, J. (2005). “Learning about Computers: An Analysis of Information Search and Technology,” Quantitative Marketing and Economics, 3(3), 207-246.

[17] Feng, Y. and Gallego,G. (1995). “Optimal Starting Times for End-of-season Sales and Optimal Stopping Times for Promotional Fares,” Management Science, 41, 1371-1391.

[18] Fisher, M.L.,Hammond, J.H., Obermeyer, W. and Raman, A. (1994). “Making Supply Meet Demand in an Uncertain World,” Harvard Business Review, 72(3), 83-93.

[19] Gallego, G. and van Ryzin, G.J. (1994). “Optimal Dynamic Pricing of Inventories with Stochastic Demand Over Finite Horizon,” Management Science, 40, 999-1020.

[20] Gonul, F. and Srinivasan, K. (1996). ‘‘Estimating the Impact of Consumer Expectations of Coupons on Purchase Behavior: A Dynamic Structural Model,” Marketing Science 15(3), 262-279.

[21] Ghemawat, P., Nueno, J.L. (2003). “ZARA: Fast Fashion,” Harvard Business School Case, 9-703-497.

[22] Heching, A., Gallego, G. and van Ryzin, G.J. (2002). “Markdown Pricing: An Empirical Analysis of Policies and Revenue Potential at one Apparel Retailer,” Journal of Revenue and Pricing Management , 1(2), 139-160.

[23] Hendel, Igal and Aviv Nevo. (2005). “Measuring the Implications of Sales and Consumer Stockpiling,” NBER Working paper, 11307.

[24] Landsberger, M. and Meilijson, I. (1985). “Inter-temporal Price Discrimination and Sales Strategy Under Incomplete Information”, Rand Journal of Economics 16(3), 424-430.

[25] Lazear, E. (1986). "Retail Pricing and Clearance Sales," American Economic Review 76(1), 14-32.

[26] Luan, J. (2005). "Optimal Inter-Release Time between Sequential Products: The Case of Theatrical Movies and DVDs," Working paper, Dartmouth College.

[27] Manski, C. (2004). "Measuring Expectations," Econometrica, 72 (5), 1329-1376.

[28] Mantrala, M.K. and Rao, S. (2001). "A Decision-Support System that Helps Retailers Decide Order Quantities and Markdowns for Fashion Goods," Interfaces, 31(3), 146-165.

[29] Melnikov, O. (2000). "Demand for Differentiated Durable Products: The Case of the U.S. Computer Printer Market," Working paper, Yale University.

[30] Nair, H. (2004). "Dynamics of Pricing in Durable Good Markets: Application to 32-bit Console Video Games," Working Paper, Stanford University.

[31] National Retail Federation (1998). "The Combined Financial. Merchandising and Operating Results of Retail Stores in 1997"

[32] Nevo, A. (2000). "A Practitioner's Guide to Estimating Random Coefficients Logit Models of Demand," Journal of Economics and Management Strategy 9(4), 513-548.

[33] Nevo, A. (2001). "Measuring Market Power in the Ready-to-Eat Cereal Industry," Econometrica, 69(2), 307-342.

[34] Newey, W. K. and West, K.D. (1987). "Hypothesis Testing with Efficient Method of Moments Estimation," International Economic Review, 28, 777-787.

[35] Pashigian, B.P. (1988). “Demand Uncertainty and Sales: A Study of Fashion and Markdown Pricing,” American Economic Review 78, 936 - 953.

[36] Pashigian, B.P. and Bowen, B. (1991). “Why are Products Sold on Sale? Explanations of Pricing Regularities,” The Quarterly Journal of Economics, 106, 1015-1038.

[37] Rust, J.(1987). “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” Econometrica, 55(5), 999-1033

[38] Rust, J.(1994). “Structural Estimation of Markov Decision Processes,” In Handbook of Econometrics, Vol. IV, edited by R.F. Engle and D.L. Mc Fadden, Elsevier Science B.V.

[39] Smith, S. and Achabal, D. (1998). “Clearance Pricing and Inventory Policies for Retail Chains,” Management Science, 44, 285-300.

[40] Song, I. and Chintagunta, P. (2003). “A Micromodel of New Product Adoption with Heterogeneous and Forwardlooking Consumers: Application to the Digital Camera Category,” Quantitative Marketing and Economics 1(4), 371-407.

[41] Stokey, N. (1979). “Intertemporal Price Discrimination,” Quarterly Journal of Economics 93(3), 355-371.

[42] Stokey, N. (1981). “Rational Expectations and Durable Goods Pricing,” The Bell Journal of Economics, 12(1), 112-128.

[43] The Wall Street Journal (2002). “Cracking the Code: How Not to Pay Retail,” November 27.

[44] Zhao, W. and Zheng, Y. (2000). "Optimal Dynamic Pricing of Perishable Assets with Nonhomogenous Demand," Management Science, 46(3), 375-388.

APPENDIX

Algorithm Used to Translate Aggregate Inventory Levels to Availability

In a multi-store retail environment with S stores, in order to calculate the availability measure corresponding to an aggregate inventory level of n units:

1) Simulate K distribution vectors consistent with an aggregate inventory level of N units

Distribute ***N*** units of inventory to S stores. For each item j=1,…, N, draw a random variable from the discrete uniform distribution with a range from 1 to S to determine which store this item would be located and form the resulting distribution vector. Repeat this step K times to simulate K vectors.

2) Calculate the Retail Distribution for each vector

For each distribution vector, calculate the number of stores with positive stock divided by the total number of stores for an aggregate inventory level of ***N*** and simulation k, as; $A_{Nk} = \sum_{s=1}^{S} I_s / S$ where I_s is the stock-in indicator for store s in distribution vector k.

3) Calculate the Expected Retail Distribution

Average A_{Nk} over K simulations to calculate the availability corresponding to an aggregate inventory level of n : $Av(N) = \sum_{k=1}^{K} A_{Nk} / K$

TABLES and FIGURES

Table 1: Own & Cross Price Elasticities - Homogenous Multinomial Logit Model

		Resulting % Change in Sales of Product														
1 % Price Decrease in Price of Product		**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**
	1	**3.51**	-0.02	-0.03	-0.02	-0.03	-0.03	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01
	2	0.00	**3.29**	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	3	0.00	0.00	**3.52**	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	4	-0.01	-0.01	-0.01	**3.37**	-0.01	-0.01	-0.01	0.00	0.00	0.00	0.00	-0.01	-0.01	0.00	-0.01
	5	-0.01	-0.01	-0.01	0.00	**2.60**	-0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	6	0.00	0.00	0.00	0.00	0.00	**2.47**	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	7	-0.02	-0.02	-0.02	-0.03	-0.03	-0.02	**4.28**	-0.01	-0.01	-0.01	-0.01	-0.02	-0.02	-0.01	-0.02
	8	0.00	0.00	0.00	0.00	0.00	0.00	0.00	**7.79**	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	9	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.02	-0.02	**7.47**	-0.01	-0.02	-0.02	-0.01	-0.02	-0.02
	10	0.00	-0.01	-0.01	-0.01	-0.01	-0.01	-0.02	-0.02	-0.02	**7.68**	-0.02	-0.02	-0.01	-0.02	-0.02
	11	-0.06	-0.05	-0.06	-0.07	-0.07	-0.06	-0.04	-0.04	-0.04	-0.04	**5.83**	-0.04	-0.04	-0.04	-0.04
	12	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	**4.53**	0.00	0.00	0.00
	13	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	**4.95**	0.00	0.00
	14	-0.02	-0.01	-0.02	-0.02	-0.02	-0.02	-0.01	-0.01	-0.01	-0.01	-0.02	-0.01	-0.01	**5.76**	-0.01
	15	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	**4.87**

Table 2: Total Revenue and Quantity Sold at Different Relative Price Points

Relative Price	Revenue	Quantity Sold
100%	57%	43%
80% - 99%	2%	2%
60% - 79%	16%	16%
40% - 59%	20%	29%
20% - 39%	5%	8%
<20%	0%	1%

Table 3: Availability Expectations Process Parameter Estimates

Parameter	Symbol	Estimate	Standard Error
Current Period Availability	β	0.939	0.006
Relative Price	γ	0.051	0.007
Time in Season	η	-0.003	0.000
Time* Rel. Price Interaction	θ	-0.002	0.001
R^2	0.992		
Root MSE	0.058		
Number of Observations	1827		

Table 4: Demand Estimates for a 2 Segment Specification

Parameter Estimates	**Segment 1**		**Segment 2**	
	Estimate	**s.e.**	**Estimate**	**s.e.**
Price Sensitivity	-0.008	0.001	-0.036	0.001
Time Preference (β_{t1})	0.110	0.015	-0.345	0.044
Time Preference (β_{t2})	-0.009	0.001	0.011	0.002
Markdown Sensitivity	0.892	0.090	0.748	0.107
Seasonality	1.198	0.059	-9.429	5.535
Segment Size Parameter ($\ln((1-\theta)/\theta)$)	-1.317	0.000	Size of Segm.1 = 79%	
Objective Function: 254. 4	# of Observations: 1932			

Table 5: Demand Estimates for the Restricted Model that Ignores Stock-out Risk

Parameter Estimates	**Segment 1**		**Segment 2**	
	Estimate	**s.e.**	**Estimate**	**s.e.**
Price Sensitivity	0.005	0.001	-0.010	0.004
Time Preference (β_{t1})	0.038	0.014	-0.193	0.122
Time Preference (β_{t2})	-0.010	0.000	0.006	0.003
Markdown Sensitivity	1.487	0.087	-1.415	1.263
Seasonality	1.466	0.058	-20.233	0.950
Segment Size Parameter (ln((1-θ)/θ)	-0.965	0.084	Size of Segm.1 = 86.3%	
Objective Function: 259.8	# of Observations: 1932			

Table 6:

Demand Estimates for the Restricted Model that Ignores Time Variation in Valuations

Parameter Estimates	Segment 1		Segment 2	
	Estimate	**s.e.**	**Estimate**	**s.e.**
Price Sensitivity	-0.010	0.001	-0.950	0.063
Markdown Sensitivity	0.308	0.090	14.199	0.106
Seasonality	0.718	0.059	-7.999	0.060
Segment Size Parameter (ln((1-θ)/θ)	-1.391		Size of Segm.1 = 80 %	
Objective Function: 292.6	# of Observations: 1932			

Table 7: Model Comparison

Model	Model 1 No Availability	Model 2 No Time Sensitivity	Benchmark Model Avail. & Time Sens.
Objective Value	259.8	292.6	254.4
DM Statistic	5.4	38.2	-
p-value	0.0201	<0.0001	-

DM Statistic is used to compare the two restrictive models against the benchmark model

Table 8: Price Elasticities Averaged Across SKUs by Period

1 % Price Decrease in Period	Resulting % Change in Sales in Period														
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**
1	**1.45**	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.06	-0.06	-0.06	-0.06
2	-0.06	**1.44**	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.06	-0.06
3	-0.05	-0.06	**1.43**	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.07	-0.06	-0.06
4	-0.05	-0.06	-0.06	**1.40**	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08	-0.08	-0.07	-0.07
5	-0.06	-0.06	-0.07	-0.08	**1.19**	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.09	-0.09	-0.09	-0.08
6	-0.05	-0.05	-0.06	-0.07	-0.08	**1.20**	-0.10	-0.10	-0.10	-0.10	-0.10	-0.09	-0.09	-0.09	-0.08
7	-0.04	-0.04	-0.05	-0.05	-0.06	-0.07	**1.20**	-0.11	-0.11	-0.10	-0.10	-0.10	-0.09	-0.09	-0.09
8	-0.03	-0.03	-0.04	-0.04	-0.05	-0.06	-0.08	**1.12**	-0.13	-0.13	-0.13	-0.12	-0.11	-0.11	-0.10
9	-0.03	-0.03	-0.03	-0.04	-0.04	-0.06	-0.08	-0.10	**0.86**	-0.18	-0.17	-0.16	-0.14	-0.13	-0.12
10	-0.01	-0.02	-0.02	-0.02	-0.02	-0.03	-0.04	-0.05	-0.07	**0.79**	-0.15	-0.15	-0.14	-0.14	-0.13
11	-0.01	-0.01	-0.01	-0.01	-0.01	-0.02	-0.02	-0.03	-0.04	-0.05	**0.80**	-0.15	-0.14	-0.13	-0.13
12	0.00	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	-0.02	-0.02	-0.03	-0.05	**0.72**	-0.14	-0.13	-0.13
13	0.00	0.00	0.00	0.00	0.00	-0.01	-0.01	-0.01	-0.01	-0.02	-0.03	-0.04	**0.65**	-0.13	-0.13
14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	-0.01	-0.01	-0.01	-0.01	-0.02	-0.03	**0.62**	-0.12
15	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	-0.01	-0.01	-0.01	-0.02	**0.60**

Table 9: Availability Elasticities Averaged Across SKUs by Period

		Resulting % Change in Sales in Period														
1 % Availability Decrease in Period		**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**
	1	**2.78**	-0.12	-0.12	-0.12	-0.12	-0.12	-0.12	-0.12	-0.12	-0.12	-0.12	-0.11	-0.11	-0.11	-0.11
	2	0	**2.58**	-0.13	-0.13	-0.13	-0.13	-0.13	-0.13	-0.13	-0.13	-0.13	-0.12	-0.12	-0.12	-0.12
	3	0	0	**2.39**	-0.13	-0.13	-0.13	-0.13	-0.13	-0.13	-0.13	-0.13	-0.12	-0.11	-0.11	-0.11
	4	0	0	0	**2.16**	-0.14	-0.14	-0.14	-0.14	-0.14	-0.13	-0.13	-0.13	-0.12	-0.12	-0.12
	5	0	0	0	0	**1.93**	-0.18	-0.18	-0.18	-0.18	-0.18	-0.17	-0.16	-0.14	-0.14	-0.14
	6	0	0	0	0	0	**1.82**	-0.19	-0.19	-0.19	-0.19	-0.18	-0.17	-0.16	-0.16	-0.15
	7	0	0	0	0	0	0	**1.61**	-0.18	-0.18	-0.17	-0.16	-0.16	-0.15	-0.14	-0.14
	8	0	0	0	0	0	0	0	**1.29**	-0.20	-0.20	-0.18	-0.16	-0.15	-0.14	-0.14
	9	0	0	0	0	0	0	0	0	**0.92**	-0.27	-0.26	-0.24	-0.22	-0.21	-0.19
	10	0	0	0	0	0	0	0	0	0	**0.76**	-0.18	-0.18	-0.18	-0.18	-0.16
	11	0	0	0	0	0	0	0	0	0	0	**0.69**	-0.13	-0.13	-0.13	-0.13
	12	0	0	0	0	0	0	0	0	0	0	0	**0.57**	-0.11	-0.11	-0.11
	13	0	0	0	0	0	0	0	0	0	0	0	0	**0.42**	-0.13	-0.13
	14	0	0	0	0	0	0	0	0	0	0	0	0	0	**0.33**	-0.10
	15	0	0	0	0	0	0	0	0	0	0	0	0	0	0	**0.31**

Table 10: Revenues from a Uniform Single Markdown Policy

		Markdown Percentage									
		5%	10%	15%	20%	25%	30%	35%	40%	45%	50%
Time Period	1	100.00	95.36	**90.57**	85.67	80.65	75.54	70.35	65.10	59.78	54.43
	2	99.87	95.39	90.76	85.98	81.10	76.10	71.02	65.86	60.63	55.35
	3	99.74	95.47	91.02	86.43	81.71	76.87	71.93	66.91	61.80	56.62
	4	99.59	95.56	91.36	86.99	82.47	77.83	73.07	68.21	63.25	58.22
	5	99.41	95.67	91.75	87.65	83.39	78.99	74.46	69.82	65.08	60.24
	6	99.23	95.83	92.25	88.47	84.53	80.44	76.20	71.84	67.36	62.78
	7	99.01	96.05	92.89	**89.55**	86.03	82.35	78.52	74.54	70.43	66.20
	8	98.72	96.21	93.52	**90.65**	87.61	84.39	81.02	77.50	73.84	70.05
	9	98.27	96.28	94.12	91.81	89.33	86.69	83.89	80.95	77.87	74.66
	10	97.60	96.12	94.51	92.77	90.89	88.87	86.73	84.45	82.04	79.52
	11	96.83	95.70	94.47	93.13	91.67	**90.11**	88.44	86.66	84.77	82.77
	12	96.11	95.26	94.32	93.30	92.19	90.98	**89.68**	88.29	86.81	85.24
	13	95.44	94.83	94.15	93.40	92.57	91.67	**90.69**	**89.63**	88.49	87.28
	14	94.79	94.37	93.89	93.36	92.77	92.12	91.41	**90.64**	**89.80**	88.91
	15	94.09	93.80	93.47	93.10	92.68	92.23	91.73	91.18	**90.59**	**89.95**
	16	93.43	93.21	92.95	92.67	92.35	92.00	91.61	91.19	**90.73**	**90.24**
	17	92.97	92.80	92.61	92.39	92.14	91.86	91.55	91.22	90.85	**90.45**
	18	92.51	92.40	92.28	92.13	91.97	91.78	91.56	91.32	91.06	90.77
	19	92.03	91.99	91.93	91.85	91.76	91.66	91.54	91.40	91.25	91.08
	20	91.58	91.56	91.53	91.49	91.44	91.39	91.32	91.24	91.15	91.04
	21	91.26	91.24	91.22	91.20	91.16	91.13	91.08	91.03	90.97	90.90
	22	91.05	91.04	91.02	91.00	90.98	90.96	90.92	90.89	90.85	90.80
	23	90.88	90.87	90.86	90.85	90.83	90.82	90.80	90.77	**90.74**	**90.71**
	24	**90.74**	**90.73**	**90.73**	**90.72**	**90.71**	**90.69**	**90.68**	**90.66**	**90.64**	**90.62**
	25	**90.65**	**90.65**	**90.64**	**90.64**	**90.63**	**90.62**	**90.62**	**90.61**	**90.60**	**90.58**
	26	**90.56**	**90.56**	**90.56**	**90.56**	**90.56**	**90.56**	**90.55**	**90.55**	**90.55**	**90.55**
	27	**90.54**	**90.54**	**90.54**	**90.54**	**90.54**	**90.54**	**90.54**	**90.54**	**90.54**	**90.54**
	28	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**
	29	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**
	30	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**	**90.53**

Table 11: Effects of a Change in the Timing of Markdowns

Change Compared to Current Policy	**Change in Total Sales**	**Change in Total Revenues**
Advance the First MD 1 period	0.34%	-3.49%
Delay the First MD 1 period	-0.49%	3.89%

Table 12: Effects of a Change in the Initial Stock

Change Compared to Current Policy	**Change in Total Sales**	**Change in Total Revenues**	**Change in Total Profits**
Limit Initial Stock by 5%	-3.54%	2.32%	8.39%
Limit Initial Stock by 10%	-8.51%	-0.82%	6.80%
Limit Initial Stock by 15%	-13.59%	-4.37%	4.46%
Limit Initial Stock by 20%	-18.67%	-8.25%	1.50%
Limit Initial Stock by 25%	-23.75%	-12.32%	-1.80%

Figure 1: Aggregate Inventory to Availability Mapping

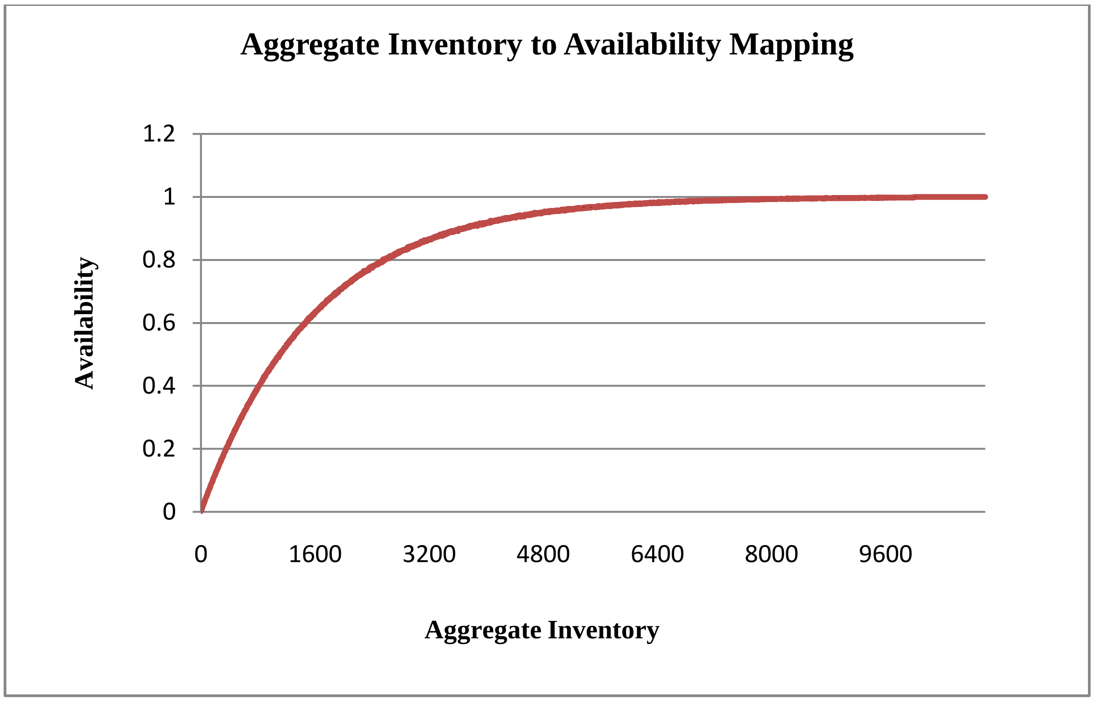

Figure 2: Histogram of Season Length across SKUs

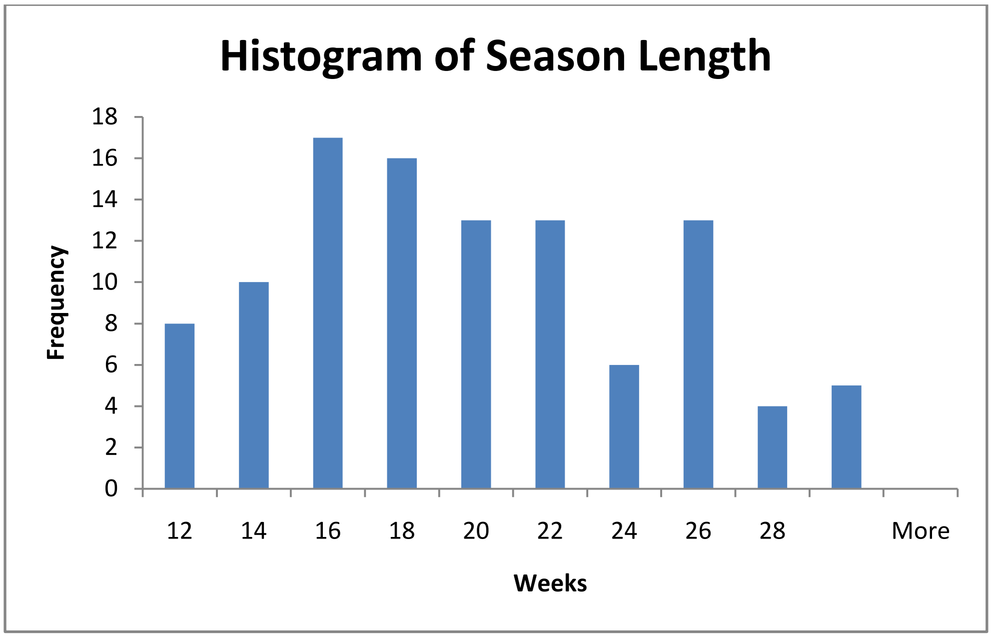

Figure 3: Sales and Prices for a Sample SKU

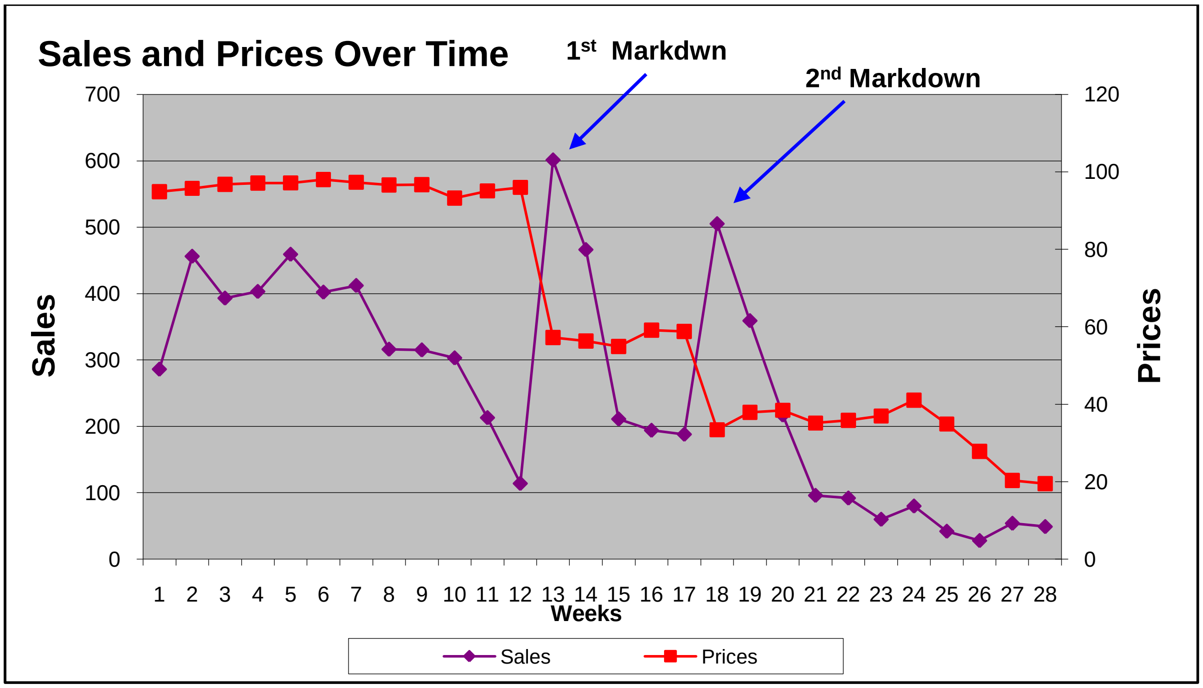

Figure 4: Change in Base Valuations over Time for a Sample SKU

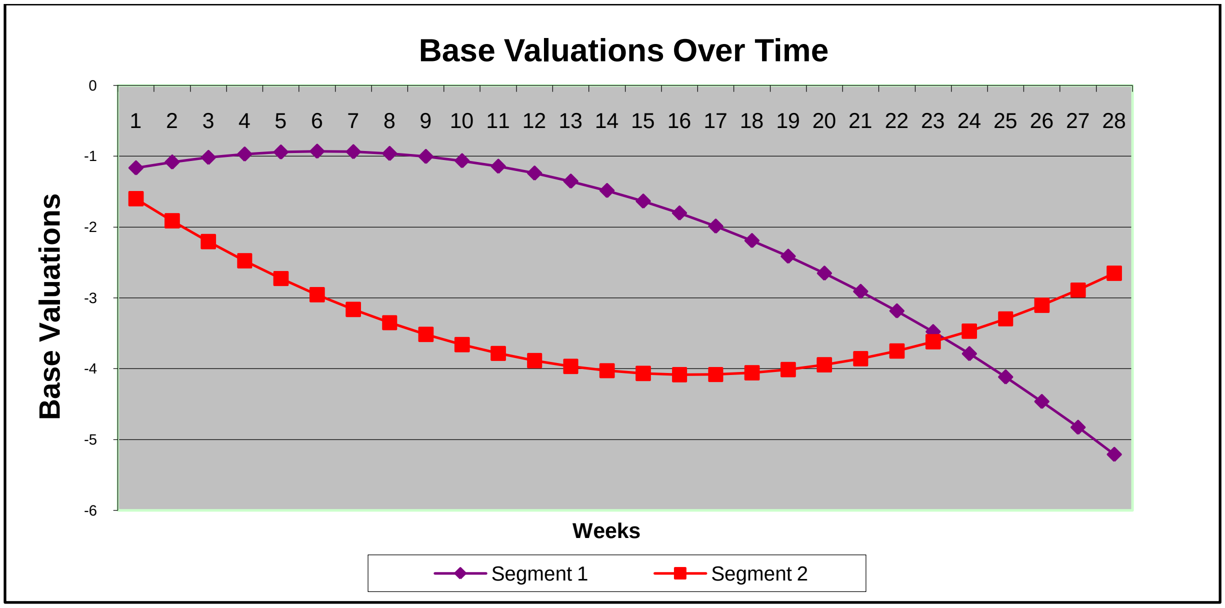

Figure 5: Evolution of the Segment Sizes over Time

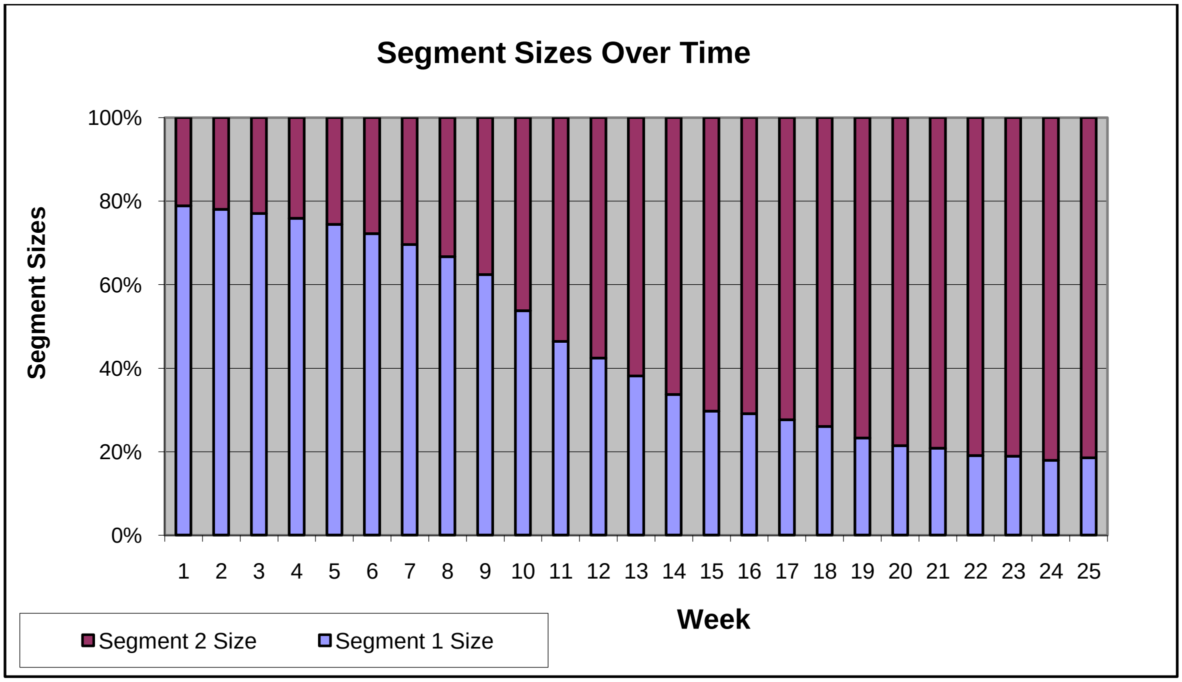

Figure 6: Evolution of the Sales Shares over Time

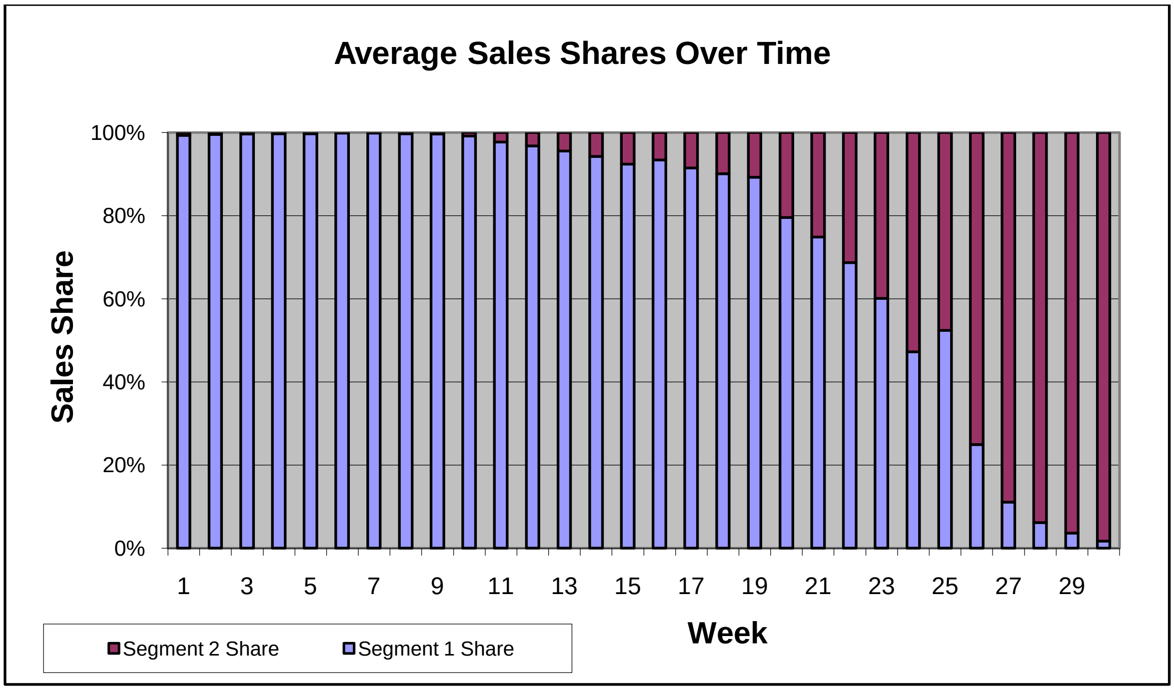

Figure 7: Simulated Sales for a Sample SKU

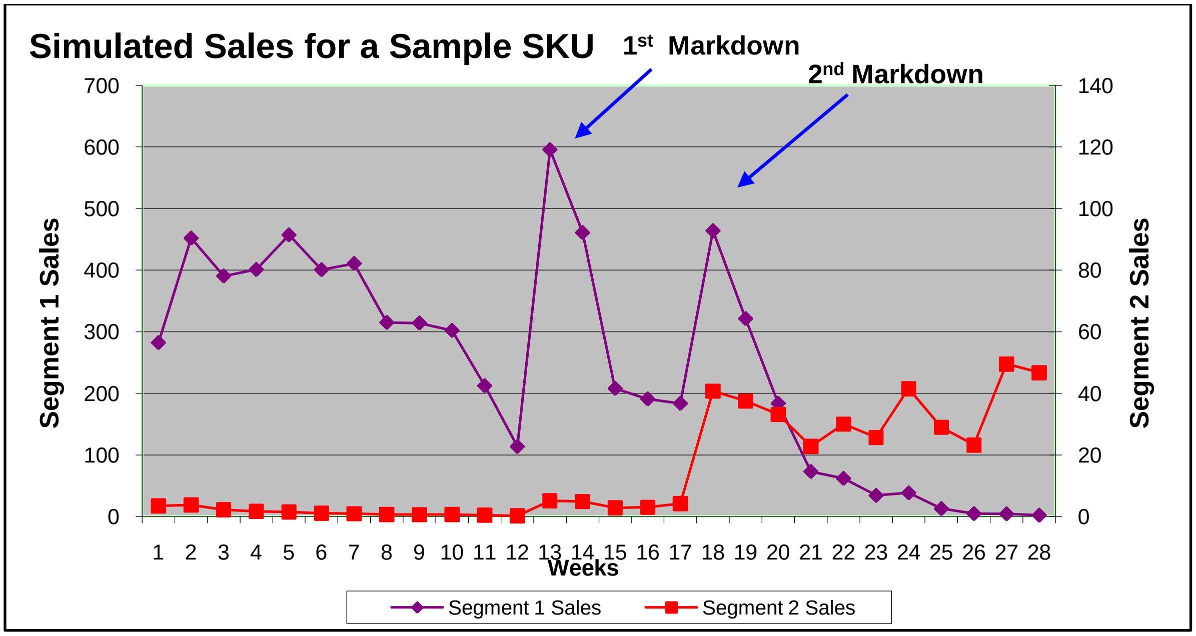

Figure 8: Average Price Elasticities over Time

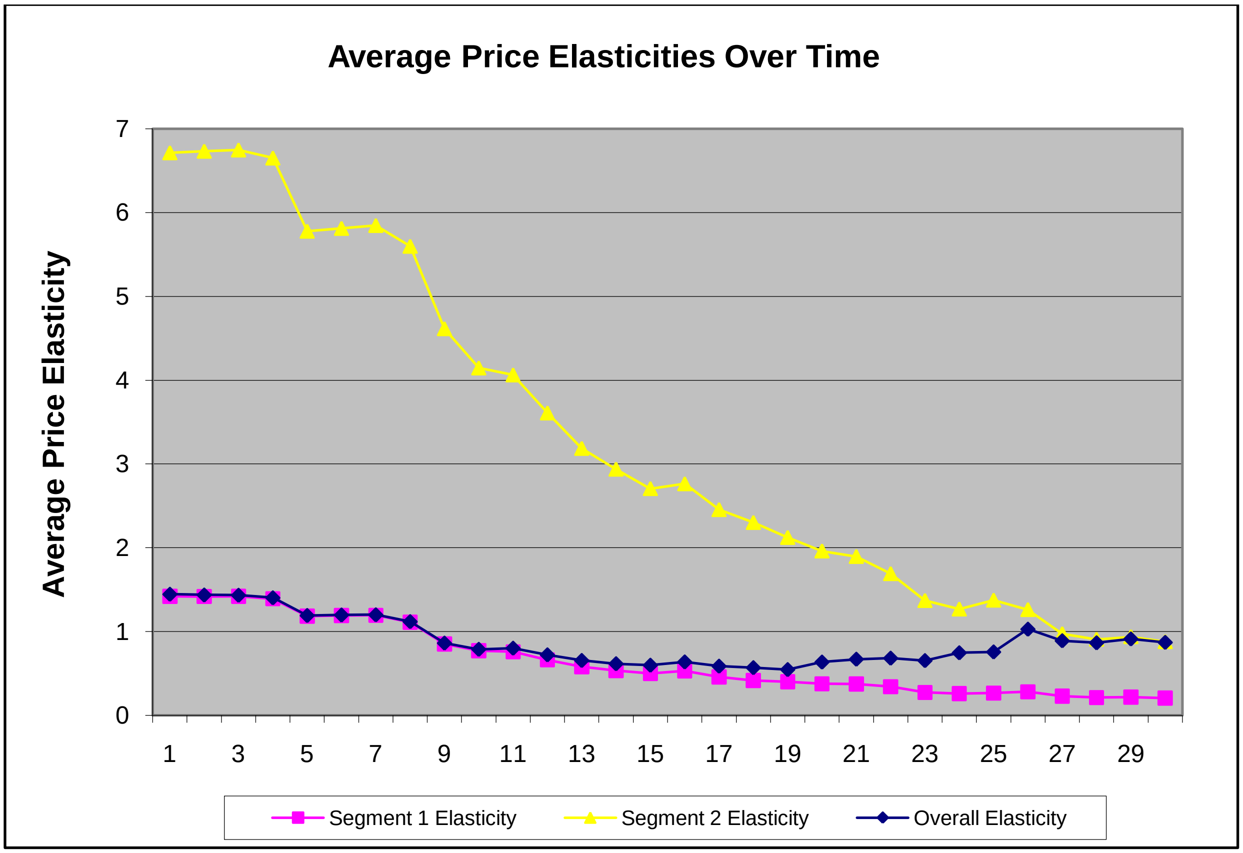

Figure 9: Impact of Strategic Consumer Behavior and Limited Availability on Retailer Revenue

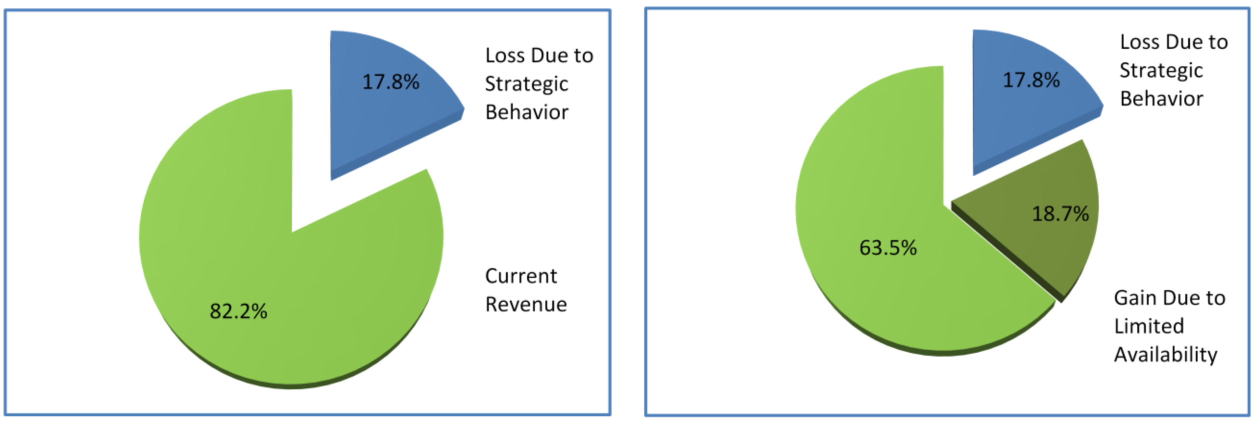

Purchase delays due to strategic behavior hurt the retailer revenues

Purchase acceleration due to stock-out risk dampens the effect of strategic waiting