Title
Essays on Manufacturer Pricing Policies When Retailers and Consumers Stockpile

Author
Qi, Huanhuan

Publication Date
2011

Peer reviewed|Thesis/dissertation

Essays on Manufacturer Pricing Policies When Retailers and Consumers Stockpile

by

Huanhuan Qi

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Engineering - Industrial Engineering and Operations Research

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:
Professor Candace Yano, Chair
Professor Zuo-Jun Max Shen
Professor Xuanming Su

Fall 2011

Abstract

Essays on Manufacturer Pricing Policies When Retailers and Consumers Stockpile

by

Huanhuan Qi

Doctor of Philosophy in Engineering - Industrial Engineering and Operations Research

University of California, Berkeley

Professor Candace Yano, Chair

This dissertation is concerned with pricing issues facing manufacturers when retailers offer periodic discounts and customers stockpile in response. Chapter 1 provides an overview of the dissertation.

In Chapter 2, we study a new Pareto-improving pricing scheme in which the manufacturer subsidizes the retailer's setup (transportation) cost in exchange for a (possibly) higher wholesale price. The retailer responds by choosing regular and discount prices and his order frequency to maximize his revenue less setup, purchasing and inventory holding costs, considering the customers' response. There are two customer segments that differ in their reservation prices and inventory holding costs. Customers make purchasing (including stockpiling) decisions to maximize their utility from consumption less purchasing and inventory holding costs. We characterize the retailer's optimal response to the manufacturer's pricing decisions and the consumers' response to the retailer's pricing schemes. We then show how to solve the manufacturer's decision problem in view of the downstream responses.

In Chapter 3, we investigate the retailer's pass–through of manufacturer trade discounts. The manufacturer offers a fixed wholesale price and periodic trade discounts. The retailer optimizes his ordering plan (including stockpiling when a trade discount is offered) and the pattern of discounts to offer to customers, seeking to maximize revenue less setup, purchasing and inventory holding costs. Customers differ in their reservation prices, and in our model, we account for the adverse effect of retail discounts on consumers reservation prices. For a given frequency and depth of the manufacturer's trade discount, we characterize the retailer's optimal discounting pattern for a given ordering schedule that spans the time between the manufacturer's trade discount offers. We solve for the retailer's jointly optimal ordering and discounting patterns by enumerating appropriate ordering schedules and optimizing the retailer's discount pattern for each.

For the models in Chapters 2 and 3, we also perform associated numerical studies which, together with our analytical results, provide insight into how both manufacturers and retailers should make decisions in these problem settings, and circumstances in which various policies are most effective in increasing profit.

Chapter 4 concludes the dissertation with a summary of contributions and key findings.

Contents

List of Figures

List of Tables

Chapter 1

INTRODUCTION

In a retail supply chain with a manufacturer, retailer and consumers, the manufacturer's pricing policy has complex interactions with many other decisions in the supply chain, such as the retailer's ordering schedule, the retailer's pricing scheme, and consumers' purchasing and stockpiling decisions, all of which ultimately affect the manufacturer's profitability. A single change in the manufacturer's pricing policy may cause complex and indirect downstream reactions, making it difficult for the manufacturer to predict the impact of his pricing changes.

This dissertation addresses two pricing issues that manufacturers may face. The first study focuses on how a manufacturer should couple a transportation subsidy with an adjusted wholesale price to maximize his profit when his own decisions affect the retailer's ordering schedule and pricing decisions in a context where customers are heterogeneous and will stockpile when the retailer offers periodic discounts. The second study focuses on a retailer's choice of ordering schedule, including stockpiling decisions, and discounting patterns, in response to a manufacturer's periodic trade discount in circumstances where consumers stockpile in response to retail discounts and retail discounts reduce consumers' willingness-to-pay. The results of our analysis provide insights into how manufacturers should design trade discounts to be most effective. More detailed descriptions are presented in the following subsections.

1.1 Improving Supply Chain Profitability via Manufacturer Pricing Policies Mitigating the Bullwhip Effect

The bullwhip effect is a well-known phenomenon and has been observed in industry for a long time. Two causes of the bullwhip effect are order batching and high-low pricing. We study manufacturer pricing policies that aim to dampen these behaviors on the part of the retailer so as to reduce the bullwhip effect and its associated costs. We consider a scenario in which transportation scale economies lead the retailer to order in batches, which then gives him an incentive to utilize temporary price discounts (high-low pricing) to clear inventory more rapidly. Customers take advantage of these discounts and may stockpile.

In the literature on order batching as it relates to the bullwhip effect, the emphasis is on smoothing demand by balancing orders across time and on synchronizing orders to improve the opportunity to use consistent information. The main reason for order batching that we consider in this study is the setup cost (fixed transportation cost per order). A few papers in the literature consider setup cost issues as they relate to the bullwhip effect, but to the best of our knowledge, none considers them in a setting where the manufacturer offers a setup cost subsidy to dampen the bullwhip effect by increasing ordering frequency and thereby improve profitability.

We study a pricing scheme in which the manufacturer combines a transportation cost subsidy with an adjusted wholesale price to achieve Pareto-improving outcomes for the manufacturer and retailer. We analyze the problem as a Stackelberg game with the manufacturer moving first, the retailer moving second by setting regular and discount prices as well as his replenishment interval, and finally, heterogeneous customer segments choosing when and how much to purchase given the retailer's prices. We characterize the retailer's optimal policy for any manufacturer pricing policy, and show that the manufacturer needs to consider at most three potentially optimal solutions. We also examine the impact of the optimal pricing policies on the bullwhip effect and profits.

1.2 Retailer's Optimal Pass-Through of Manufacturer's Trade Discounts When Retail Discounts Affect Reservation Prices and Stockpiling

Manufacturers offer trade (wholesale) discounts to increase market share and profitability, but retailers do not necessarily pass the full amount to customers. Empirical research has found various pass-through rates in the market, ranging from as low as about 20% (Besanko et al. 2005) to more than 100%. Researchers have investigated determinants of the retailer's pass-through rate. But to the best of our knowledge, none of the previous studies considers the fact that retail discounts reduce consumers' willingness-to-pay (which we refer to as the (negative) brand equity effect). Some papers on the effects of retail discounts investigate how and to what extent they affect the manufacturer's long-term profitability. We incorporate the brand equity effect in the retailer's optimal pricing decisions.

We model and analyze the effect of retail discounts on customers' reservation prices on the retailer's optimal pass–through of trade discounts when both the retailer and customers may stockpile in response to the discounts. We show that the decline in the customers' willingness to pay as a consequence of retail discounts leads to a threshold effect for the manufacturer's trade discount. More specifically, the decline in the customer's willingness to pay due to retail discounts sometimes makes it necessary for the manufacturer to offer a trade discount exceeding a threshold before the retailer exhibits any response to the discount. This threshold effect does not arise when only stockpiling behavior is considered. We show that the retailer's pass through of the trade discount may be non–monotonic due to the effects of the retailer's and customers' stockpiling. Results from a numerical study enable us to characterize combinations of factors that lead to low or high pass–through rates. The interactions among the factors in influencing the pass–through rate are complex. Our analysis helps to explain phenomena observed in practice, and provides manufacturers insight into the multi-dimensional effect of trade discounts on retailer and consumer behavior, and on profit.

In both studies, we develop new models of complex interactions among a manufacturer, retailer and customers where pricing decisions are involved, where the retailer makes operational decisions (such as ordering and stockpiling), and where common behavioral phenomena among consumers (e.g., stockpiling in response to a discount and/or decreased willingness to pay due to retail discounting). We derive detailed characterizations of each party's optimal decisions and the interactions among the parties' decisions. With these characterizations, we are able to obtain insights, both from the analysis directly and from extensive numerical studies. The insights can provide guidance to manufacturers who face related decisions, and help to provide explanations for phenomena observed in practice.

The first study appears in Chapter 2 and the second study appears in Chapter 3.

Chapter 4 contains conclusions, including a discussion of contributions of the research and a summary of key findings.

Chapter 2

Improving Supply Chain Profitability via Manufacturer Pricing Policies Mitigating the Bullwhip Effect

2.1 Introduction

Our work was motivated by discussions with both manufacturers and retailers regarding the deleterious effects of significant transportation scale economies along the supply chain. Typically, large retail chains bear the cost of inbound transportation because the most economical choice offered by most manufacturers is a wholesale pricing option that excludes transportation. Due to transportation scale economies, retailers purchase in large batches and transport the goods to warehouses for eventual distribution to retail stores. When a retailer orders in large batches, he has an incentive to offer temporary discounts to clear his inventory more quickly. Although the temporary discounts may allow the retailer to sell to customers who are unwilling to purchase at the regular price, they may also cause customers to stockpile, which reduces the burden of holding inventory for the retailer. Therefore, retailers may order even less frequently and in even larger batches, thereby causing further difficulties for the manufacturer in managing capacity and planning production. Such order batching and temporary retail discounts (high-low pricing) are two causes of the well-known bullwhip effect (Lee et al. 1997).

Transportation scale economies are unlikely to disappear any time soon, so we seek a manufacturer pricing policy that benefits both the manufacturer and retailer vis-a-vis the status quo of the retailer bearing the full cost of inbound transportation. Manufacturers typically offer a few pricing options, from a low unit price without transportation provided to a high unit price with transportation provided. We focus on pricing policies that are equally simple because they are easy to implement and can account for typical laws prohibiting price discrimination. We present and analyze a pricing scheme in which the manufacturer seeks to improve his profitability by partially subsidizing the retailer's transportation cost in exchange for an increase (perhaps zero) in the price per unit. The manufacturer could implement the transportation subsidy by offering to take responsibility for the transportation and charging the retailer a transportation fee that is less than what the retailer currently incurs. The manufacturer designs the pricing scheme to maximize his own profit, recognizing that the retailer will optimize his own ordering and pricing policies, and consequently, retail demands and customer purchasing patterns may also change. We seek solutions that are Pareto-improving for the manufacturer and retailer.

For the special case of a single customer segment, we show that a transportation subsidy is always optimal. For more general cases,, we use numerical examples to explore the impact of the manufacturer's optimal pricing policy on the retailer's replenishment interval and pricing policy, and their combined impact on the bullwhip effect. We also provide insights into conditions that have the greatest potential for profit improvements for the manufacturer and retailer. In particular, sizable gains in both parties' profits arise when the manufacturer chooses a transportation subsidy and adjusted wholesale price that either (i) induce the retailer to lower his prices and thereby reach market segments that were not profitable for the retailer under the manufacturer's initial price structure, or (ii) enable the retailer to order frequently enough that he can utilize a policy with smaller temporary discount and yet increase sales.

The remainder of this chapter is organized as follows. Section 2.2 reviews the relevant literature. In Section 2.3, we formulate the decision problems facing customers, the retailer and the manufacturer, and analyze the customer's problem. We analyze the retailer's and manufacturer's problems in Sections 2.4 and 2.5, respectively. Section 2.6 closes the chapter with a discussion of managerial insights.

2.2 Literature Review

The literature on two-part tariffs, the bullwhip effect, customers' response to price promotions, and retailers' price promotion strategies are all pertinent to our work; we discuss relevant articles below. The literature on the latter two topics is vast, so we limit our discussion to the most relevant papers. Raju (1995) provides a nice overview of various aspects of both trade promotions and retail promotions. See Lal et al. (1996) for a foundational model of manufacturer promotions and retailer forward-buying. For models involving (normative) coordination of manufacturer promotion and production decisions, see Sogomonian and Tang (1993) and references therein.

2.2.1 Two-Part Tariffs

In the model presented in Section 2.3, the retailer faces a two-part tariff with a fixed transportation cost per order either borne internally or paid to a transporter and constant wholesale price per unit paid to the manufacturer. Two-part tariffs have been studied extensively in the economics, marketing and operations management literatures (see, for example, Schmalensee 1981 and references in Sudhir and Datta 2009). A two-part tariff is a special type of quantity discount, and there is an extensive literature on various types of quantity discounts. Dolan (1987) provides a nice discussion of the various motivations for quantity discounts. Jeuland and Shugan 1983, Lal and Staelin 1984, Dada and Srikanth 1987 and Weng 1995, among others, offer approaches that a seller can use to structure discounts to improve supply chain performance.

We are not aware of any research in which one of the manufacturer's (seller's) decisions is a possible subsidy of a cost normally incurred by the retailer. Interestingly, however, Dolan (1987) points out that some sellers offer f.o.b pricing at the *destination* (i.e., inbound freight is included in the wholesale price) and then use quantity discounts as a mechanism for sharing the transportation cost savings (owing to larger shipments) with the retailer and simultaneously shifting the inventory costs to him. In some sense, our analysis provides a mechanism for addressing these tradeoffs, but in our context, the manufacturer generally prefers smaller, more frequent orders, because they reduce the required inventory accumulation in advance of each shipment. Many manufacturers prefer a steady stream of small orders rather than infrequent large orders because of production-related issues as well as the costs of holding inventory, and the manufacturer's inventory holding costs in our model can serve as a proxy for both effects. Clearly, an adjustment in the two-part tariff has the potential for improving supply chain profitability; our concern is how to accomplish this while accounting for the potentially complex reactions of the retailer and customers.

2.2.2 Bullwhip Effect

The bullwhip effect has been observed in industry for a long time. Researchers have described both operational and behavioral causes of the bullwhip effect. Lee et al. (1997) articulated four operational causes of the bullwhip effect: demand signal processing (distortion of demand information when orders are transmitted upstream), shortage gaming (retailer's tendency to inflate orders to gain a larger allocation when supply is scarce), order batching and price fluctuations. Both order batching and price fluctuations arise in our model, and we discuss them further later in this section. Note also that order batching indirectly affects demand signal processing, as infrequent orders lead to greater uncertainty about the order quantity and longer information delays. For a recent survey on the bullwhip effect, see Miragliotta (2006).

Behavioral causes arise from suboptimal decisions and can be explained by the bounded rationality of decision makers, particularly the failure to adequately account for feedback effects, time delays, and in-transit inventory (Sterman 1989; Croson et al. 2004).

2.2.2.1 Order Batching

Lee et al. (1997) indicate that order frequency and the evenness of order arrivals over time are key factors that determine the magnitude of the bullwhip effect. There is a sizable literature, too extensive to review in detail here, that examines the effect of different ordering policies (sometimes in conjunction with different forecasting strategies and under different information regimes). For recent examples, see Warburton (2004) and Wright and Yuan (2007). There is also a stream of research that emphasizes coordination of the timing of orders, either across supply chain members at the same echelon (such as retailers ordering from the same manufacturer; cf. Cachon 1999) or across stages in the supply chain (e.g., Johansson et al. 2000). In these papers, the emphasis is on smoothing demand by balancing orders across time and on synchronizing orders to improve the opportunity to use consistent information. Finally, a few papers consider the impact of standard batch sizes (of which customers must order a multiple) on the bullwhip effect (e.g., Riddalls and Bennett 2001, Holland and Sodhi 2004 and Potter and Disney 2006). Not surprisingly, larger standard batch sizes tend to exacerbate the bullwhip effect due to increased demand distortion.

The main reason for order batching in our model is setup costs. Although a few papers in the literature (all mentioned in Section 2.2.3) consider setups costs, to the best of our knowledge, none considers them in a setting where two or more supply chain members explicitly seek to modify the allocation of setup costs among the parties to improve profitability.

2.2.3 Consumer's Reaction to Price Promotions

There is a vast literature on the impact of price promotions on consumer purchasing decisions such as brand choice and purchase acceleration (i.e., stockpiling). In a variety of empirical studies, researchers have come to somewhat different conclusions about which of these effects is most prominent. In broad terms, however, they generally conclude that brand switching constitutes the main effect, followed by purchase acceleration (see Gupta 1988 and

Bell et al. 1999 for examples). Currim and Schneider (1991) propose a taxonomy of purchasing strategies, categorizing customers according to their propensity to switch brands and/or to accelerate purchases when a promotion is offered, and demonstrate their validity for ground coffee purchases.

Evidence is mixed on whether consumer stockpiling in reaction to promotions leads to an overall increase in consumption. A study by Chandon and Wansink (2002) suggests that stockpiling does increase consumption for both high- and low-convenience products, but Ailawadi and Neslin (1998), Bell et al. (2002) and Ailawadi et al. (2007) report evidence that stockpiling causes an increase in consumption in some product categories, but not in others. Nijs et al. (2001) present empirical results indicating that price promotions do not influence total category sales. Researchers have studied whether purchase acceleration during promotions is offset by lower purchases post-promotion. Hendel and Nevo (2003) present results suggesting that post-promotion dips do occur. They (2004) also provide a survey of research on intertemporal substitution for storable products.

A stream of empirical research examines the effect of promotions on stockpiling behavior in particular. Beasley (1998) presents results of a study which indicates that deal-proneness of the household, the consumer's level of inventory and the depth of the discount affect the incidence of stockpiling behavior. Two studies by Meyer and Assuncao (1990, 1993) suggest that the stockpiling decision depends on the observed price of the good, the distribution of future prices and the customer's on-hand inventory. A study by Aggarwal and Vaidyanathan (1993) indicates that short-term promotions accelerate purchases but longer-term promotions (such as manufacturer coupons) do not. Mela et al. (1998) report evidence that increasing promotional activity over the years has led to customers purchasing more during promotions and less during non-promotional periods.

Krishna (1994) develops a normative model of consumer purchase behavior in the presence of deals whose timing is uncertain and shows that the optimal purchase quantity increases with the uncertainty in timing. Related work has been done by Golabi (1985) and Helsen and Schmittlein (1992), among others. Bucovetsky (1983) presents a normative model of customer stockpiling and shows that the customers' ability to stockpile combined with search costs that limit their visits to different retailers can lead to an equilibrium strategy for the retailers in which there is price dispersion, even when both retailers and customers are homogeneous.

2.2.4 Retailer's Optimal Price Promotion Strategies

There is extensive research that attempts to explain retailers' incentives for offering price promotions. Blattberg et al. (1981) suggest that promotions are a means to transfer inventory costs to the customer. They develop a model to determine the optimal discount and optimal reorder cycle for the retailer. Jeuland and Narasimhan (1985) analyze a variant of the Blattberg et al. model that considers price-sensitive demand and customers whose holding costs are correlated with their purchase intensity. They show that the latter allows the retailer to price discriminate, and that price discrimination is profitable if high demanders have high inventory holding costs and low demanders have inventory holding costs small enough that their forward buying during promotions compensates for the opportunistic behavior of the high demanders. Analysis by Kopalle et al. (1996) indicates that asymmetry in customers' (demand) response around a reference price may make temporary price promotions optimal for the retailer.

Achabal et al. (2001) develop a normative model for the retailer's promotion frequency, regular and promotional prices, and inventory levels in promotional and non-promotional weeks. Each promotion involves a fixed cost (e.g., for advertising) in addition to a price decrease. Demand depends upon the price reduction, the frequency of promotions (due to the impact on purchase acceleration), inventory levels, and seasonal effects. In their model, purchase acceleration necessitates higher retail inventory levels, which is the reverse of the effect in our model.

The Blattberg et al. (1981) paper mentioned earlier considers a situation in which a retailer sells to two customer segments, only one of which will stockpile. The retailer optimizes his replenishment and pricing policy considering the consequent revenue, per-order setup costs and inventory holding costs. We generalize the model to allow both customer segments to stockpile and also include the manufacturer as another decision-maker who makes pricing decisions in view of the retailer's response, which depends upon the customer's purchasing and stockpiling choices.

We do not consider competitive effects in our model. There is a vast literature on promotions in a competitive context, but very little of it considers the effects of customer stockpiling. For examples of papers on this topic, see Kopalle et al. (1999) and Bell et al. (2002).

2.3 Problem Statement and Formulation

We model a scenario with two customer segments of arbitrary sizes, arbitrary but different reservation prices and possibly different holding cost rates. We assume that each customer consumes one unit per unit time (deterministically), provided that he has access to a unit whose gross cost (including the cost of holding inventory from purchase to consumption) does not exceed his reservation price. (It is straightforward to accommodate segments with different consumption rates.) If a customer does not have access to such a unit, he foregoes consumption of the product, possibly purchasing a less expensive (default) substitute whose effects are not considered in the model. Products that have this characteristic include bottled water (the default is tap water) or a national brand product whose consumption remains stable over time (such as toothpaste or laundry detergent) where the default might be a store-brand product. Using the aforementioned demand constructs, we are able to model a situation in which the total demand for the product is price-elastic and both the total demand and the profile of demand (as a function of time) seen at the retail level may differ depending upon the manufacturer's and retailer's decisions. We assume that the presence of customer stockpiles (where applicable) does not change the basic consumption rates. (This is easily justified for products such as toothpaste but there is evidence of small increases in consumption for other products when customers stockpile those products, as mentioned earlier.)

We model the case of a single retailer. Although one motivation for this assumption is tractability, it turns out that the trajectory of total inventory over time in a system consisting of a warehouse and the set of retail stores that it supports is quite similar to the trajectory for a single mega-store that is directly supplied by the manufacturer and supplies the same set of customers. This correspondence occurs because in both systems, the retailer receives orders in batches from the manufacturer in the same way and sells goods to customers at the same pace.

We also assume that customers shop frequently enough that they can take advantage of all promotions if they so choose. Many consumers purchase groceries frequently (e.g., once a week), and most retail chains hold prices of non-perishable goods constant for a week, so our assumption allows us to capture forward-buying behavior reasonably accurately. For simplicity, we assume that the retailer offers a single price discount at the beginning of each replenishment cycle (upon receipt of a shipment from the manufacturer). It can be shown that if the retailer offers a discount only once during each replenishment cycle, this is the optimal timing.

Customer segment i, i=1,2, has λ_i customers, each of whom has a reservation price r_i and holding cost rate h_i per unit time. We assume without loss of generality that $r_1 > r_2$. Although customer heterogeneity can be modeled via a single segment with a range of reservation prices, such a representation leads to extremely messy expressions for quantities purchased at the discount price and necessitates optimization of the regular price over the entire range of reservation prices. Our representation of customer heterogeneity avoids both of these complications.

The manufacturer has a constant unit production cost c (which we normalize to zero) and an annual inventory holding cost per unit h_m. We also assume that the manufacturer produces at a finite rate P while operating, and that his production rate is adequate to satisfy his demand.

The current transportation (setup) cost borne by the retailer is K per shipment and the current wholesale price is w per unit. (Shipment quantity constraints can be accommodated by limiting the replenishment interval.) In practice, K is either a cost borne internally by the retailer if he uses his own truck fleet or a cost paid by the retailer to a third-party transporter, and in the medium-term, the associated per-shipment costs are not controllable. We also assume that the status quo value of w is exogenous. For a variety of reasons, including competitive factors and price discrimination laws, the manufacturer may not be able to optimize w for a given K, if he knows the value of K at all. (In particular, the retailer cannot offer different wholesale prices simply because the retailers face different transportation costs.) Furthermore, in many practical settings, the value of w is negotiated. One might view our model as one in which the manufacturer is negotiating a transportation subsidy and wholesale price together, starting from any arbitrary status quo situation. An offer of this form is a reasonable *quid pro quo* that has the same type of structure as price structures commonly offered in practice, as mentioned earlier. We note, however, that our analysis can be used to optimize w for any fixed K, and provides insights for this case, as well.

The manufacturer chooses a (possibly subsidized) cost per shipment, K' and a wholesale price, w'. The retailer has an annual holding cost per unit h_r. The retailer chooses an "everyday" price, p, the discount, Δ, and the replenishment interval, T. The customers choose how much to purchase when the discount price is offered and whether to purchase at other times. We assume that none of the parties has a strict limit on inventory storage. To avoid trivial solutions, we assume that: (1) customers have finite, positive holding cost rates; and (2) $r_2 > c$, so there exists a discount price that allows the retailer to profitably sell to segment 2.

We analyze the problem as a Stackelberg game with the manufacturer as the first-mover, the retailer as the second-mover, and the customer as the third (and last) mover. The manufacturer and retailer must anticipate how the downstream player(s) will respond.

2.3.1 Customer's Problem

Given the retailer's decisions T, p and Δ and his own parameters, the customer seeks to minimize his total purchase and holding costs, with the stipulation that he will not purchase if the gross cost of a unit, i.e., the unit purchase cost plus the cost of holding the unit from purchase to consumption, exceeds his reservation price. In addition, customers do not stockpile more than what they can consume during the retailer's replenishment cycle, T, simply because the retailer will offer the same discount at the beginning of the next cycle.

Before analyzing the customer's decisions, we first examine whether we can limit the range of regular prices that the retailer needs to consider. Clearly $p < r_2$ is suboptimal because the retailer could charge up to r_2 without sacrificing any demand. If $p > r_1$, all customers will buy only at the discounted price (if at all), and the retailer would achieve the same or better result by setting $p = r_1$ with an appropriately-selected discount. Finally, any other price in the open interval (r_1, r_2) is suboptimal because raising the price to r_1 would not result in any loss of demand. Therefore, the optimal regular price is either r_1 or r_2. We now analyze the customer's purchasing and stockpiling behavior for these two cases.

2.3.1.1 $p = r_1$

When the regular price is r_1, the discount may not be large enough to entice customers in segment 2 to buy. If the discount price $r_1 - \Delta$ exceeds r_2, customers in segment 2 will not purchase at all. Customers in segment 1 stockpile to satisfy their consumption for a duration Δ/h_1, after which they purchase at the regular price for immediate consumption. As noted earlier, customers stockpile to satisfy demand for a duration of at most T. Thus, segment 1 stockpiles to satisfy demand for a duration $t_1 = \min\{\Delta/h_1, T\}$. Observe that it is not advantageous for the retailer to choose $\Delta > h_1 T$ because this will not induce customers in segment 1 to buy more and simply reduces the price on all units sold to that segment. Segment 2 does not purchase at all, so revenue from that segment is unaffected.

On the other hand, if the discount price, $r_1 - \Delta$, is less than r_2, then customers in segment 1 stockpile in the same way as described above, but the situation is different for customers in segment 2 because they are willing to pay a gross cost of at most r_2 per unit. The surplus per unit of a customer in segment 2 when purchasing at the discount price is $r_2 - (r_1 - \Delta)$, and this is the maximum that such a customer is willing to spend on holding a single unit until consumption. Therefore, customers in segment 2 stockpile to satisfy consumption for a duration $t_2 = \min\{[r_2 - (r_1 - \Delta)]/h_2, T\}$.

2.3.1.2 $p = r_2$

When the regular price is r_2 and a discount Δ is offered, then the two segments stockpile to satisfy demand for a duration $t_i = \min\{\Delta/h_i, T\}$ units, $i = 1, 2$. When their stockpiles are depleted, they purchase at price r_2 to satisfy consumption needs on a just-in-time basis.

From the above, we can write a formula for each customer's stockpiling duration that considers both cases: $t_i = \min(T, \frac{[\Delta - (p - r_i)^+]^+}{h_i})$. This formula explicitly accounts for T as the upper bound on each customer's stockpiling duration. We initially assume, however, that $t_i \leq T$, $i = 1, 2$ (i.e., the constraint is not binding), because this is a reasonable assumption for a wide range of products. Our rationale is as follows. Due to internal material handling costs associated with small order quantities (i.e., the non-reducible part of the setup cost incurred by the retailer for each shipment), for most non-perishable, non-bulky products with low to moderate demands, even very large retailers place orders only a few times a year, so T may be as long as several months. (As an example, the authors are familiar with a drug store chain with over 500 stores whose system-wide demand for a typical stockkeeping unit totals only several pallets a year, partly because some of the items are quite small (e.g., dental floss) and partly because customer demands are divided among a great diversity of products within a product category.) Furthermore, it is unusual in practice for customers to stockpile for many months ahead, if only due to space constraints. (Indeed, the literature suggests that customers often underbuy when large purchases would be suggested by purely economic considerations; see Meyer and Assancao 1990.) Our initial assumption captures situations where changes in the manufacturer's pricing policy have a reasonable chance of affecting the degree of customer stockpiling. Nevertheless, we provide analysis in Appendix A to show what happens when customers would be willing to stockpile extensively. Although the algebraic details differ, the essential structural results and insights carry over to the more general situation.

2.3.2 Retailer's Problem

Now we formulate the retailer's problem of maximizing his profit per unit time. We often express revenue as the difference between revenue at full price and the total discount for all units purchased at the discount price. For notational simplicity, we omit the primes on K' and w' as they are constants in this section. As stated earlier, the retailer sets $p = r_2$ or r_1. If $p = r_1$, the retailer can choose a small discount and sell only to segment 1, or he can choose $\Delta > r_1 - r_2$ and sell to both segments. We discuss the various pricing options in turn.

2.3.2.1 Option 1: $p = r_2$

When $p = r_2$, both segments purchase to satisfy consumption throughout the cycle. In addition, both groups stockpile when a discount is offered, with $t_i = \Delta/h_i$ for $i = 1, 2$. If the customer segment's respective stockpile is depleted before the end of the retailer's replenishment cycle, customers in that segment continue to purchase at the regular price for the remainder of the cycle. Thus, the gross margin per cycle is $(r_2 - w)(\lambda_1 + \lambda_2)T - \Delta(\lambda_1 t_1 + \lambda_2 t_2)$. The retailer sells $\lambda_1 t_1 + \lambda_2 t_2$ units instantaneously at the discount price at the beginning of each cycle. Then, when customers in one segment have depleted their stockpiles, the retailer starts to deplete inventory at that segment's demand rate and then finally starts to deplete inventory at the sum of the demand rates when customers in the other segment have depleted their stockpiles. Therefore, the area under the retailer's inventory curve for one cycle is $\frac{h_r}{2}((\lambda_1 + \lambda_2)T^2 - \lambda_1 t_1^2 - \lambda_2 t_2^2)$ (see inventory diagram in Figure 2.1).

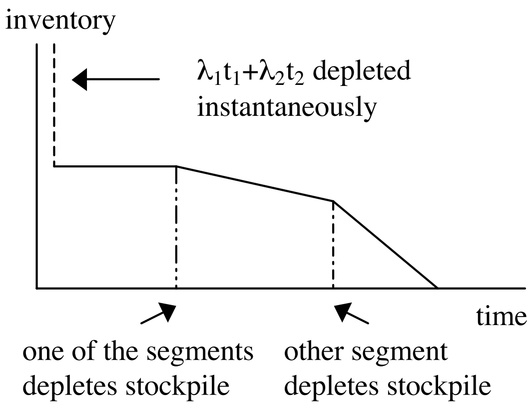

Figure 2.1: Inventory Diagram for Option 1

After substituting $t_i = \Delta/h_i$, the objective can be written as

$$\Pi_1(\Delta, T) = (r_2 - w)(\lambda_1 + \lambda_2) - \frac{K + (\sum \frac{\lambda_i}{h_i} - \frac{h_r}{2} \sum \frac{\lambda_i}{h_i^2})\Delta^2}{T} - \frac{h_r}{2}(\lambda_1 + \lambda_2)T. \tag{2.1}$$

2.3.2.2 Option 2: $p = r_1$ and $\Delta \leq r_1 - r_2$

When $p = r_1$ and $\Delta \leq r_1 - r_2$, segment 2 does not purchase at all. Segment 1 stockpiles to satisfy consumption for a duration $t_1 = \Delta/h_1$. The retailer sells $\lambda_1 t_1$ units instantaneously at the beginning of each cycle, and then starts to deplete inventory again when customers in segment 1 have depleted their stockpiles. The retailer's gross margin per cycle is $(r_1 - w)\lambda_1 - \Delta\lambda_1 t_1/T$ and the area under the retailer's inventory curve for one cycle is $\lambda_1 T^2 - \lambda_1 t_1^2$.

After substituting for t_1, the retailer's objective can be written as

$$(r_1 - w)\lambda_1 - \frac{(\frac{\lambda_1}{h_1} - \frac{h_r}{2}\frac{\lambda_1}{h_1^2})\Delta^2 + K}{T} - \frac{\lambda_1 h_r}{2}T. \tag{2.2}$$

2.3.2.3 Option 3: $p = r_1$, $\Delta > r_1 - r_2$

Similarly to the previous case, only segment 1 purchases at the regular price, but the discount is large enough that segment 2 will purchase at the discount price and then stockpile the purchase for future consumption. Hence we have $t_1 = \Delta/h_1$ and $t_2 = \frac{\Delta-(r_1-r_2)}{h_2}$. Substituting for t_1 and t_2, we can write the retailer's profit function as:

$$\Pi_3(\Delta, T) = (r_1 - w)\lambda_1 - \frac{K + g(\Delta)}{T} - \frac{h_r\lambda_1}{2}T, \tag{2.3}$$

$$\text{where } g(\Delta) = (\sum \frac{\lambda_i}{h_i} - \frac{h_r}{2}\frac{\lambda_1}{h_1^2})\Delta^2 - \frac{\lambda_2}{h_2}(r_1 + r_1 - r_2 - w)\Delta + (r_1 - w)(r_1 - r_2)\frac{\lambda_2}{h_2}. \tag{2.4}$$

Note that $g(\Delta)$ reflects all of the effects of the discount: reduction in revenue from segment 1, revenue from segment 2, and the savings in inventory holding costs due to customer stockpiling.

We provide a full analysis of the retailer's problem in Section 2.4.

2.3.3 Manufacturer's Problem

The manufacturer maximizes his profit by choosing a price pair (K', w') to offer to the retailer. Under the status quo, the retailer pays a transportation cost of K per shipment either borne internally or paid to a third party and a wholesale price of w per unit. Unlike the retailer, who can optimize his profit without regard for the impact of his choices on the manufacturer, the manufacturer needs to ensure that the retailer will accept the new arrangement. Thus, the manufacturer must account for the retailer's participation constraint.

In its most general form, the manufacturer's objective function is quite complex. As we will show in the next section, the set of potentially optimal solutions for the retailer is small and highly structured, and this allows for a significant simplification of the problem(s) that the manufacturer needs to solve. We use the term *dominant* to refer to these potentially optimal solutions. Here, we present the manufacturer's objective function in shorthand notation:

$$\begin{aligned}\max_{K',w'} \quad & R(T^*(K',w'), p^*(K',w'), \Delta^*(K',w')) - cD(T^*(K',w'), p^*(K',w'), \Delta^*(K',w')) \\ - \quad & h_m I(T^*(K',w'), p^*(K',w'), \Delta^*(K',w')) - (K - K')/T^*(K',w')\end{aligned}$$

where R, D and I are the manufacturer's revenue rate, (aggregate) average demand rate and average inventory levels, respectively, and p^*, Δ^* and T^* are the retailer's optimal decisions for the given (K', w'). The last term is the transportation subsidy per unit time. We analyze the manufacturer's problem in Section 2.5.

2.4 Solving the Retailer's Problem

For each profit function in (2.1),(2.2) and (2.3), we first optimize Δ for an arbitrary fixed T and then use these results to find the corresponding optimal T. These results, combined with other results obtained later in this section, reduce the set of candidate optima significantly.

Throughout our analysis, we make comparisons where dominance depends upon one expression being less (greater) than another. For ease of exposition, we omit the points of equality, and the reader should recognize that the equality could be associated with either the "less than" or "greater than" with no loss of optimality, and when equality holds, weak dominance applies. We also use the notation $\Pi_i(\cdot)$ to refer both to a profit function and the corresponding policy; in all cases, the meaning should be clear from the context.

2.4.1 Optimizing Δ for Each Option

Characteristics of the optimal Δ for a fixed T depend heavily upon the signs of the coefficients of Δ^2 in (2.1), (2.2) and g. Let α denote $\sum \frac{\lambda_i}{h_i} - \frac{h_r}{2} \sum \frac{\lambda_i}{h_i^2}$ (the coefficient of Δ^2 in (2.1)), β denote $\frac{\lambda_1}{h_1} - \frac{h_r}{2} \frac{\lambda_1}{h_1^2}$ (i.e., the coefficient of Δ^2 in (2.2)) and γ denote $\sum \frac{\lambda_i}{h_i} - \frac{h_r}{2} \frac{\lambda_1}{h_1^2}$ (i.e., the coefficient of Δ^2 in g). These expressions are functions only of the problem data, so their signs can be determined in advance.

The value of α is positive if

$$h_r < 2 \sum \frac{\lambda_i}{h_i} \Big/ \sum \frac{\lambda_i}{h_i^2}$$

To provide some intuition regarding this condition, in the special case of $h_1 = h_2 = h$, the right hand side of the above expression is equal to $2h$. Loosely speaking, the relation is a condition on h_r relative to the (weighted average) holding costs of the two segments.

The condition $\beta > 0$ is equivalent to $h_r < 2h_1$.

The value of γ is positive if

$$h_r < 2h_1 + \frac{2h_1^2 \lambda_2}{h_2 \lambda_1}$$

which again is a condition on the relationship between h_r and a function of the holding costs of the two customer segments. To provide some intuition, in the special case where $h_2 = h_1$ and $\lambda_1 = \lambda_2$, the right hand side of the above expression is equal to $4h_1$.

Observe that $\alpha > 0$ implies $\gamma > 0$, $\gamma < 0$ implies $\alpha < 0$, $\gamma < 0$ implies $\beta < 0$ and $\gamma < 0$ implies $\alpha < 0$. Thus, the only possible combinations of signs are shown in Table 1.

Case	α	β	γ
I	+	-	+
II	-	-	+
III	+	+	+
IV	-	+	+
V	-	-	-

Table 2.1: Combinations of Signs of Coefficients of Δ^2

By analyzing terms in the profit functions that depend upon Δ, we are able to determine the optimal Δ for each profit function. We summarize the results for Cases I through IV in Table 2.1. We use two pieces of shorthand notation in the Table, Δ_0 and Δ^{UB}, which we define next. Δ_0 is the stationary point of g. Δ^{UB} denotes the optimal discount for situations where $\alpha < 0$ (i.e., Cases II and IV). In these situations, the optimal solution for Option 1 has the property that at least one of the segments stockpiles for the entire replenishment cycle. (Technically, Δ^{UB}

	$\alpha > 0$	$\alpha < 0$
$\beta < 0$ $(h_r > 2h_1)$	Case I $\Pi_1(0), \Pi_2(r_1 - r_2), \Pi_3(\Delta_0)$	Case II $\Pi_1(\Delta^{UB}), \Pi_2(r_1 - r_2), \Pi_3(\Delta_0)$
$\beta > 0$ $(h_r < 2h_1)$	Case III $\Pi_1(0), \Pi_2(0), \Pi_3(\Delta_0)$	Case IV $\Pi_1(\Delta^{UB}), \Pi_2(0), \Pi_3(\Delta_0)$

Table 2.2: Potentially Optimal Policies Considering Optimal Discounts

is not an upper bound, but the nomenclature serves as a reminder that the discount is fairly large.) Detailed derivations of the results in the table can be found in Appendix D. The results in Table 2.2 rely only on the functional forms of the terms involving Δ. We can eliminate other policies due to (economic) dominance of one profit function over the other (see Appendix D for details).

Additional results can be derived for Cases III and IV, where $\beta > 0$ so for Π_2, $\Delta^* = 0$. We can show that $\Pi_2(0)$ dominates $\Pi_3(\Delta_0)$ if w' exceeds a threshold which we call w_2 and the reverse holds otherwise. (The threshold w_2 is the value of w' that equates $g(\Delta_0)$ to 0; $g(\Delta_0) <> 0$ for $w' <> w_2$.) Thus, if there exists no $\Delta > r_1 - r_2$ (the range of Δ values for which Π_3 is applicable) such that $g(\Delta) < 0$, then $\Pi_3(\Delta_0)$ is dominated by $\Pi_2(0)$ and the converse holds if any such feasible Δ exists. Thus, for any set of problem parameters, the retailer only needs to compare two potentially optimal policies, as shown in Table 2.3. (Whether Π_2 or Π_3 needs to be considered depends upon the value of w.)

	$\alpha > 0$	$\alpha < 0$
$\beta < 0$ $(h_r > 2h_1)$	Case I $\Pi_1(0), \Pi_3(\Delta_0)$	Case II $\Pi_1(\Delta^{UB}),\Pi_3(\Delta_0)$
$\beta > 0$ $(h_r < 2h_1)$	Case III $\Pi_1(0)$, and $\Pi_2(0)$ or $\Pi_3(\Delta_0)$	Case IV $\Pi_1(\Delta^{UB})$, and $\Pi_2(0)$ or $\Pi_3(\Delta_0)$

Table 2.3: Reduced Set of Potentially Optimal Policies

Case V is the only case with $\gamma < 0$, which corresponds to situations in which the retailer's holding cost is quite high in comparison to the customers'. In Appendix D, utilizing our earlier analysis characterizing how the profit functions change with Δ, we show that only $\Pi_1(\Delta^{UB})$, $\Pi_2(r_1 - r_2)$ and $\Pi_3(\Delta_{ub})$ need to be considered, where Δ_{ub} is the optimal solution applicable to Option 3 which is analogous to Δ^{UB} for Option 1. We also show that $\Pi_2(r_1 - r_2)$ is dominated by $\Pi_3(\Delta_{ub})$, hence only $\Pi_1(\Delta^{UB})$ and $\Pi_3(\Delta_{ub})$ remain as potentially optimal solutions.

To summarize, the retailer must evaluate the solutions listed in Table 2.3 for Cases I, II, III, and IV, and $\Pi_1(\Delta^{UB})$ and $\Pi_3(\Delta_{ub})$ for Case V, and choose the best one.

We now seek to identify regions in the two-dimensional space of w' and K' in which the retailer prefers each of these policies. As before, we drop the primes on w' and K' as they are constants in the retailer's problem. We present the analysis for cases with $\alpha > 0$ in detail and later explain what differs for analogous cases with $\alpha < 0$.

2.4.2 Retailer's Dominant Policy Regions

We divide our analysis according to the cases in Table 2.3, first considering cases with $\alpha > 0$.

2.4.2.1 Case I: $\alpha > 0, \beta < 0$ ($\sum \frac{\lambda_i}{h_i} - \frac{h_r}{2} \sum \frac{\lambda_i}{h_i^2} > 0, h_r > 2h_1$)

The profit functions for the two dominant choices are: -0.05in

$$\Pi_1(0,T) = (r_2 - w)(\lambda_1 + \lambda_2) - \frac{K}{T} - \frac{h_r}{2}(\lambda_1 + \lambda_2)T,$$
$$\Pi_3(\Delta_0,T) = (r_1 - w)\lambda_1 - \frac{K + g(\Delta_0)}{T} - \frac{h_r\lambda_1}{2}T.$$

Maximizing $\Pi_1(0,T)$ with respect to T gives $T^* = \sqrt{\frac{2K}{(\lambda_1+\lambda_2)h_r}}$. Substituting this expression for T^* into the objective, we have

$$\Pi_1^*(\Delta = 0) = \Pi_1(\Delta = 0, T^*) = (r_2 - w)(\lambda_1 + \lambda_2) - \sqrt{2h_r(\lambda_1 + \lambda_2)}\sqrt{K}. \tag{2.5}$$

Similarly, maximizing $\Pi_3(\Delta_0,T)$ gives $T^* = \sqrt{\frac{2(K+g(\Delta_0))}{h_r\lambda_1}}$ (if $K + g(\Delta_0) \geq 0$), so

$$\Pi_3^*(\Delta_0) = \Pi_3(\Delta_0, T^*) = (r_1 - w)\lambda_1 - \sqrt{2h_r\lambda_1}\sqrt{K + g(\Delta_0)}. \tag{2.6}$$

Both profit functions are convex decreasing in K, and $\Pi_1^*(\Delta = 0)$ decreases more rapidly than $\Pi_3^*(\Delta_0)$ because $\lambda_2 \geq 0$ and $g(\Delta_0) < 0$. (For a proof that $g(\Delta_0) < 0$ in this case, see Appendix F.) Hence if $(r_2 - w)(\lambda_1 + \lambda_2) \leq (r_1 - w)\lambda_1$, then $\Pi_3^*(\Delta_0)$ dominates $\Pi_1^*(\Delta = 0)$ for all K, or equivalently, there is a threshold $w_1 \equiv \frac{r_2(\lambda_1+\lambda_2)-r_1\lambda_1}{\lambda_2}$ such that if $w \geq w_1$, $\Pi_3^*(\Delta_0)$ dominates. On the other hand, if $w < w_1$, $\Pi_1^*(\Delta = 0)$ is dominant to the southwest, and $\Pi_3^*(\Delta_0)$ is dominant to the northeast, of the switching curve which is defined by points where the two profit functions are equal. Figure 2.2 shows a typical dominance map for this case. The switching curve is not necessarily monotonic, but, it is straightforward to show it is unimodal and the stationary point (with respect to K) is negative, so the function is monotonically decreasing over the relevant range.

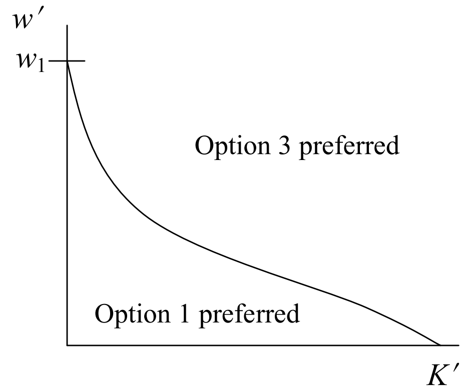

Figure 2.2: Dominance Map 1 (Case I)

Thus far we have assumed that $K + g(\Delta_0) \geq 0$. In the unusual case where $K + g(\Delta_0) <$

0, the benefits of offering the optimal discount, Δ_0, are so large that they outweigh the setup cost, so the retailer has an incentive to order extremely frequently. But because the manufacturer will be worse off as T goes to zero, he will choose w and K to prevent the retailer from choosing this option.

2.4.2.2 Case III: $\alpha > 0, \beta > 0$ $(\sum \frac{\lambda_i}{h_i} - \frac{h_r}{2} \sum \frac{\lambda_i}{h_i^2} > 0, h_r \leq 2h_1)$

Case III is more complicated as we have to consider all three profit functions. However, as mentioned earlier, if $g(\Delta_0) > 0$ at the given w, $\Pi_3^*(\Delta_0)$ is dominated by $\Pi_2^*(\Delta = 0)$ and the reverse holds if $g(\Delta_0) < 0$ at the given w. With some analysis (see Appendix G for details), we can show that the dominance map for this case has the form shown in Figure 2.3, where w_2 is the value of w at which $g(\Delta_0) = 0$. In particular, we show that such a threshold, which separates the regions in which Option 2 and Option 3 are dominant (for large K') exists, and that the curve shown in the diagram is continuous and monotonically decreasing.

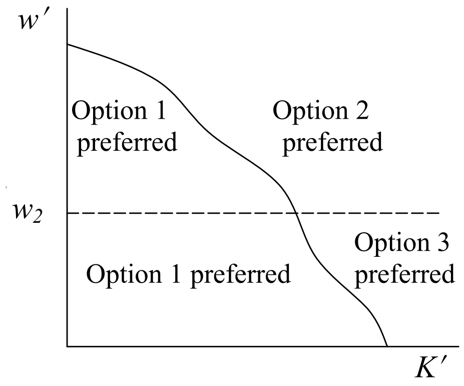

Figure 2.3: Dominance Map 2 (Case III)

2.4.2.3 Cases II and IV: $\alpha < 0$

The condition $\alpha < 0$ implies that the retailer's holding cost is large in comparison to the "weighted average" holding costs of the two customer segments, so Δ is at an upper limit rather than at a lower limit. The analysis for Cases II and IV is otherwise analogous to that for Cases I and III, respectively, replacing $\Pi_1(0)$ by $\Pi_1(\Delta^{UB})$. The objective function takes different forms depending upon which customer segment has the higher holding cost. Details appear in Appendix A in the subsection titled "Analysis of Π_1 for Large Δ."

The resulting optimal profit functions differ from $\Pi_1^*(\Delta = 0)$ obtained in Case I (cf. (2.5)) only in the coefficient of $\sqrt{K'}$, and in both cases, the coefficients are simply constants that depend upon the problem parameters. Thus, we can define switching curves analogous to those derived for Cases I and III. The dominance maps have qualitatively the same structure as those for Cases I and III, respectively. As in the cases analyzed earlier, there is a threshold value of w that plays a role in the dominance map.

2.4.2.4 Implications

Here, we discuss implications of Dominance Maps 1 and 2 for Cases I and III, respectively, but qualitatively similar conclusions can be drawn for Cases II and IV. We wish to answer the following question: What does the manufacturer have to do induce the retailer to choose another pricing option that may be advantageous for the manufacturer as well?

For the situation illustrated in Dominance Map 1, $\Pi_1(0)$ and $\Pi_3(\Delta_0)$ (with their corresponding T^* values) are candidate optimal solutions. If the current K and w are to the northeast of the switching curve, the manufacturer can induce the retailer to switch by choosing w' and K' to the southwest of the switching curve. The retailer will then prefer $\Pi_1(0)$, i.e., an everyday low price policy under which he sells to *both* segments rather than only one segment. Both the retailer and manufacturer may gain due to greater market penetration. Furthermore, because the retailer would then be holding more inventory (as $\Delta = 0$) and because K' is relatively small, he orders frequently in small quantities, so the manufacturer may also reduce his holding costs.

For situations corresponding to Dominance Map 2, the manufacturer may shift to a high value of w', which induces the retailer to choose an "every day high price" policy; thus, he sells only to customers in Segment 1 and those customers do not stockpile. However, the numerical results reported later in this chapter suggest that such a shift is rarely preferred by the manufacturer. Instead, the manufacturer prefers to induce the retailer to shift his pricing strategy to entice Segment 2 customers to purchase for at least part of the cycle. This may require a transportation subsidy, a reduction in the wholesale price, or possibly both. Either shift in the retailer's pricing strategy enables the retailer to sell to segment 2, which may lead to a substantial increase in the retailer's profits, and possibly also the manufacturer's. Of course, the manufacturer will only choose a transportation subsidy and adjustment in the wholesale price that ultimately benefit himself. We explore these phenomena further after analyzing the manufacturer's problem.

2.5 Manufacturer's Problem

The dominance maps reveal that for any set of problem parameters, the retailer will choose among at most three options. Thus, to solve the manufacturer's problem for each dominance map, we only need to solve one constrained optimization for each possible retailer response, where within each problem, the constraints consist of the retailer's incentive compatibility (IC) constraint and his participation constraint. In modeling the retailer's participation constraint that the manufacturer must consider, we assume that the status quo values of K and w are prespecified. The manufacturer chooses the pricing solution with the highest profit. (By definition, one of these solutions must satisfy his own participation constraint.)

Because the optimization problems involve two nonlinear constraints containing square roots, solutions for most of the problems cannot be obtained in closed form. Consequently, in the remainder of this section, we focus on problem formulations. Results that simplify the analysis and sketches of the solution procedures can be found in Appendix H. In general, the problems can be solved by careful application of standard nonlinear optimization techniques. We discuss each dominance map in turn. Note that for some problem parameters there may be no feasible solution to the manufacturer's problem in one or more regions of the dominance map; such alternatives may be eliminated from consideration. In this section, we use the following additional notation: Δ_0' is the optimal value of Δ_0 at $w = w'$ and t_2' is the value of t_2 at $\Delta = \Delta_0'$.

Below, we index options in the same way as the retailer's pricing options (cf. Section 2.3.2).

Dominance Map 1: $\Pi_1(0)$ versus $\Pi_3(\Delta_0')$

Option 1: Retailer chooses $\Pi_1(0)$ with $T^* = \sqrt{\frac{2K'}{(\lambda_1+\lambda_2)h_r}}$

After substituting for T^*, the manufacturer's objective becomes

$$\Pi_1^m = (\lambda_1+\lambda_2)w' - \frac{K\sqrt{(\lambda_1+\lambda_2)h_r/2}}{\sqrt{K'}} + \frac{\sqrt{\lambda_1+\lambda_2}[h_rP-(\lambda_1+\lambda_2)h_m]}{\sqrt{2h_r}}\sqrt{K'} \tag{2.7}$$

The participation and incentive compatibility constraints are, respectively,

$$\begin{aligned} w' &\leq \left[r_1(\lambda_1+\lambda_2) - \sqrt{2h_r(\lambda_1+\lambda_2)}\sqrt{K'} - \Pi_0\right]/(\lambda_1+\lambda_2), \text{ and} \\ w' &\leq \left[r_2(\lambda_1+\lambda_2) - r_1\lambda_1 - \sqrt{2h_r(\lambda_1+\lambda_2)}\sqrt{K'} + \sqrt{2h_r\lambda_1}\sqrt{K'+g(\Delta_0')}\right]/\lambda_2, \end{aligned} \tag{2.8}$$

where Π_0 denotes the retailer's current profit.

Option 3: Retailer chooses $\Pi_3(\Delta_0)$ with $T^* = \sqrt{\frac{2[K'+g(\Delta_0')]}{\lambda_1 h_r}}$

After substituting for T^*, Δ_0' and t_2', the manufacturer's objective becomes:

$$\Pi_3^m = \lambda_1 w' + \frac{\sqrt{h_r\lambda_1}[K'-K+f(w')]}{\sqrt{2[K'+g(\Delta_0')]}} - \frac{h_m\lambda_1\sqrt{2[K'+g(\Delta_0')]}}{2h_r} - \frac{h_m\lambda_1\lambda_2(c_1w'+c_2)}{P} \tag{2.9}$$

where $f(w')$ is a quadratic function of w' in which the coefficients depend only on the problem data and c_1 and c_2 are constants that depend only on on the problem data (see Appendix H for details).

The IC constraint is the same as given in (2.8) but with the inequality reversed. The participation constraint is $(r_1 - w')\lambda_1 - \sqrt{2h_r\lambda_1}\sqrt{K'+g(\Delta_0')} \geq \Pi_0$.

Dominance Map 2: $\Pi_1(0)$ versus $\Pi_2(0)$ or $\Pi_3(\Delta_0')$

The relevant objectives and participation constraints were presented earlier. The IC constraints only require that $w' > w_2$ ($< w_2$) if the retailer chooses Π_2^* (Π_3^*).

Option 1: Retailer chooses $\Pi_1(0)$ with $T^* = \sqrt{\frac{2K}{(\lambda_1+\lambda_2)h_r}}$

The objective and participation constraint are the same as for Dominance Map 1. There are two different IC constraints depending upon whether $w >< w_2$.

For $w' > w_2$, the IC constraint is:

$$w' \leq \left[r_2(\lambda_1+\lambda_2) - r_1\lambda_1 - \sqrt{2h_r(\lambda_1+\lambda_2)}\sqrt{K} + \sqrt{2h_r\lambda_1}\sqrt{K}\right]/\lambda_2. \tag{2.10}$$

For $w' < w_2$, the IC constraint is the same as given for Dominance Map 1 when the retailer chooses $\Pi_1(0)$.

Option 2: Retailer chooses $\Pi_2(0)$ with $T^* = \sqrt{\frac{2K'}{\lambda_1 h_r}}$

After substituting for T^*, the manufacturer's objective becomes:

$$\Pi_2^m = \lambda_1 w' - \frac{K\sqrt{\lambda_1 h_r/2}}{\sqrt{K'}} + \frac{\sqrt{\lambda_1+\lambda_2}(h_rP-\lambda_1 h_m)}{\sqrt{2h_r}}\sqrt{K'}. \tag{2.11}$$

The retailer's participation constraint is $(r_1 - w')\lambda_1 - \sqrt{2h_r\lambda_1}\sqrt{K'} \geq \Pi_0$ and the IC constraint is the same as that given in (2.10) but with the inequality reversed.

For Cases II and IV, the formulations can be constructed similarly.

Before closing this section, we note that if the manufacturer wishes to optimize w' for any fixed K (e.g., the current value), he still needs to account for the retailer's choices as reflected in the dominance maps. The problem, of course, is easier because he only needs to identify the relevant threshold value of w and solve the appropriate constrained optimization problems for w' above and below that threshold. Likewise, the manufacturer can optimize K' for any fixed w, which would be relevant in competitive situations where the manufacturer is a price-taker.

2.5.1 The Single-Segment Problem

Due to the complex form of the manufacturer's problem, we cannot obtain closed-form solutions for the case of two customer segments. To obtain additional insights, we analyze the single-segment case. The retailer only has one pricing option, Π_2, so the retailer's objective function reduces to

$$\Pi_2(\Delta, T) = (r_1 - w')\lambda_1 - \frac{K' + \frac{\lambda_1}{h_1}(1 - \frac{h_r}{2h_1})\Delta^2}{T} - \frac{h_r}{2}\lambda_1 T.$$

Clearly, the retailer's optimal discount, Δ^*, depends on the sign of $1 - \frac{h_r}{2h_1}$. We consider the two possibilities in turn.

2.5.1.1 Case A: $1 - \frac{h_r}{2h_1} \geq 0$

When $1 - \frac{h_r}{2h_1} \geq 0$, the coefficient of Δ in the retailer's objective function is negative and hence $\Delta^* = 0$. With $\Delta^* = 0$, the retailer's first order necessary condition with respect to T gives

$$T^* = \sqrt{\frac{2K'}{\lambda_1 h_r}}.$$

Substituting for T^* in the manufacturer's objective, we get

$$\Pi_2^m = \lambda_1 w' - \frac{(K - K')\sqrt{\lambda_1 h_r}}{\sqrt{2K'}} - \frac{h_m \lambda_1^2 \sqrt{K'}}{P\sqrt{2\lambda_1 h_r}}. \tag{2.12}$$

Because the retailer has only one pricing option in the single-segment problem, IC constraints are unnecessary. The retailer's participation constraint is binding, as we explain next. From (2.12), observe that for any K', the manufacturer's objective is increasing in w' so he would like to increase w' as much as possible for any arbitrary K'. The retailer's objective evaluated at T^*, i.e., $(r_1 - w')\lambda_1 - \sqrt{2h_r\lambda_1 K'}$, is decreasing in both w' and K'. Thus, for any K' the manufacturer will choose a w' that necessarily makes the retailer's participation constraint binding.

Applying the participation constraint $\Pi_2^* = \Pi_0$, where Π_0 is the retailer's profit under the current pricing regime, we get

$$w'^* = w + \frac{\sqrt{2h_r\lambda_1}(\sqrt{K} - \sqrt{K'})}{\lambda_1}. \tag{2.13}$$

Taking the first derivative of Π_1^m with respect to K' and substituting for w'^* and then for Π_0, and performing a few additional algebraic manipulations, we find that K'^* satisfies

$$\frac{K'^*}{K} = \left[1 + \frac{h_m}{h_r}\frac{\lambda_1}{P}\right]^{-1}. \tag{2.14}$$

We can obtain w'^* by substituting K'^* back into (2.13).

The formula in (2.14) suggests that it is always optimal for the manufacturer to subsidize transportation and the optimal fractional transportation subsidy increases with h_m/h_r and λ_1/P. As h_m/h_r increases, the relative economic disadvantage of the manufacturer holding inventory vis-a-vis the retailer increases. As λ_1/P increases, the manufacturer's holding costs incurred while accumulating stock prior to each shipment increases. Both of these effects increase the manufacturer's incentive to provide a larger transportation subsidy. The formula also makes clear the essence of the tradeoff: the manufacturer's holding cost is proportional to $h_m\lambda_1/P$ while the retailer's holding cost is proportional to h_r. (Recall that $\Delta = 0$ in this case, so the retailer does not transfer any of the holding costs to the customers.) So the fractional transportation subsidy depends directly on the ratio of these two cost expressions.

2.5.1.2 Case B: $1 - \frac{h_r}{2h_1} < 0$

Following analysis similar to that for Case A, we can derive the optimal solution to the retailer's problem as $\Delta^* = h_1 T^*$ and $T^* = \sqrt{\frac{K'}{\lambda_1 h_1}}$. Then continuing as in Case A, we obtain

$$\frac{K'^*}{K} = \left[1 + \frac{h_m \lambda_1}{2h_1 P}\right]^{-1}. \tag{2.15}$$

As in (2.14), the optimal fractional transportation subsidy increases with λ_1/P, but instead of increasing with h_m/h_r, the subsidy increases with $h_m/2h_1$. This relationship arises because T^* is a function of h_1 instead of h_r, as it was in Case A. In Case B, the retailer holds no inventory because the customers stockpile to satisfy their demand for the entire cycle, so the key tradeoff is between the manufacturer's holding cost h_m and the customer's holding cost h_1. As in Case A, the retailer's participation constraint is binding and a transportation subsidy is always optimal.

In both cases, the retailer's participation constraint is binding, so the retailer's profit is unaffected by the manufacturer's pricing decision, and the manufacturer simply seeks a (K', w') to induce the retailer to choose Δ and T that optimize the total system profit. Because the equilibrium solution maximizes system-wide profit and the retailer's profit is fixed, the manufacturer has no incentive to deviate from the stated solution. Thus, the manufacturer applies either (2.14) or (2.15), as appropriate, just once, not repeatedly.

The single-segment problem does not capture the impact of discounts in reaching a larger customer base, which is a key reason why retailers may use high-low pricing. We explore this issue through a computational study for scenarios with two segments. Here, too, we will see that it may be beneficial for the manufacturer to offer a transportation subsidy. The results also provide insights into the interactions between the transportation subsidy and the pricing policy.

2.6 Numerical Results

In this section, we explore characteristics of the manufacturer's optimal pricing scheme for an array of problems. We found equilibrium solutions for the following combinations of parameters which would be typical for a product sold at the drug store chain that motivated our work, but they also represent a range of scenarios that lead to varied optimal pricing decisions.: $r_1 = 5, r_2 = 3.5$ or 4.5, $\lambda_1 = 3000, \lambda_2 = 1500, 3000$ or 15000, $P = 20(\lambda_1 + \lambda_2)$, $K = 500$ and $w = 2.5$. We set $h_1 = 0.2r_1$ or $0.5r_1$, $h_2 = 0.2r_2$ or $0.5r_2$, $h_r = 0.5h_1, h_1$, or $2.5h_1$, and $h_m = 0.2h_r, 0.5h_r$ or h_r. We explain our parameter choices below. We consider both situations where the two segments have similar ($r_1 = 5, r_2 = 4.5$) and quite different ($r_1 = 5$ and $r_2 = 3.5$) reservation prices. Similarly, the number of customers in segment 2 may be much larger than in segment 1 (because typically, many more customers would be willing to pay a moderate price rather than a high price) but we also allow it to be smaller. The production rate is set so that the manufacturer could supply *all potential* customers of 20 retailers that are equivalent in size to the retailer in our model. (The actual demand from each retailer may be far less so the manufacturer could then supply many more retailers.) We chose a single base transportation cost, $K = 500$, which is roughly the cost of a short-haul trip of a few hours, which would be typical in our motivating scenario.

The customers' annual per unit holding costs correspond to annual inventory holding cost rates (percentages) of 20% and 50% assuming their purchase cost is equal to their reservation price. (The actual price may be less, depending upon the retailer's pricing policy.) These relatively high holding cost rates reflect the customers' tendency to stockpile less than what economic considerations (i.e., their opportunity cost of capital and opportunity cost of storage space) would suggest; this is consistent with empirical evidence on stockpiling behavior. The retailer's holding costs range from 50% of segment 1's holding cost up to 2.5 times segment 1's holding cost. At the low end of the range, the retailer prefers to hold inventory and avoid offering discounts, while at the high end of the spectrum, the retailer has a strong incentive to transfer inventory to the customers by offering discounts. We set the manufacturer's holding costs to be less than or equal to those of the retailer because manufacturer's storage facilities tend to be in lower-cost locations than retail distribution centers. These combinations yield a total of 216 problem instances. There are 6, 66, 126, and 18 instances that satisfy the conditions of Cases I, II, III and IV, respectively (see Table 3 for the conditions defining the cases). It turns out that for some of these cases, the ranges of parameter values satisfying the corresponding conditions on α and β are quite small. The reasons are as follows. The expression for α can be rewritten as $\frac{\lambda_1}{h_1}(1 - \frac{h_r}{2h_1}) + \frac{\lambda_2}{h_2}(1 - \frac{h_r}{2h_2})$. So, if β (i.e., the first parenthetical expression in the formula) is negative, the first term in the expression is negative and α is positive only if h_2 lies in a restricted range (holding other parameters constant). This is why only a few problem instances satisfy the conditions of Case I. (We generated additional problem instances satisfying the conditions of Case 1; their solutions have characteristics that are similar to those in our initial set.) The number of problem instances satisfying the conditions of Case IV are limited for similar reasons, but the conditions are not as restrictive as they are for Case I.

We found equilibria for the manufacturer's current pricing scheme ("before") and when he can choose K' and w' ("after"). We categorize the results according to the *retailer's* pricing strategies in the "before" and "after" scenarios. There are five pricing schemes: (i) $p = r_2, \Delta = 0$ (everyday low price—EDLP); (ii) $p = r_2, \Delta > 0$ (low price with discounts—LPD); (iii) $p = r_1, \Delta = 0$ (everyday high price—EDHP); (iv) $p = r_1, 0 < \Delta \leq r_1 - r_2$ (high price with shallow discounts—HPSD); and (v) $p = r_1, \Delta > r_1 - r_2$ (high price with deep discounts—HPDD). Table 4 shows the average profit improvement for the manufacturer and the number of instances for each before–after combination.

The manufacturer achieves substantial profit improvement when his new pricing policy reduces the retailer's operating costs enough to allow him to embrace Segment 2 as customers. This occurs, for example, when the retailer switches from EDHP to EDLP, EDHP to HPDD, HPSD to LPD or HPSD to HPDD. The best opportunities arise when (i) products have high price elasticity, where small changes in the price would bring in a relatively large incremental revenue (i.e., a large segment 2 with a reservation price fairly close to that of segment 1), or (2) when the retailer has high implicit holding costs—-either his own holding costs if he is holding most of the inventory, or the cost of the discounts to induce customers to stockpile to relieve him of the holding costs. In the latter situation, a transportation subsidy relieves the retailer of considerable holding costs.

	After				
Before	EDLP	LPD	EDHP	HPSD	HPDD
EDLP	0% (90)†	—	—	—	—
LPD	—	0% (15)	—	—	—
EDHP	36% (21)	—	0% (24)	—	39% (3)
HPSD	—	20% (3)	—	0% (12)	52% (9)
HPDD		—	—	—	1.6% (39)

† number of problem instances

Table 2.4: Manufacturer's Profit Increase from New Pricing Scheme for Various Changes in Retailer's Pricing Strategy

There are also some instances (see the bottom right cell in the table) where adjustments in a HPDD policy are beneficial to the manufacturer. In 3 of these 39 instances, the retailer's replenishment cycle is long and Segment 2 purchases (only at the discount price) to fulfill demand for only a small portion of the replenishment cycle. To induce Segment 2 to purchase more, the retailer would have to offer an extremely large discount. By offering a transportation subsidy, the manufacturer enables the retailer to reduce his replenishment cycle to the point where he can adjust the discount in such a way that he can entice Segment 2 to purchase to satisfy demand for a much larger portion of the replenishment cycle. For these 3 instances, the improvement in the manufacturer's profit averages 21%. These instances have relatively high values of h_2 so the customer's stockpiling duration is short. As such, changes in the cycle duration may have a strong impact on total sales to Segment 2.

Recall that the manufacturer chooses Pareto-improving K' and w', so the retailer is never worse off. Table 2.5 shows the improvement in the retailer's profit as a function of the price strategy change induced by the manufacturer. In the cells with positive changes, retailer's profit increases due to greater market penetration (more or all of Segment 2's demand). Some profit gains of the retailer are quite substantial and although they accrue to the retailer in our model, if retailers perceive that a manufacturer offers a price structure that allows retailers to earn more profit when operational costs such as transportation and inventory are considered, that manufacturer's products will be more competitive in the long term.

We now turn to an evaluation of consumer surplus. For several of the changes in the retailer's pricing strategy (EDHP to EDLP, EDHP to EDHP, EDHP to HPDD, and HPSD to HPDD), the surplus of each segment is guaranteed to improve or remain the same. Here, we discuss circumstances where there is potential for the net surplus to decline. For all instances where the retailer retained the LPD or HPSD policy, there was no change in the consumer surplus because the retailer did not change his prices. In all 39 instances where the retailer retained the HPDD policy, consumer surplus increased, sometimes manifold. Thus, although it is possible for consumer surplus to decline when the manufacturer induces the retailer to switch to a new pricing policy, this appears not to have a deleterious effect on consumers, and indeed, often makes them better off.

Thus far, we have focused on the effects of the optimal pricing policy on the retailer's and manufacturer's profits and consumer surplus. We next report on the extent to which the profit-optimizing solutions dampen two causes of the bullwhip effect: order batching and high-low pricing. A detailed examination of "before" and "after" solutions reveals that, when it is

	After				
Before	EDLP	LPD	EDHP	HPSD	HPDD
EDLP	0% (90)†	—	—	—	—
LPD	—	0% (15)	—	—	—
EDHP	45% (21)	—	0% (24)	—	60% (3)
HPSD	—	62% (3)	—	0% (12)	47% (9)
HPDD	—	—	—	—	0% (39)

† number of problem instances

Table 2.5: Retailer's Profit Increase for Various Changes in His Pricing Strategy Induced by the Manufacturer's Pricing Policy

optimal for the manufacturer to change K, two types of changes are common:
(1) T decreases, w decreases (from r_1 to r_2) and Δ remains unchanged (virtually always zero both before and after)—about 61% of the changes; and
(2) T decreases, w is unchanged and Δ increases—about 33% of the changes
Less common is the following type of change:
(3) T decreases, w decreases (from r_1 to r_2) and Δ decreases—about 6% of the changes.
Thus, in all cases where $K' < K$, T decreases, with a typical reduction of 30% in our set of numerical examples (details omitted). From the manufacturer's perspective, the retailer's replenishment interval is a key metric of the bullwhip effect—the impact of high-low pricing on the manufacturer is fully reflected in the retailer's replenishment interval. Thus, although Δ increases for instances in category (2) above, T decreases despite the increase in Δ because this change makes both parties better off. As such, there may be an increase in the degree of price fluctuations even when the bullwhip effect—as observed by the manufacturer—has declined. The reason why this occurs is that the price fluctuations at the retailer occur *within* the retailer's replenishment cycle and not from one replenishment cycle to the next. As such, care must be taken in claiming that price fluctuations exacerbate the bullwhip effect. Instead, a change in the manufacturer's pricing policy that induces the retailer to make changes in either pricing or operational policies that ultimately lead to shorter retailer replenishment cycles will dampen the bullwhip effect, even if intermediate symptoms such as price fluctuations appear to worsen.

Thus far, we have focused on problem in stances in which it is optimal for the manufacturer to change K. There are three other instances in which it is optimal for the manufacturer to keep K constant but to reduce w. In these instances, the reduction in w allows the retailer to sell to both segments profitably so the retailer lowers his regular price. Furthermore and interestingly, the retailer also increases Δ and by doing so, shifts the inventory burden almost entirely to the customers. Consequently, the retailer also *increases* T (by about 30%). Thus, the bullwhip effect (as measured by T) increases, but the manufacturer's profitability increases (by about 80% in these instances) because two segments are purchasing instead of one. Thus, if the manufacturer seeks to optimize K and w, he may keep K constant and reduce w. In such cases, it is obvious that the bullwhip effect will worsen. If the manufacturer is concerned primarily about the bullwhip effect, he can simply retain his current pricing scheme. However, by doing so, he may be foregoing a substantial increase in profit and market share.

To summarize, in about 80% of our problem instances, K and w do not change when the manufacturer attempts to optimize them because of his own participation constraint. However, when he stands to benefit from a change, it is quite likely that the bullwhip effect—as reflected in the retailer's replenishment interval—decreases and the profitability of the supply chain may increase substantially, with the vast majority of the benefit accruing to the manufacturer. In many instances, this occurs even when retail price fluctuations become worse. We note that our numerical results are based on values of K that are consistent with the cost of fairly short trips of 200 miles or less (as in our motivating example). Thus, the manufacturer's latitude for subsidizing transportation is relatively limited. When the initial K is larger, we would expect a greater percentage of cases in which a transportation subsidy would be advantageous

2.7 Conclusions

We have analyzed a three-stage supply chain consisting of a manufacturer, retailer, and two segments of end-consumers who may differ in their reservation prices and their propensity to stockpile. The retailer incurs a transportation (setup) cost for each order as well as inventory holding costs. He therefore orders in batches, which gives him an incentive to offer discounts to clear his inventory more rapidly. The retailer decides regular and discount prices and his replenishment cycle.

We analyze a three-stage Stackelberg game in which the manufacturer can subsidize the retailer's transportation cost and adjust the wholesale price. The manufacturer earns revenue from the units sold at his selected wholesale price and incurs holding costs while accumulating inventory prior to each shipment and his share of the transportation costs. The manufacturer and retailer seek to optimize their respective profits. We fully characterize the retailer's optimal response to the manufacturer's pricing policy, which provides clear insights on how the transportation subsidy and adjusted wholesale price affect the retailer's pricing strategy and replenishment cycle.

The manufacturer stands to gain the most when, under the current costs the retailer finds it unprofitable to sell to the segment with the lower reservation price. By subsidizing the transportation cost, the manufacturer makes it possible for the retailer to increase sales to the segment with the lower reservation price without having to offer a large discount. This occurs because, when the retailer's replenishment cycles are short, these customers may be willing to stockpile to satisfy demand for most or all of the replenishment cycle even when offered a modest discount, whereas they are unwilling to stockpile to the same extent as a fraction of the replenishment cycle when the cycle is long. The more modest discount also leads to greater revenue from the segment of customers with the higher reservation price because they purchase a greater percentage of their demand at full price. These two effects can lead to a marked increase in profits. Thus, the impact of a transportation subsidy is complex, indirect, and not intuitive at the outset. Our numerical results also show that the transportation subsidies, when it is optimal to offer them, have a direct effect on the bullwhip effect—as reflected in the retailer's replenishment interval—even when the base transportation cost is relatively small, and often even when they lead to greater fluctuations in the retail price.

Although we have considered deterministic demand here, when demand is stochastic, shorter replenishment cycles contribute to reducing the adverse effects of demand signal processing (another cause of the bullwhip effect) because orders are transmitted more frequently, thereby reducing the horizon over which the retailer needs to forecast demand. Because near-term forecasts tend to be more accurate than forecasts of demand father into the future, the magnitude of the forecast error decreases more than proportionally as the replenishment cycle decreases.

Manufacturers can increase their market penetration by reducing the total cost incurred by a retailer for offering their products. This includes not just the wholesale price but also logistics costs, including the cost of holding inventory. Retailers may bear the inventory holding costs directly, or they may do so indirectly via revenue lost when they offer discounts to clear inventory more quickly. Large retailers such as Wal-Mart have pressured their suppliers to reduce logistics-related costs borne by the retailer, but manufacturers may benefit by finding new ways to reduce these costs. Here, we have considered one method that relies on a transportation cost subsidy, but the transportation subsidy also has other benefits for the retailer. Further research is needed to explore other means of reducing retailer-borne logistics costs while considering the array of other decisions made in a supply chain, such as pricing and replenishment decisions, and their complex interactions.

Chapter 3

Retailer's Optimal Pass-Through of Manufacturer Trade Discounts When Retail Discounts Affect Reservation Prices and Stockpiling

3.1 Introduction

Manufacturers offer discounts, known as trade discounts, to retailers, hoping that retailers will pass the savings on to their customers, thereby increasing the manufacturer's market share and possibly also increase his profit. Empirical research (e.g., Armstrong 1991 and Besanko et al. 2005) suggests, however, that manufacturers rarely pass through the entire discount to customers, and sometimes do not pass through any of it at all.

In this chapter, we focus on low- to moderately-priced, discretionary, consumable product that customers will purchase if the price is "low enough," but for which they will forego consumption (or consume an alternate default product) otherwise. For such products, retail discounts have a strong impact on demand. We study the retailer's problem of determining his optimal pass–through strategy when intermittent retail price discounts affect customer's reservation prices and may lead them to stockpile. Intermittent price discounts are known to have an adverse effect on reservation prices (or synonymously, willingness-to-pay)—we refer to this as the (negative) brand equity effect— but they allow retailers to reach customers who are unwilling to buy at the regular price. Furthermore, customers—both those who purchase regularly and those who purchase when periodic price discounts are offered—may stockpile, which serves to reduce the retailer's costs by shifting the cost of holding inventory from the retailer to customers. When a manufacturer offers a trade discount, the retailer may choose to offer even deeper discounts, which will magnify both the beneficial and adverse effects mentioned above. We seek a strategy for the retailer that achieves the best tradeoff. We also discuss implications for manufacturer's choice of trade discounts in view of the retailer's optimal reaction to these discounts.

We show that, not surprisingly, considering the brand equity effect causes the retailer to discount less than he would when considering customer stockpiling alone. What is surprising, however, is that the brand equity effect sometimes causes the retailer to change the structure of his discounting policy vis-a-vis situations with customer stockpiling alone. We also show that there may be a threshold effect: in some circumstances, if the trade discount does not exceed a threshold, the retailer does not pass through any of the discount, and above the threshold, the retailer passes through a portion of the trade discount. The threshold effect arises when the retailer discounts affect brand equity but does not arise when only stockpiling behavior is considered. Furthermore, rarely does the retailer pass on a substantial portion of the discount. The retailer's reaction to a trade discount depends, in part, on the retailer's (optimal) discounting strategy in the absence of trade discounts. Consequently, in the presence of two common phenomena—-negative brand equity effects due to periodic retail discounts and stockpiling by customers when a discount is offered—manufacturers need to understand the retailer's discounting strategy in the absence of a trade discount and to evaluate how large a trade discount will be needed to induce a reaction on the part of the retailer that is profitable for the manufacturer.

The remainder of this chapter is organized as follows. In Section 3.2, we provide a review of the literature. In Section 3.3, we provide a formal statement of the problem and formulations of the customers' problem of when and how much to purchase and a general version of the retailer's problem of choosing a discounting and ordering strategy. In Section 3.4, we derive more detailed results and a solution procedure for the retailer's problem under the assumption of a uniform distribution of customer reservation prices. Section 3.5 presents the results of an extensive numerical study that examines how the retailer's pass–through changes with the manufacturer's trade discount, and characteristics of situations that lead to high or low pass-through rates. Conclusions appear in Section 3.6.

3.2 Literature Review

We review the literature on several topics that are closely related to our study: (i) formation of customers' reference prices and how reference prices affect purchasing decisions, (ii) consumers' reaction to promotions (such as stockpiling and brand switching), (iii) how retail price promotions affect immediate and long-term demand, and (iv) retailers' optimal pass–through of manufacturers' trade deals.

3.2.1 Reference Prices

Reference price is the standard against which consumers evaluate purchase prices (Monroe 1973); a positive difference between the reference price and the purchase price is considered a gain and a negative difference is considered a loss (Winer 1986). The concept of reference prices stems from adaptation-level theory (Helson 1964) which holds that people judge stimuli based on standards shaped by their prior exposure to similar stimuli. A substantial amount of empirical work has included reference prices in consumer-choice models and their inclusion generally improves the predictive ability of the models (see, for example, Guadagni and Little 1983). We discuss the literature on reference price formation and its effect on purchasing behavior.

Conceptually, many researchers regard reference price as being a function of some type of weighted average of past prices that consumers have observed, plus perhaps other factors, but different mathematical representations have been used. Lattin and Bucklin (1989) develop a model of reference prices that includes two effects: (i) the effect of past prices, represented as an exponentially weighted average of these prices and (ii) promotion reference effects, by which they mean the disparity between consumers' expectation of whether a promotion will be offered in the store and the actual situation they observe. By combining these effects, they significantly improve the predictive performance of their brand choice model over their baseline model without reference price effects. Mayhew and Winer (1992) utilize both internal and external reference prices in discrete choice models to estimate purchase probabilities. Internal reference prices are based on customer's memories of actual prices or other internalized price concepts, while external reference prices are observed prices, such as displayed "regular prices." In their model, the former is represented by the price paid by the consumer at the last purchase occasion and the external reference price is represented by the displayed regular and reduced prices. Their results show that both types of reference prices play significant roles in explaining purchase probabilities.

Other researchers have extended reference price models to include other factors. Winer (1986) includes posted price, price trend and market share in his model to estimate reference prices. Using data on coffee purchases, Winer finds that the reference price effect is significant and the predictive performance of his brand choice model with the additional predictors of reference price is substantially better than the model without these factors. Kalwani et al. (1990) include frequency of sales promotions and price trends over time. They report that frequent promotions lower consumers' reference prices, more deal-prone customers have lower reference prices, and people expect to pay different prices at different stores. They also find that consumers have asymmetric responses to perceived gains and losses vis-a-vis their reference price (i.e., sensitivity of demand to a price loss is greater than that to a price gain).

Briesch et al. (1997) develop a model to estimate consumers' utility for various brands that includes some contextual factors (i.e., whether the brand is featured or displayed), and households' loyalty toward each brand along with reference prices. They evaluate five different price representations: (i) current price of a randomly chosen brand, (ii) current price of the last-chosen brand, (iii) exponentially-weighted average of past retail prices of brands chosen by a household, (iv) exponentially-weighted average of a brand's own retail prices, and (v) the combination of a brand's immediate past price, price trend, deal frequency, deal proneness of the household. Their estimates of utility are used in their multinomial logit model to predict the probability that a particular brand will be chosen by a household at each purchase occasion. The authors conclude that the model in which reference price is represented as an exponentially weighted average of a brand's own retail prices appears to be the best overall.

Another stream of work investigates how reference prices affect consumers' purchasing decisions, such as brand choice, purchase quantity, and purchase timing decisions. In general, a positive difference between the reference price and the purchase price increases the customer's willingness to pay for the item (brand), and a negative difference reduces it (Winer 1986, Kalyanaram and Winer 1995, Bell and Lattin 2000). Krishnamurthi et al. (1992) find that a positive difference between reference price and observed price has a positive effect on consumers' brand-choice and purchase-quantity decisions for both loyal customers and switchers. Bell and Bucklin (1999) find that consumers stockpile if they perceive a price gain (i.e., a positive difference between the reference price and the actual price) and postpone a purchase if they perceive a loss (negative difference between the reference price and the observed price). They also observed that the effect of reference prices on purchase quantity and purchase probability is moderated by household inventory positions (Krishnamurthi et al. 1992, Bell and Bucklin 1999).

3.2.2 Effects of Price Promotions

In the literature, various impacts of price promotions on consumer purchasing decisions, such as brand choice, stockpiling and consumption acceleration, have been reported. Researchers report that brand switching constitutes the main effect, followed by stockpiling (Gupta 1988 and Bell et al. 1999). Our model is a single-product model, so we do not discuss the literature on brand choice here. We discuss literature on the two other effects next, and then turn to the literature on the long-term effects of price promotions on sales volume.

One stream of empirical research examines the effect of promotions on consumers' stockpiling behavior. Beasley (1998) shows that consumers' stockpiling decisions depend on their household deal-proneness, inventory level and the depth of discount. Meyer and Assuncao (1990, 1993) find that the the observed price of the good, the distribution of future prices and the consumers' inventory level affect the incidence of consumers' stockpiling behavior. A study by Aggarwal and Vaidyanathan (1993) suggests that short-term promotions encourage stockpiling but long-term promotions (for example, manufacturer's coupons) do not. The study by Mela et al. (1998) indicates that increasing promotional activity over the years has cultivated consumers' tendency to purchase more during promotions and less during non-promotional periods.

In another stream of literature, researchers report mixed results on whether promotions lead to accelerated consumption. Chandon and Wansink (2002) report that stockpiling increases consumption for both high- and low-convenience products. Ailawadi and Neslin (1998), Bell et al. (2002) and Ailawadi et al. (2007) report that stockpiling causes an increase in consumption in some product categories, but not in others. Researchers have also investigated whether consumers' stockpiling during promotions is offset by lower purchases after promotions. Hendel and Nevo (2003) report evidence that post-promotion dips do occur.

Researchers have investigated the long-term effects of promotions and conclude that promotions have essentially no persistent impact on sales volume in a mature market. Dekimpe et al. (1999) utilized impulse-response functions to analyze the effects of promotions and found that both category and brand sales are stable, and demand returns to a baseline level after promotional effects fade away in the catsup, liquid detergent and yogurt markets. Their results show long-term effects on promotions on sales in the soup category, but with relatively low persistence. Nijs et al. (2001) studied the effects of consumer price promotions on category demands across 560 consumer product categories and found little persistent or long-term effects of promotions. Their results indicate that price-promotion effects typically last for about 10 weeks and the long-term impact converges to zero in 98% of the 560 product categories. Persistence models have their methodological roots in econometrics and time series analysis and have been used to study the long-run effects of various marketing activities on market performance (see Dekimpe and Hassens 2005). Pauwels et al. (2002) use persistence models to capture the long-term effects of promotions and conclude that promotion effects are virtually absent. In only 1 out of 29 cases did they detect a permanent promotion effect. Differences in persistence of price promotions across different product categories and brands may be due to product characteristics: mature markets are less likely to be permanently affected by marketing actions because consumers have become habituated to the usual promotional patterns (Bronnenberg et al. 2000).

3.2.3 Retailer's Pass–Through Rate

The literature includes both analytical and empirical studies of the retailer's pass-through of manufacturer's promotions. In empirical studies, Chevalier and Curhan (1976), Walters (1989) and Armstrong (1991) all find substantial variations in pass-through rates. They find not only values less than 100%, as expected, but also rather surprising values greater than 100%.

In more recent empirical studies, Besanko et al. (2005) report on pass–through behavior of a major U.S. supermarket chain for 78 products across 11 categories and find that, on average, the pass–through rates are more than 60% for 9 out of 11 categories, but the rate is as low as 22% in the toothpaste category and as high as 558% in the beer category. They observe that brands with larger market shares or that contribute more to retailer profits have higher pass–through rates. Nijs et al. (2010) investigate pass-through from wholesaler to retailer and retailer to consumer, and observe large variances in pass-through at each level. They find retail pass-through to be positively correlated with retailer size, price elasticity, and negatively correlated with package size and wholesaler's deal frequency. Meza and Sudhir (2006) explore variations in the retailer's pass-through over time, which may be important for manufacturers who produce products that have strong seasonality in demand. They find that during regular demand seasons, only loss-leaders receive high pass-through, and during peak-demand seasons, both loss-leaders and regular products receive a considerable amount of pass-through. Ailawadi and Harlam (2009) investigate the magnitude of, and drivers of the differences in, the retailer's pass-through. Their results indicate that private labels, manufacturers with high market share, and categories with certain characteristics (high sales volume, high promotion elasticity, low margin, and low concentration) typically receive larger pass-through. Pauwels (2006) finds that retailers offer higher pass-through for leading brands and high-revenue categories. His findings also indicate that when a manufacturer offers a trade discount, its competitors also reduce their wholesale prices, and the retailer adjusts competing brands' retail prices, thus the effectiveness of a manufacturer's promotion is reduced.

Some researchers have studied the determinants of the retailer's pass-through rate using analytical models and suggest that the pass–through rate depends on characteristics of the consumers' purchasing behavior, such as consumers' willingness-to-pay and propensity to switch brands or stores, and the performance of the product, such as its profit margin and demand at the regular price. Tyagi (1999) offers differences in responsiveness of a firm's marginal revenue to a change in its price as an explanation for differences in pass–through rates. A profit-maximizing retailer equates marginal revenue with marginal cost, so when he receives a discount from the manufacturer, he adjusts his retail price to match the marginal revenue with the reduced marginal cost. If the retailer's marginal revenue is very responsive (or not very responsive) to a change in retail price, the retailer must reduce its price by less (more) than the amount of the reduction in his marginal cost.

Kumar et al. (2001) suggest that the retailer offers a lower pass–through rate when many customers are willing to pay the regular price and when they have a high cost of searching for deals elsewhere. Kim and Staelin (1999) analyze the retailer's pass–through rate in a competitive environment by including the effects of brand switching, store switching and category expansion (i.e., new demand from consumers who did not purchase products in this category previously) on demand. Comparative statics results indicate that the pass–through rate increases with the customers' propensity to switch stores or brands. Besanko et al. (2005) and Moorthy (2005) examine both own-brand pass–through (i.e., the pass–through of a manufacturer's discount as a retail discount on his own brand) and cross-brand pass–through (i.e., reduction in the retail price of competing brand(s) due to a manufacturer's discount) and find both negative and positive cross-brand pass–through in addition to significant own-brand pass–through. Their research generally indicates that brands with larger profit margins have higher pass–through rates and are more likely to have a positive pass–through.

To the best of our knowledge, there is no research that considers (negative) brand equity effects when analyzing the retailer's optimal pass–through of trade discounts. Most of the previous work on the effects of retail price discounting focuses on how it hurts the manufacturer's long-term profitability. Here, we incorporate its effects into the retailer's price discounting decisions and explore how this affects the pass–through of manufacturer's trade discounts—and thus, also, how the manufacturer should structure trade discounts.

3.3 Problem Statement and Formulation

We model a retail channel that consists of a single manufacturer, a single retailer and heterogeneous consumers. In our model, the retailer makes decisions about temporary retail discounts as well as a procurement plan (i.e., how much to order and when). Grocery chains and other similar retailers typically make these decisions on a periodic (e.g., weekly) basis, so we use a discrete time framework in which the basic time unit is chosen consistently with the periodicity of the discounting and procurement decisions. We assume that the basic time unit is determined by institutional arrangements that exist due to historical, competitive, and/or practical considerations, and therefore will not be influenced by small to modest changes in the retailer's discounting decisions. The retailer's pass–through rate is implicit in his discounting decisions. We elaborate on the retailer's problem later in this section.

The manufacturer offers a constant regular wholesale price, w, but offers a per unit discount of Δ^M every T periods. We assume that w and T are given, but the analysis can be performed for any desired values of w and T. We initially assume that Δ^M is given, but later in the chapter, we examine how its value affects the retailer's decisions and the consequent pass–through rate.

The total number of potential customers in the market is Γ and each customer consumes one unit of the product per period if he has a unit available. Customers are heterogeneous with respect to their reservation prices. The distribution of the customers' reservation prices in the absence of retail discounts is denoted by $F(\cdot)$ which we assume is continuous and differentiable; the density is $f(\cdot)$. We assume that if the retailer offers discounts, the entire distribution of reservation prices shifts to the left (downward) by an increment that depends upon the retailer's discounting pattern. The new distribution function of the customers' reservation prices is denoted by $\tilde{F}(\cdot)$ which we refer to as the *adjusted reservation price distribution*; $\tilde{f}(\cdot)$ is the p.d.f. We explain how the customer's reservation prices are affected by the retailer's discounts in Section 3.3.1.

In our model, all customers have the same holding cost, h per unit per period, which is incurred on average inventory. Each customer consumes one unit of the product per unit time if he has access to a unit whose *gross cost*, i.e., the actual purchase cost plus the cost of holding inventory from purchase to consumption, does not exceed his reservation price. If a customer does not have access to such a unit, he foregoes consumption of the product. When the retailer offers a discount, customers may stockpile. So at various points in time, customers may be depleting their stockpile, they may purchase one unit per period and consume it immediately, or they may not be consuming at all (if the gross cost is too high). We assume that the presence of customer stockpiles (where applicable) do not change the basic consumption rate. We also assume that customers shop frequently enough to take advantage of all discounts if they so choose. Many consumers purchase groceries frequently (e.g., once a week) and most retail chains hold prices for non-perishable goods constant for a week (although they may differ from week to week), so our assumption allows us to capture the forward-buying behavior of customers reasonably accurately.

The retailer pays a fixed transportation cost K for each replenishment and incurs an inventory holding cost of h_r per unit per period (incurred on average inventory). We assume that the regular price, p, is fixed. The retailer chooses the timing and magnitude (or synonymously, depth) of retail discounts, and his procurement policy to optimize his profit (revenue less the sum of variable unit costs, transportation and inventory holding costs) per unit time, taking into account the customers' response. We assume the retailer's discounting and replenishment pattern repeats after T periods, the periodicity of the manufacturer's discount. Within the T periods are two types of replenishment cycles. One type of cycle, which we call a *high* cycle, has a duration denoted by T^H and starts when the manufacturer offers a discount and the retailer places an order at the discounted price. The other type of cycle, which we call a *low* cycle, is a regular replenishment cycle with a duration denoted by T^L in which no discount is offered by the manufacturer, but the retailer may choose to offer a discount. There may be one or more low cycles following each high cycle, so we have $T = T^H + NT^L$ for some integer $N \geq 0$.

We assume that the retailer offers a single discount ($\Delta^H \geq 0$ in high cycles and $\Delta_i^L \geq 0$ in the i^{th} low cycle after each high cycle, $\forall i \in \{1, ..., N\}$) at the beginning of each of his replenishment cycles (immediately after receipt of a shipment from the manufacturer). It can be shown that if the retailer offers a discount only once during each of his replenishment intervals, it is optimal to schedule the discount offering at the beginning of his replenishment cycle because this timing induces customers to purchase a stockpile immediately and thereby leads to the greatest reduction in the retailer's inventory holding costs. By offering a discount, retailer not only clears his inventory more quickly, but he also attracts customers who would not purchase at the regular price. On the other hand, he sacrifices some profit because the customers' reservation prices decline. This is the fundamental tradeoff that the retailer faces.

We analyze the problem as a Stackelberg game with the retailer as the leader. In the next subsection, we analyze the customer's problem of when and how much to buy for a given discounting strategy chosen by the retailer.

3.3.1 Customer's Problem

Given the fixed regular price p, the retailer's decisions regarding cycle lengths T^H and T^L and discount magnitudes Δ^H and Δ_i^L, $i \in \{1 \dots N\}$, as well as his own parameters, the customer seeks to maximize his net utility (utility from consumption less purchase costs and inventory holding costs), where the utility is equal to his reservation price. Krishna and Johar (1996) suggest that the average of the prices offered by the retailer (across all periods) is a good estimate of the customer's reservation price. Their findings suggest a simple model of the impact of discounts on reservation prices in which each customer's reservation price declines from his original reservation price by mT, which is the average depth of discount offered over the T-period cycle. (We also refer to m as the *loss of brand equity* due to retailer discounting.) Other researchers have proposed and tested other functional forms to capture the impact of discounts on reservation prices and more generally, on the propensity of the customer to purchase. Many of these functional forms are quite complicated, so we elected to use a simple model that captures the first-order effects of retail discounts on reservation prices. Letting R denote the random variable for the adjusted reservation price and r the observed value, the distribution of R is $\tilde{F}(r) = F(r + m)$, and the density is $\tilde{f}(r) = \tilde{F}'(r)$.

Given each customer's adjusted reservation price, he will not purchase if the *gross cost* of a unit, i.e., the unit purchase cost plus the cost of holding the unit from purchase to consumption, exceeds his reservation price. Customers behave strategically and may stockpile when the retailer offers a discount. A customer whose reservation price r is higher than the regular price p, referred to as a *regular* customer, will stockpile to satisfy his consumption for a duration Δ^C/h, where Δ^C is the discount offered in the current cycle, at which point he is indifferent between (i) a unit purchased at the beginning of the replenishment cycle and held until it is consumed and (ii) a unit purchased at the regular price for immediate consumption.

However, if Δ/h exceeds the duration of the current cycle, consumers may stockpile to satisfy demand only up to the end of the cycle if the retailer offers another discount at the beginning of the next cycle. Let T^C, Δ^C and Δ^N represent the duration of the current cycle, the discount in the current cycle and the discount in the next cycle, respectively. Then consumers do not stockpile for consumption in the next cycle if $\Delta^C - \Delta^N < hT^C$, i.e., if the incremental discount offered in the current cycle is not large enough to offset the incremental holding cost due to an early purchase. Although there are instances in which consumers would be willing to stockpile to satisfy demand for longer than T^C, empirical research (e.g., Meyer and Assancao 1990) indicates that consumers stockpile less than their economic tradeoffs would dictate. Possible reasons for this behavior include storage limitations, cash flow limitations, etc. As such, we assume that discounts satisfy $\Delta^C - \Delta^N < hT^C$ and leave other cases for future research. (Our preliminary analysis indicates that these other cases are quite complicated and it is difficult to obtain any insights from them.)

Similarly, customers whose reservation price r is below the regular price but above the discount price, referred to as *discount* customers, will stockpile to satisfy their consumption for a duration of $[\Delta^C - (p-r)]/h$, which is always less than T^C if $\Delta^C - \Delta^N < hT^C$. Unlike the regular customers, after their stockpiles are depleted, the discount customers forego consumption until the next time the discount price is below their respective reservation prices.

3.3.2 Retailer's Problem

The manufacturer offers a trade discounts every T periods, and within each such T-period manufacturer-discount cycle, the retailer has a high cycle in which the retailer receives a discount from the manufacturer, and N low cycles in which the retailer does not receive a discount from the manufacturer. (One can think of "high" and "low" as characterizing the retailer's likely order quantities.) Let Π^H and Π_i^L represent the *total profit* in the high cycle and the i^{th} low cycle, respectively. The retailer's objective is to maximize his profit per unit time:

$$\Pi = (\Pi^H + \sum_{i=1}^{N} \Pi_i^L)/T \qquad if \;\; N \geq 1$$
$$\Pi = \Pi^H/T \qquad if \;\; N = 0.$$

Recall that the retailer needs to choose Δ^H, Δ_i^L, T^H, T^L and N. The values of T^H, T^L and N are integral, but we treat the Δs as continuous variables.

In the remainder of this section, we derive the components of the retailer's objective function. We focus on the retailer's total profit in a high cycle first; the profit in low cycles can be derived in a similar fashion.

If the retailer offers a discount $\Delta^H \leq hT^H$ at the beginning of a high cycle, then each *regular* customer will stockpile to satisfy consumption for a duration $\frac{\Delta^H}{h}$. Similarly, the *discount* customers will stockpile to satisfy consumption for a duration $[\Delta^H - (p-r)]/h$. Hence the total quantity purchased at a discount at the beginning of a high cycle, denoted by d_1^H, is

$$d_1^H = \Gamma[(1 - \tilde{F}(p))\frac{\Delta^H}{h} + \int_{p-\Delta^H}^{p} \frac{\Delta^H - (p-r)}{h} \tilde{f}(r)dr] \tag{3.1}$$

The first term is the number of customers with a reservation price higher than p multiplied by the quantity that each customer purchases at a discount of Δ^H. The second term is the total quantity stockpiled by customers whose reservation price is between $p - \Delta^H$ and p; it accounts for the fact that each such customer stockpiles a quantity that depends upon his reservation price, r.

From the above, we can express the total quantity purchased at a discount at the beginning of a high cycle as:

$$\tilde{d}_1^H = \Gamma[(1 - \tilde{F}(p))T^H + \int_{p-\Delta^H}^{p} \min\{\frac{\Delta^H - (p-r)}{h}, T^H\}\tilde{f}(r)dr] \tag{3.2}$$

Customers with reservation prices higher than p start to purchase one unit just-in-time when their stockpiles are depleted. Therefore, the demand during the remainder of a high cycle, denoted by d_2^H can be expressed as:

$$d_2^H = \Gamma \tilde{F}(p)(T^H - \frac{\Delta^H}{h})^+. \tag{3.3}$$

The retailer's profit in a high cycle is equal to the total revenue from stockpiled quantities and from demand from regular customers during the remainder of the cycle, less variable unit costs, inventory holding costs and the transportation cost per order. Therefore, the retailer's total profit is:

$$\Pi^H(\Delta^H, \Delta_i^L, T^H, T^L, N) = (p - w - \Delta^H + \Delta^m)d_1^H + (p - w + \Delta^m)d_2^H - hrT^H d_2^H/2 - K. \tag{3.4}$$

where d_1^H is defined in (3.1) and d_2^H is defined in (3.3).

Total profit in the i^{th} low cycle, Π_i^L, is analogous to that of the high cycle. So the total quantity purchased at a discount of Δ_i^L at the beginning of the i^{th} low cycle is

$$d_{1i}^L = [\frac{\Delta_i^L}{h}(1 - \tilde{F}(p)) + \int_{p-\Delta_i^L}^{p} \frac{\Delta_i^L - (p-r)}{h} \tilde{f}(r)dr]\Gamma. \tag{3.5}$$

The total demand during the remainder of a low cycle is

$$d_{2i}^L = (T^L - \frac{\Delta_i^L}{h})^+ \tilde{F}(p)\Gamma. \tag{3.6}$$

Therefore, the total profit in the i^{th} low cycle is:

$$\Pi^L(\Delta^H, \Delta_i^L, T^H, T^L, N) = (p - w - \Delta_i^L)d_{1i}^L + (p - w)d_{2i}^L - hrT^L d_{2i}^L/2 - K, \tag{3.7}$$

where d_{1i}^L is defined in (3.5) and d_{2i}^L is defined in (3.6).

In summary, the retailer seeks to maximize his profit per unit time:

$$\Pi = [\Pi^H(\Delta^H, \Delta_i^L, T^H, T^L, N) + \sum_{i=1}^{N} \Pi^L(\Delta^H, \Delta_i^L, T^H, T^L, N)]/T.$$

3.3.2.1 Special Case: No Regular Customers

In this subsection, we analyze the special case in which the maximum adjusted reservation price, $\overline{R} - m$, is below the regular price, so no customers purchase at the regular price.

The derivations of retailer's demand and profit rate parallel those above. The retailer's demand at the beginning of high and low cycles are, respectively:

$$d_1^H = \int_{p-\Delta^H}^{\overline{R}-m} \frac{\Delta^H - (p-r)}{h} \tilde{f}(r)dr] \tag{3.8}$$

$$d_{1i}^L = \int_{p-\Delta_i^L}^{\overline{R}-m} \frac{\Delta^L - (p-r)}{h} \tilde{f}(r)dr] \tag{3.9}$$

There is no demand at the regular price and the retailer does not need to hold any inventory, hence the retailer's revenue only comes from sales at a discount. The retailer's total profit in a high and a low cycle are, respectively:

$$\Pi^H(\Delta^H, \Delta_i^L, T^H, T^L, N) = (p - w - \Delta^H + \Delta^m)d_1^H - K \tag{3.10}$$
$$\Pi^L(\Delta^H, \Delta_i^L, T^H, T^L, N) = (p - w - \Delta_i^L)d_{1i}^L - K \tag{3.11}$$

The retailer maximizes his total profit per unit time:

$$\Pi = [\Pi^H(\Delta^H, \Delta_i^L, T^H, T^L, N) + \sum_{i=1}^{N} \Pi^L(\Delta^H, \Delta_i^L, T^H, T^L, N)]/T$$

.

In the next section, we show how to solve the retailer's problem for fixed T^H, T^L and N assuming (for analytic tractability) that the distribution of customers' reservation prices is uniform. With this, it is straightforward to enumerate combinations of T^H, T^L and N satisfying $T^H + NT^L = T$, to solve the retailer's discounting and procurement problem for each, and to find the combination that leads to the best objective value. By enumerating different values of T^H, in particular, we are able to explore the interplay between the manufacturer's depth of trade discount and the retailer's propensity to forward-buy, and their combined effect on the retailer's optimal discounting policy. We defer the discussion of the T^H, T^L, N values to Section 3.5. We first discuss the the general case in which there are regular customers and then briefly discuss differences for the special case with no regular customers.

3.4 Retailer's Problem For Uniformly Distributed Reservation Prices

To gain some insight into the qualitative structure of the retailer's optimal strategy, in this section, we analyze the problem under the assumption that the customers' reservation prices follow a Uniform distribution on $[0, R]$. Note that the main effect of the "shape" of the distribution of reservation prices is on the partitioning of customers among three categories: (i) those whose reservation prices exceed the regular price; (ii) those whose reservation prices are between the discount price and the regular price; and (iii) those whose reservation price is below the discount price (and who will never buy). For different reservation price distributions, the relative sizes of these groups will differ, but the qualitative effects of the retailer's decisions remain the same irrespective of the distribution of reservation prices. Under the distributional assumption stated above, the expressions for demands in (3.1), (3.5) and (3.6) can be simplified to:

$$\begin{aligned}
d_1^H &= \Delta^H[(\overline{R} - p - m)^+ + \Delta^H/2]\frac{\Gamma}{h\overline{R}}, && (3.12)\\
d_{1i}^L &= \Delta_i^L[(\overline{R} - p - m)^+ + \Delta_i^L/2]\frac{\Gamma}{h\overline{R}}, && (3.13)\\
d_{2i}^L &= (hT^L - \Delta_i^L)(\overline{R} - p - m)^+\frac{\Gamma}{h\overline{R}}. && (3.14)
\end{aligned}$$

The retailer's profit functions in high and low cycles, respectively, can be rewritten as:

$$\begin{aligned}
\Pi^H(\Delta^H, \Delta_i^L, T^H, T^L, N) &= a^H(\Delta^H)^3 + b^H(\Delta^H)^2 + c^H\Delta^H + \gamma^H, && (3.15)\\
\Pi^L(\Delta^H, \Delta_i^L, T^H, T^L, N) &= a^L(\Delta_i^L)^3 + b^L(\Delta_i^L)^2 + c^H\Delta_i^L + \gamma^L, && (3.16)
\end{aligned}$$

where

$$
\begin{aligned}
a^H &= a^L = -0.5\frac{\Gamma}{h\overline{\overline{R}}},\\
b^H &= [0.5(p-w+\Delta^M) - (\overline{R}-p-m)^+]\frac{\Gamma}{h\overline{\overline{R}}},\\
c^H &= 0.5h_rT^H(\overline{R}-p-m)^+\frac{\Gamma}{h\overline{\overline{R}}},\\
\gamma^H &= (p-w+\Delta^M-0.5h_rT^H)hT^H(\overline{R}-p-m)^+\frac{\Gamma}{h\overline{\overline{R}}} - K,
\end{aligned}
$$

$$
\begin{aligned}
b^L &= [0.5(p-w) - (\overline{R}-p-m)^+]\frac{\Gamma}{h\overline{\overline{R}}},\\
c^L &= 0.5h_rT^L(\overline{R}-p-m)^+,\\
\gamma^L &= (p-w-0.5h_rT^L)hT^L(\overline{R}-p-m)^+\frac{\Gamma}{h\overline{\overline{R}}} - K
\end{aligned}
$$

In the remainder of this section, we provide relevant structural results and a sketch of the solution procedure, deferring details to the Appendices.

3.4.1 Optimizing the Δ Values

The optimization problem is:

$$
\begin{aligned}
&\max_{\Delta^H,\Delta_1^L,\ldots,\Delta_N^L} \Pi = \Pi^H(\Delta^H) + \sum_{i=1}^{N}\Pi^L(\Delta_i^L)\\
&\text{subject to } 0 \le \Delta^H \le hT^H\\
&\qquad\qquad\quad 0 \le \Delta_i^L \le hT^L,\ i = 1,\ldots,N
\end{aligned}
$$

The constraints ensure that customers do not stockpile more than they will consume before the beginning of the retailer's next replenishment cycle, which is consistent with our earlier assumptions.

Finding the optimal solution is complicated by the fact that the discount in each cycle has an effect on brand equity and hence affects the profit in all cycles. Consequently, the retailer's objective is not always jointly concave in the Δ values. To deal with this complication, we solve the problem using a nested optimization approach. In the "inner" optimization, we optimize Δ^H and the Δ_i^Ls for a fixed (negative) brand equity effect, m. To do so, we first obtain results that allow us to jointly optimize the Δ_i^L values. These results allow us to collapse the Δ_i^L decisions into a single decision variable. With this, the problem of jointly optimizing Δ^H and the Δ_i^Ls for a fixed m reduces to a single-variable optimization problem. Then, in the "outer" optimization, we optimize m, utilizing expressions for the optimal Δ^H and Δ_i^Ls as functions of m. The shape of the objective as a function of m is fairly well behaved (in the worst case, a cubic function), but the solution depends upon the signs and relative values of the coefficients in the relevant function.

In the discussion that follows, we assume $N \ge 1$; the special case of $N = 0$ (i.e., no low cycles) is discussed in Appendix I. Here, we provide an overview of the key structural results that provide the foundation for solution procedure(s); finer details and proofs are relegated to the Appendices.

Inner Optimization Problem

In the inner optimization problem, the sum of the discounts is fixed. We represent the average discount per period by m (which is also the loss of brand equity), so the total discount across all periods in a cycle is mT. Therefore, for any fixed m, the problem can be written as:

$$\max_{\Delta^H, \Delta_1^L, \ldots, \Delta_N^L} \Pi = \Pi^H(\Delta^H) + \sum_{i=1}^{N} \Pi^L(\Delta_i^L)$$

$$\text{subject to } \Delta^H + \sum_{i=1}^{N} \Delta_i^L = mT$$

$$0 \le \Delta^H \le hT^H$$

$$0 \le \Delta_i^L \le hT^L, \ i = 1, \ldots, N$$

Both Π^H and Π^L are cubic functions (of Δ^H and Δ^L, respectively) and differ only in their coefficients. Given that $c^H \ge 0$ and $c^L \ge 0$, Π^H and Π^L can take one of two possible functional forms: (i) convex increasing then concave (and unimodal) and (ii) strictly concave. Examples of the functional forms of Π^H are shown in Figure 3.1, and the forms of Π^L are exactly the same. Although these two functional forms are unimodal, we cannot guarantee that the objective is jointly unimodal in the discounts. We can, however, obtain a partial characterization of the solution for the low cycles, as stated in Proposition 1 below.

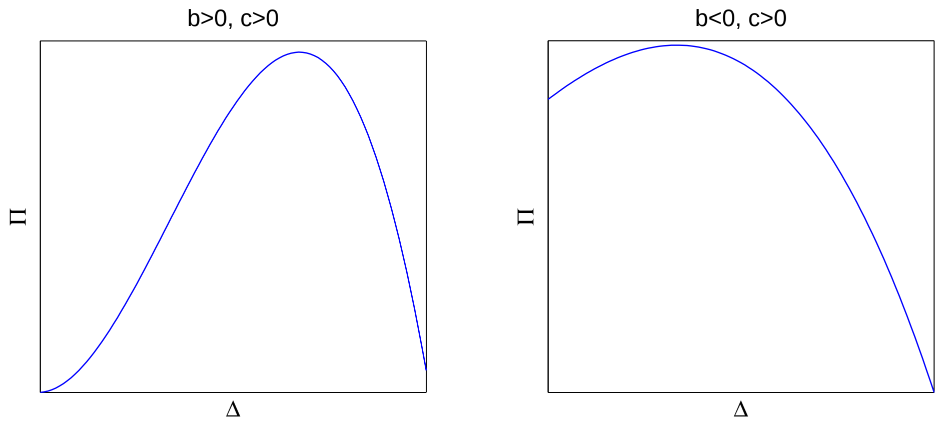

Figure 3.1: Examples of Profit Functions for a Replenishment Cycle

Proposition 1. *Let $(\Delta^{H*}, \Delta_1^{L*}, ..., \Delta_N^{L*})$ be the optimal solution to the retailer's problem in which there is a high cycle followed by N low cycles. For any low cycles i and j such that the discounts Δ_i^{L*} and Δ_j^{L*} are positive, $\Delta_i^{L*} = \Delta_j^{L*}(= \Delta^{L*})$.*

Proof: See Appendix F.

The intuition underlying the proposition can be explained as follows. For some combinations of parameter values, Π^L may be convex for Δ^L below a threshold. Because we wish to maximize the objective in the N-dimensional space of Δ_i^L values, boundary solutions (i.e., with one or more of the Δ_i^L values equal to zero) may be optimal. On the other hand, when Π^L is concave increasing, equal and positive values of Δ_i^L are optimal. For any fixed number, n, of positive Δ_i^Ls that are equal to each other, optimizing the Δ_i^L values reduces to a single variable problem. Note also that Δ^H can be expressed as $mT - \sum_{i=1}^{N} \Delta_i^L$, so the problem of jointly optimizing all of the Δ values for a given n remains a single-variable problem but we need to consider all possible number of positive Δ_i^L values, i.e., $n = 0, ..., N$, to find the optimal solution.

Intuitively, the convexity of Π^L (as function of Δ^L) below a threshold can be explained by the fact that the number of units sold at a discount is convex increasing (roughly quadratic) as the discount increases. Not only do customers who were already purchasing choose to buy more, but customers who did not purchase at smaller discounts now choose to purchase. The retailer thus sells more units at a discount when offering one large discount than when offering two smaller discounts whose sum is equal to the single large discount. Therefore, holding the loss of brand equity constant, if the volume increase outweighs the reduction in the unit profit margin, the retailer may be better off offering a mix of large and zero discounts instead of equalizing discounts across low cycles. Therefore, when Π_i^L has a convex region, the discounts in the low cycles are not necessarily equal to each other despite the fact that we have the same economic parameters in all low cycles. On the other hand, in the absence of the brand equity effect, the retailer faces independent and identical decisions in each low cycle and therefore offers the same discount in all low cycles. The convexity that may arise due to the brand equity effect thus provides a reason for fluctuations in retail prices even during time intervals when the manufacturer's wholesale price is constant.

We now provide results that relate decisions in the high and low cycles.

Proposition 2. *Let* $(\Delta^{H*}, \Delta_1^{L*}, ..., \Delta_n^{L*})$ *be the optimal solution to the retailer's problem. If* $T^H \geq T^L$ *and* $\Delta^M \geq 0$*, then* $\Delta^{H*} \geq \Delta_i^{L*}, \forall i = 1, ..., N$*, that is, the retailer's discount in high cycles is greater than or equal to the discount in low cycles.*

Proof: See Appendix J.

Propositions 2 shows that manufacturer discounts cause the retailer to offer the same or higher discount if the retailer chooses a longer ordering interval. We explore the question of what fraction of the manufacturer's discount the retailer passes on to consumers in our numerical study.

To summarize, for a fixed m, we can reduce the problem of finding Δ^H and the Δ_i^L values to a unidimensional problem by utilizing the fact that all positive discounts in low cycles are equal. If $b^L < 0$, then Π^L is initially convex and then concave, and we need to solve problems with $n = 1, \ldots, N$ positive discounts and choose the best one. If $b^L \geq 0$, then Π^L is concave, hence we only need to solve the inner problem with $n^* = N$ positive discounts.

For a fixed n, we show in Appendix K that the retailer's objective as a function of the (positive) discount in the applicable low cycles has one of four forms: (i) convex increasing; (ii) convex decreasing and then increasing; (iii) initially concave then convex; (iv) concave decreasing then convex. Which form is pertinent depends upon whether certain coefficients, which can be computed from the problem data, are positive or negative.

3.4.1.1 Outer Optimization Problem: Optimizing m

In this section, we analyze the outer problem (i.e., optimizing m) for a fixed n and provide an overview of properties of the objective function that form the foundation for identifying the optimal solution. Details appear in Appendix L. The relevant functions are cubic or quadratic expressions, so the optimal solution is either a stationary point (if it is the global maximum) or a boundary solution with Δ^L at its lower or upper limit. The characteristics of these functions (e.g., whether the cubic functions are initially concave then convex or the reverse, and whether the quadratic functions are convex or concave) depend on the coefficients which, in turn, depend upon the problem parameters.

We refer to the coefficients as $\alpha_1, \alpha_2, \alpha_3$ and α_4 and the objective function has the form $\alpha_1 m^3 + \alpha_2 m^2 + \alpha_3 m + \alpha_4$. The expressions for the coefficients can be found in Appendix L. Under the strategy $(\Delta^H = mT, \Delta^L = 0)$, α_1 is a constant and the other coefficients are linear functions of Δ^M. Under the strategy $(\Delta^{L^*} > 0, \Delta^H = mT - n\Delta^{L^*})$, α_1 is again a constant, α_2 is a linear function of Δ^M, and α_3 and α_4 are both equal to zero. The "shapes" of the functions depend upon the signs (negative, zero, or positive) of the coefficients. Typically, α_1, the coefficient on the cubic term, is negative, and it primarily captures the effect of discounting on overall revenue. If $\alpha_2 > 0$ at a given value of Δ^M, then, roughly speaking, the retailer gains more from selling to discount customers than he loses due to the reduced margin on units sold to regular customers. If $\alpha_3 > 0$ at a given value of Δ^M, then the retailer's savings in inventory costs from discounting is greater than the reduction in profit due to the fact that the (negative) brand equity effect reduces the number of regular customers. If α_2 is positive and α_3 is negative (or the reverse), then the retailer faces competing forces when optimizing his discounts. If both α_2 and α_3 are negative, then the retailer has little incentive to offer discounts at the given value of Δ^M. In most cases, the stationary point can be expressed in closed form, but sometimes a unidimensional numerical search is required. Four prototypical shapes of the objective function are shown in Figure 3.2.

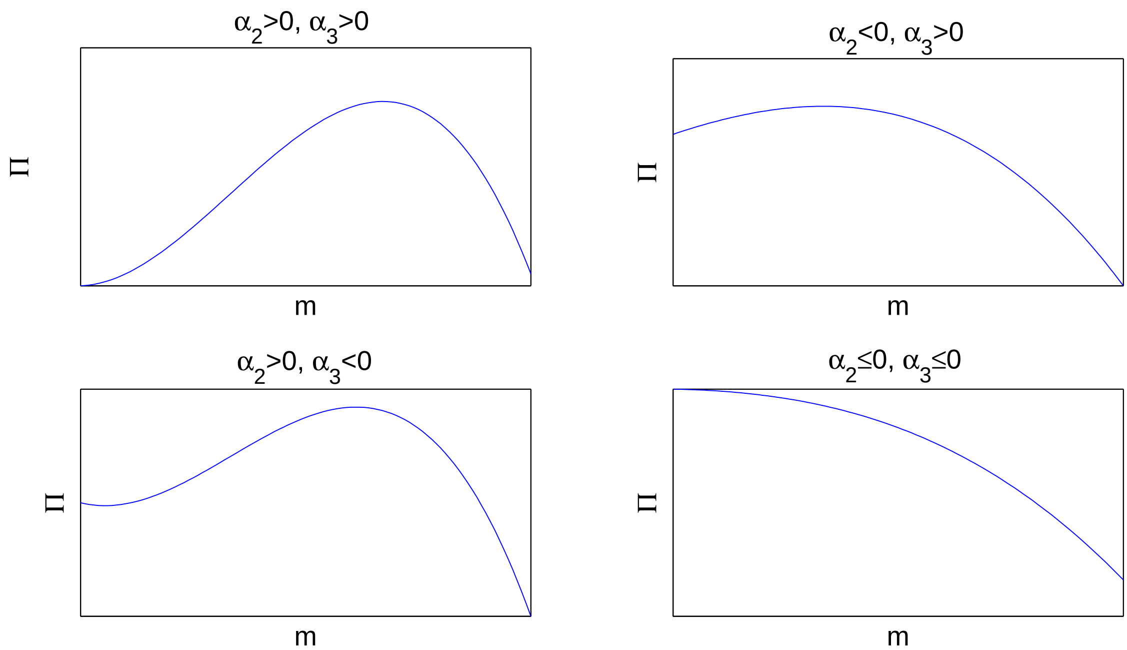

Figure 3.2: Retailer's Objective as a Function of m

When we combine the facts that the coefficients change with Δ^M with the fact that the retailer's profit functions are quadratic or cubic functions of m, we can infer that the optimal value of m may not change smoothly as Δ^M changes. In particular, it is possible for the optimal retailer discount to be zero for small values of Δ^M, eventually becoming positive at some threshold, and then increasing in a convex, concave, or even a discontinuous fashion as

Δ^M increases further, due to the the manner in which Δ^M affects α_2 and α_3. So the analytical results indicate there may be threshold effects, and that the pass–through percentage may be increasing, decreasing, or even fluctuating as a function of Δ^M. We observe all of these phenomena in the numerical results that we report later.

3.4.2 Special Case: No Regular Customers

Analysis of the special case of no regular customers (i.e. $m > \overline{R} - p$) is analogous to that of the case with regular customers discussed above. Under the assumption of uniformly distributed reservation prices, the retailer's profit functions in high and low cycles, respectively, can be written as:

$$\begin{aligned} \Pi^H(\Delta^H, \Delta_i^L, T^H, T^L, N) &= a^H(\Delta^H)^3 + b^H(\Delta^H)^2 + c^H\Delta^H + \gamma^H && (3.17) \\ \Pi^L(\Delta^H, \Delta_i^L, T^H, T^L, N) &= a^L(\Delta_i^L)^3 + b^L(\Delta_i^L)^2 + c^H\Delta_i^L + \gamma^L && (3.18) \end{aligned}$$

where the coefficients are:

$$\begin{aligned} a^H &= a^L = -0.5\frac{\Gamma}{h\overline{R}}, \\ b^H &= [0.5(p - w + \Delta^M) - (\overline{R} - p - m)]\frac{\Gamma}{h\overline{R}}, \\ c^H &= [(R - p - m)(p - w + \Delta^M) - 0.5(R - p - m)^2]\frac{\Gamma}{h\overline{R}}, \\ \gamma^H &= 0.5(p - w + \Delta^M)(R - p - m)^2\frac{\Gamma}{h\overline{R}} - K, \\ b^L &= [0.5(p - w) - (\overline{R} - p - m)]\frac{\Gamma}{h\overline{R}}, \\ c^L &= [(R - p - m)(p - w) - 0.5(R - p - m)^2]\frac{\Gamma}{h\overline{R}}, \\ \gamma^L &= 0.5(p - w)(R - p - m)^2\frac{\Gamma}{h\overline{R}} - K. \end{aligned}$$

Note that a^H, b^H, a^L, and b^L are the same whether $m \leq \overline{R} - p$ or $m > \overline{R} - p$; only the other coefficients differ. Propositions 1 and 2 still hold and the problem can be solved in a similar fashion. Details appear in Appendix K.

Likewise, the solution to the outer optimization problem parallels that for the case of $m \leq \overline{R} - p$. (See Appendix L for details.) What is different is that when $m \leq \overline{R} - p$, the brand equity effect only affects the number of regular customers and does not affect the number of discount customers who purchase at any particular discount. In contrast, when $m > \overline{R} - p$, the retailer does not have any regular customers but brand equity affects the number of discount customers purchasing at a discount (i.e., $d_1^H = \frac{\Gamma}{\overline{R}} \int_{p-\Delta^H}^{\overline{R}-m} \frac{\Delta^H - (p-r)}{h} dr$). The retailer must make the following tradeoff: on one hand, by offering a larger discount, the retailer can reach customers with lower reservation prices, and on the other hand, the reservation prices of all customers decline so the number of customers whose reservation prices are above the discount price also declines. In addition, the retailer also reduces his profit margin on all units sold when he offers a deeper discount (because no units are purchased at the regular price in this case).

The models in Sections 3.4.1 and 3.4.2 include the negative effects of discounting on brand equity (willingness to pay). In the next section, we present an analysis for the case in which there is no brand equity effect due to retail discounting.

3.4.3 Model Without a Brand Equity Effect

In the model with no brand equity effect, the retailer's objective has the same form as that in the model with brand equity effect, and it can be written as:

$$\Pi = (\Pi^H + \sum_{i=1}^{N} \Pi_i^L)/T \qquad if \;\; N \geq 1$$
$$\Pi = \Pi^H/T \qquad if \;\; N = 0.$$

The only component of the retailer's objective that is different from the model with a brand equity effect is that the distribution of the customers' reservation prices is the unadjusted distribution, $f(\cdot)$, instead of the adjusted distribution, $\tilde{f}(\cdot)$.

The expressions for demand are analogous to the expressions for demand in (3.12), (3.13) and (3.14) and can be written as:

$$d_1^H = \Delta^H[(\overline{R} - p) + \Delta^H/2]\frac{\Gamma}{h\overline{\overline{R}}}, \tag{3.19}$$
$$d_{1i}^L = \Delta_i^L[(\overline{R} - p) + \Delta_i^L/2]\frac{\Gamma}{h\overline{\overline{R}}}, \tag{3.20}$$
$$d_{2i}^L = (hT^L - \Delta_i^L)(\overline{R} - p)\frac{\Gamma}{h\overline{\overline{R}}}. \tag{3.21}$$

The retailer' profit functions in high and low cycles, respectively, are

$$\Pi^H(\Delta^H, \Delta_i^L, T^H, T^L, N) = a^H(\Delta^H)^3 + b^H(\Delta^H)^2 + c^H\Delta^H + \gamma^H \tag{3.22}$$
$$\Pi^L(\Delta^H, \Delta_i^L, T^H, T^L, N) = a^L(\Delta_i^L)^3 + b^L(\Delta_i^L)^2 + c^H\Delta_i^L + \gamma^L \tag{3.23}$$

where the coefficients are:

$$\begin{aligned}
a^H &= a^L = -0.5\frac{\Gamma}{h\overline{\overline{R}}}, \\
b^H &= [0.5(p - w + \Delta^M) - (\overline{R} - p)]\frac{\Gamma}{h\overline{\overline{R}}}, \\
c^H &= 0.5h_rT^H(\overline{R} - p)\frac{\Gamma}{h\overline{\overline{R}}}, \\
\gamma^H &= (p - w + \Delta^M - 0.5h_rT^H)hT^H(\overline{R} - p)\frac{\Gamma}{h\overline{\overline{R}}} - K, \\
b^L &= [0.5(p - w) - (\overline{R} - p)]\frac{\Gamma}{h\overline{\overline{R}}}, \\
c^L &= 0.5h_rT^L(\overline{R} - p), \\
\gamma^L &= (p - w - 0.5h_rT^L)hT^L(\overline{R} - p)\frac{\Gamma}{h\overline{\overline{R}}} - K.
\end{aligned}$$

When there is no brand equity effect, the discounts in the high and low cycles have no interactions, so the retailer can maximize Π^H and Π^L independently. Π^H and Π^L are unimodal in Δ^H and Δ^L, respectively, over the relevant range, and both have a local maximum defined by the corresponding stationary point. The optimal discount in a high cycle is $\Delta^{H*} = \frac{-b^H - \sqrt{{b^H}^2 - 4a^Hc^H}}{2a^H}$

and the optimal discount in a low cycle is $\Delta^{L*} = \frac{-b^L - \sqrt{{b^L}^2 - 4a^L c^L}}{2a^L}$.

Observe that b^H is linearly increasing in Δ^M, a^H is a negative constant and c^H is a positive constant. So from the expression for Δ^{H*}, we can see that the retailer's discount in a high cycle is strictly increasing in Δ^M.

3.5 Numerical Study

In this section, we report on optimal solutions for problems covering a wide range of parameter combinations. The focus of the study is on how various problem parameters, especially the manufacturer's discount, affect the retailer's pass–through. Our study focuses on situations where there is a brand equity effect. Results in Section 3.4.3 show that the relationships are much simpler in the absence of a brand equity effect.

The problem parameters are: $\Gamma = 1000$; $\overline{R} = 3$ and 6; $p = 0.25\overline{R}, 0.5\overline{R}$, $0.75\overline{R}$ and $0.95\overline{R}$; $w = 0.4p$, $0.6p$, $0.8p$ and $0.9p$; $T = 4$, 12 and 24; $h_r = 0.3w/52$, $w/52$ and $4w/52$; $h = 0.02p$, $0.05p$, $0.3p$ and $0.5p$; and $K = 100$. These parameter combinations yield a total number of 1152 instances.

The choices of $\overline{R}$ are arbitrary, as they are simply scale parameters, but we chose the other parameters for the following reasons:

• The three values of p represent situations with 75%, 50%, 25% and 5% regular customers, respectively, which covers a wide range of willingness-to-pay scenarios vis-a-vis the regular price.

• The four ratios of w to p capture the range or retail margins that are typical for consumable consumer goods.

• The three values of the customer's holding cost rate (h) represent situations in which the customer will stockpile one unit (one week's usage) for discounts of 2%, 5%, 30% and 50%, respectively. At the low end of this spectrum, the customer is quite prone to stockpiling whereas at the high end, the customer is reluctant to stockpile.

• The four values of the retailer's holding cost rate (h_r) reflect annual holding cost rates of 30% (typical for non-perishable foods), 100% (appropriate for foods requiring refrigeration or freezing) and 400% (appropriate for situations where products have short life cycles or retailer shelf space is extremely limited).

• The values of T were chosen so that several combinations of T^H, T^L and N are feasible.

• The value of K is chosen so that the retailer's economic order quantity ranges from one week of supply to 22 weeks of supply. We only use a single K value because the variation in h_r yields a wide range of values for the retailer's economic order quantity.

As mentioned in the previous sections, we find the retailer's optimal cycle lengths, T^{H*} and T^{L*}, and the optimal number of low cycles, N^*, within each T–period trade discount cycle by enumeration. Under the framework of integral cycle lengths, we enumerate T^H in $[1, \ldots, T]$; then we find the set of integer divisors of $T - T^H$ (denoted $d^{\mathcal{N}}$). We enumerate N in $d^{\mathcal{N}}$ and for each relevant N, the length of each low cycle, T^L, is set equal to $(T - T^H)/N$.

For each parameter combination, we determine the retailer's optimal response for Δ^M equal to $0.1w$, $0.2w$, $0.3w$, $0.4w$, $0.5w$, $0.6w$, $0.7w$, $0.8w$ and $0.9w$. Values at the upper end of this range are not common in practice, but we include them for completeness. In the next subsection, we analyze how the manufacturer's discount affects the retailer's optimal discount and ordering schedule, and in the subsequent subsection, we examine under what parameter combinations the retailer tends to pass a large (or small) proportion of the manufacturer's discount on to consumers.

3.5.1 How the Manufacturer's Discount Affects the Retailer's Optimal Discount and Ordering Schedule

The analytical results in the previous section raised the possibility that the retailer's discount in the high cycles (and therefore also the pass–through rate) may increase, decrease, or fluctuate as Δ^M increases. We categorize the observed relationships into six patterns: (i) increasing; (ii) decreasing; (iii) unimodal with a local minimum; (iv) unimodal with a local maximum; (v) fluctuating; and (vi) no change. Out of 1152 instances, the number of instances that fall into patterns (i) through (vi) are 467, 39, 15, 126, 36 and 469, respectively. The proportional split of problem instances among the patterns is largely a consequence of our choice of parameters. Our main concern was to have enough instances so that we could relate the patterns to systematic characteristics of the mix of problem parameters. Below, we discuss these relationships.

In some instances with Pattern (i), Δ^{H^*} is *strictly* monotonically increasing in Δ^M over the entire range of Δ^M values that we considered. These situations tend to arise when the retailer is already offering a discount even in the absence of a manufacturer discount. If the manufacturer offers a discount, the retailer sometimes chooses to increase his discount. If he does so, the discount tends to increase monotonically with the manufacturer's discount, at least up to some threshold where it is no longer advantageous for the retailer. This pattern arises commonly when the retailer has small regular demand (in which case the reduction in profit on sales to regular customers is small), the consumers have a small holding cost (so small changes in the discount can induce substantial stockpiling, which represents additional demand for the retailer), and the manufacturer discounts infrequently (so the retailer stockpiles large quantities when the manufacturer offers discount and therefore has a greater incentive to clear inventory by discounting).

In other instances in which we observe Pattern (i), Δ^{H^*} is a discontinuous function of Δ^M: it remains zero until Δ^M reaches a threshold, and then jumps from zero to a positive value at the threshold, and may or may not increase further beyond the threshold. The threshold effect is due to the fact that for some parameter combinations, the retailer's objective is a cubic function of the discount, and Δ^M affects the shape of the cubic function. Up to some threshold value of Δ^M the optimum is zero, but when Δ^M reaches the threshold, the α coefficients in the retailer's objective function change enough to alter the shape of the function so that the optimum shifts to a positive value. These instances tend to arise when the retailer has small regular demand and the manufacturer discounts frequently, so the retailer does not bear much expense in holding inventory and therefore has little incentive to offer discounts. The retailer chooses to offer a zero discount unless Δ^M is large enough to make it worthwhile to switch to a different strategy that focuses greater penetration among the discount customers. When the retailer makes this switch, the high trade discount provides him an adequate margin despite the fact that he is offering deep discounts to his customers.

In Pattern (ii), Δ^{H^*} is non-increasing in Δ^M and strictly decreasing for at least one Δ^M in the range of values that we consider. In the problem instances that give rise to this pattern, the manufacturer offers a trade discount frequently and the the retailer typically has a large regular demand and moderate holding costs. When the manufacturer offers a small trade discount, the retailer must trade off the benefits of inventory reduction (by inducing customer stockpiling) with the loss of revenue from reduced brand equity. The retailer's moderate holding costs combined with frequent trade discounts give him a fairly small incentive to clear inventory. On the other hand, the reduction in profit on sales to the large portion of customers who would have been willing to pay full price can be substantial. The retailer may choose to offer a small discount but is not inclined to discount very deeply. When the manufacturer offers a large trade discount, the discount serves to subsidize the retailer's cost of holding inventory (which is still modest due to the retailer's moderate holding cost and the frequent trade discounts). The retailer chooses to stockpile and then sell these items at (close to) the regular price. Hence the retailer decreases his discount as Δ^M increases.

In Pattern (iii), the retailer's discount first decreases in Δ^M (for the same reasons noted for Pattern (ii)) and then increases. The primary difference between problem instances that result in Pattern (iii) versus Pattern (ii) is that the retailer has smaller regular demand in the former, so for sufficiently large values of Δ^M, the retailer is willing to forego some of the profit margin on sales to regular customers in exchange for the opportunity to sell to the much larger number of customers who are willing to purchase only at a relatively deep discount.

In Pattern (iv), the retailer's discount initially increases and then decreases as the manufacturer's discount increases. In the problem instances that generate this pattern, the manufacturer offers a trade discount infrequently. As Δ^M increases, the retailer increases his stockpiling and simultaneously increases his discount to transfer part or all of his stockpile to customers. For a sufficiently high Δ^M, the retailer opts to purchase only when the manufacturer offers a discount. For Δ^M above this threshold, the incremental trade discount provides a subsidy for the retailer's cost of holding inventory, so the retailer's need to discount declines. When the subsidy is sufficiently high, the retailer prefers to reduce his discount in order to recapture sales at the full-price. This occurs because a reduction in his discount increases the proportion of regular customers, and regular customers purchase to satisfy consumption throughout the entire cycle whereas discount customers purchase to satisfy consumption for only a portion of the cycle. Furthermore, due to the high trade discount, the full-price sales provide the retailer a substantial net margin. Thus, for a sufficiently high Δ^M, the retailer employs a different discounting strategy than when Δ^M is small, and this leads to the pattern of retail discounts initially increasing and then decreasing as the trade discount increases.

The fluctuations observed in retailer's discount in Pattern (v) as Δ^M increases are due to the requirement that the retailer's cycle lengths be integral. In the instances leading to Pattern (v), the retailer has a large regular demand and high holding costs. Because the retailer must hold inventory for the regular customers (who purchase just-in-time during the latter part of the cycle) and his high holding costs make it costly to do so, the retailer uses the strategy of coupling increases in stockpiling with increases in the discount. However, due to integrality of cycle lengths, as Δ^M increases, T^{H^*} does not strictly increase. For values of Δ^M at which T^{H^*} increases, the retailer's pass–through increases as well; for values of Δ^M at which T^{H^*} is unchanged, the retailer's discount decreases for the same reason as noted for Pattern (ii). As such, we attribute these fluctuations to "rounding error": if the retailer could choose non-integral cycle lengths, the discounts would likely change in a smoother fashion.

Among instances that lead to Pattern (vi), Δ^{H^*} is constant at zero in 398 out of 469 instances and is constant at a positive value in 71 out of 469 instances. Situations with $\Delta^* = 0$ tend to arise when the customers have a large holding cost, the retailer has low holding cost and moderate to large regular demand, and manufacturer offers a discount frequently. In these circumstances, the retailer needs to hold inventory for the regular customers but it is relatively inexpensive for him to do so because he has a low holding cost rate and does not need to stockpile large quantities because the manufacturer offers discounts frequently. Because customers have a high holding cost, significant discounts would be needed to induce the relatively small number of discount customers to purchase very much. None of these factors weighs in favor of the retailer offering discounts.

Two different combinations of factors lead to situations in which the retailer offers a positive discount but does not increase it as Δ^M increases. The first combination of factors is a very large proportion of regular customers and a very small customer holding cost. The retailer offers a small discount in all of these instances but can generate considerable sales to discount customers due to their small holding cost. The retailer does not increase his discount as Δ^M increases because doing so would cause him to sacrifice considerable profit on sales to regular customers without compensating profit from discount customers. The second combination of factors is a small proportion of regular customers and a large initial profit margin $(p - w)$. In these circumstances, the retailer chooses deep discounts, foregoing the small volume of potential full-price sales, in order to reach a substantial portion of the discount customers who will not otherwise buy. The large initial profit margin allows the retailer to offer deep discounts. The manufacturer's trade discount causes the retailer to make purchases only when trade discounts are offered, but because the retailer would offer deep discounts even in the absence of a trade discount, an increase in the trade discount is insufficient incentive for the retailer to offer even deeper discounts. (Beyond a threshold, additional retail discounts only decrease per-unit margin while not increasing unit sales.)

Although the retailer's optimal discount may change in different ways as Δ^M increase, the retailer's ordering schedule changes in a consistent way: the retailer chooses larger values of T^H and in concert, stockpiles more. The number of low cycles decreases correspondingly, but the lengths of the low cycles remain constant in the vast majority of problem instances and nearly constant (i.e., a mix of two consecutive integer values) in the rest of the instances. Recall that T^H and T^L need to be integers in our model. So, the differences that we observe can be viewed as "rounding error" due to the integrality constraints; the retailer probably would have chosen equal non-integral values of T^L if such an option were available.

We now turn to a discussion of a more commonly-reported metric, the pass–through rate.

3.5.2 Retailer's Pass–Through Rate

The retailer's pass–through rate is traditionally defined as the ratio of the retailer's discount to the manufacturer's trade discount. Figure 3.3 shows the cumulative percentage of problem instances with 0, 25%, 50%, 75%, 100%, 200%, 300%, 400% and over 400% pass–through rates for $\Delta^M = 0.1w, 0.3w, 0.5w, 0.7w$ and $0.9w$. As shown in the figure, for all Δ^M values, roughly 40% of instances have a zero pass–through rate and the majority of the instances have pass–through rates below 25%. Thus, broadly speaking, pass–through rates are quite low, but interestingly, a considerable portion (up to 20%) of instances have pass–through rates above 100% when Δ^M is small. The range and mix of pass–through rates observed in our numerical study (based on our analytical model) is consistent with what researchers have reported in empirical studies.

In a small portion of problem instances, the pass–through rate is as high as 300% or 400%. In these problem instances, retailer's profit margin is high in the absence of a trade discount, but the fraction of customers who are willing to pay the regular price is low. So the retailer can "afford" to implement a high–pass through rate because he can maintain a relatively large net margin while simultaneously reaching customers who are unwilling to pay an amount close to the regular price. We observe analogous situations in practice for products that are frequently offered on a "buy-one-get-one-free" basis: retail margins are quite high at the regular price and not many customers are willing to pay the full price, so manufacturers have to invest a lot of money to induce retailers to offer deeply-discounted promotional prices.

In our model framework, the retailer may offer a discount even in the absence of a trade discount. Thus far in this section, we have reported on the traditional pass–through rate metric, but a more accurate metric would be the ratio of the *incremental* discount that the retail offers

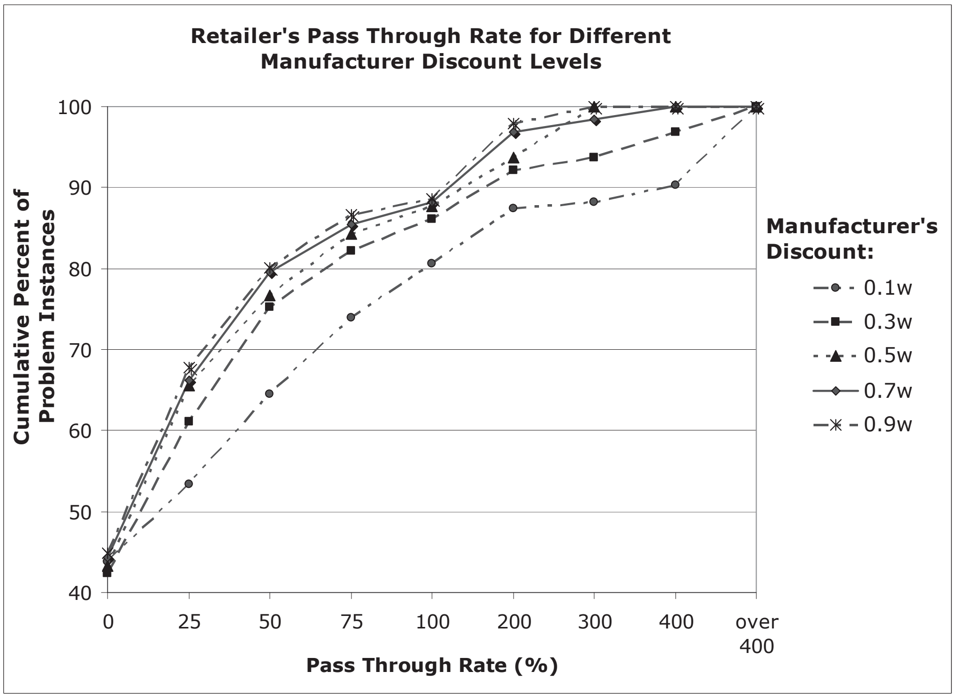

Figure 3.3: Distribution of Retailer's Pass–Through Rates for Various Trade Discount Levels

due to the trade discount to the trade discount. Figure 3.4 shows the cumulative percentage of problem instances with -5%, 0%, 10% 15%, 20%, 25%, 30%, 40%, 50% and 100% incremental pass through rates.

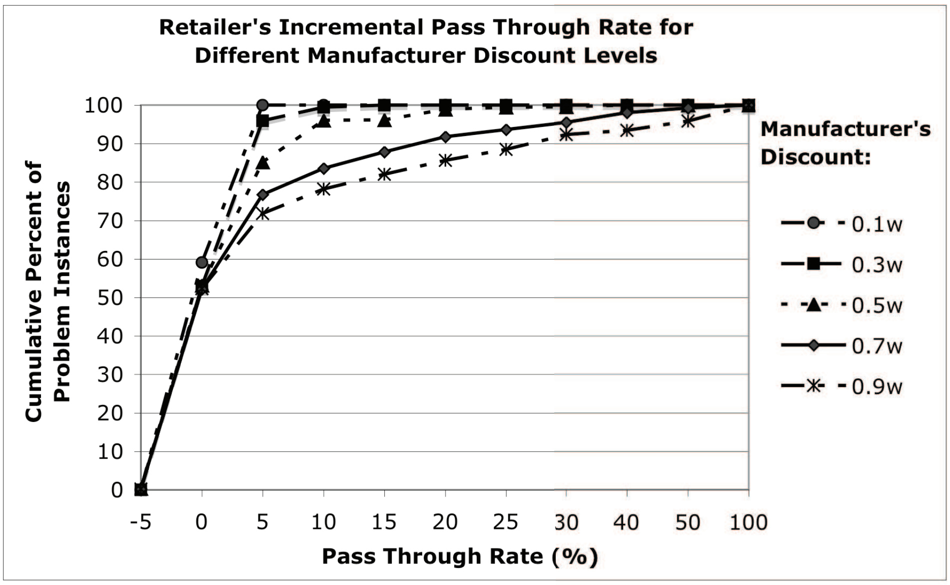

Figure 3.4: Distribution of Retailer's Incremental Pass–Through Rates for Various Trade Discount Levels

At all discount levels that we consider, the incremental pass–through rate is zero or negative in about 50% of the problem instances, and problem instances with incremental pass-through rates of more than 10% are rare. For example, when the trade discount is 50% of the wholesale price, the retailer passes through 20% or more of the discount in only about 1% of problem instances. Although the overall incremental pass–through rates are quite low overall, interestingly, they are higher for large manufacturer trade discounts than they are for small trade discounts; the opposite relationship applies to the traditional pass–through rate. These results suggest that manufacturer's face extremely difficult tradeoffs: either the manufacturer can keep trade discounts low (or zero) and get minimal reaction from the retailer, or he can invest large sums recognizing that very large trade discounts are required to get a meaningful incremental reaction from the retailer. These results also raise the question of the choice of metrics to provide economic incentives to decision–makers that are aligned with the manufacturer's goals.

As expected, both the traditional pass–through rate and the incremental pass–through rate generally decline as the trade discount increases, modulo the effects of the constraints that T^H and T^L be integral, as mentioned in the last subsection. This is true even in cases where Δ^{H*} increases as Δ^M increases, because the former almost always increases more slowly than the latter, and it increases much more slowly or declines for sufficiently large values of Δ^M.

3.6 Conclusions

In this chapter, we investigate the retailer's optimal pass–through of manufacturer trade discounts. We study a scenario with three types of strategic players: a manufacturer, a retailer and customers. The manufacturer offers a constant regular wholesale price but offers periodic trade discounts to the retailer in the hope of achieving higher market penetration and possibly earning a higher profit. The retailer offers a constant regular price and decides the timing and depth of discounts, which may coincide with trade discounts but may also be offered at other times. The retailer chooses his ordering and discounting patterns to maximize his profit per unit time, and stockpiles when the manufacturer offers a discount if the benefit of stockpiling outweighs the cost of holding inventory. Customers, who are heterogeneous with respect to their reservation prices, stockpile when the retailer offers a discount. They wish to consume the product at a constant rate but purchase only if their gross cost—the actual purchase price plus the cost of holding the item until it is consumed—does not exceed their reservation price. In addition, it is well known that retail discounts erode brand equity and therefore also reservation prices, and we include this effect in our model.

To the best of our knowledge, our study is the first to incorporate the effects of retailer and consumer stockpiling, along with the (negative) effects on brand equity of retail discounts, in a retailer's decision model for choosing discounting and ordering patterns when a manufacturer offers periodic trade discounts. We develop a full characterization of the retailer's optimal discounting policy for a given ordering schedule and we use this to determine a jointly optimal discounting and ordering plan assuming that orders can be placed at discrete points in time (e.g., weekly). Results from our analytical developments and a numerical study provide a number of important managerial insights, including:

• The combination of brand equity effects (due to retailer discounting) and retailer and consumer stockpiling may lead to threshold effects for the manufacturer: it may be necessary for the manufacturer to offer a discount above a threshold before there is any reaction on the part of the retailer. These situations tend to arise when the retailer has small regular demand and the manufacturer discounts frequently so the retailer's stockpiling quantities are relatively small and his incentive to clear inventory more rapidly by discounting is small. In these circumstances, it may be unprofitable for the retailer to increase his discount further unless the trade discount is sufficiently high, due to the loss of customers who are willing to pay full price. We show that the threshold effect is due to the brand equity effect; it does not arise in the absence of this effect.

• Numerical results indicate that the retailer's discount may not be monotonic in the trade discount. At low or moderate trade discounts, the retailer may prefer to focus on sales volume and discount fairly deeply to reach more discount customers. On the other hand, when the trade discount is high, the retailer may opt to offer a low discount so as to bolster the number of customers who are willing to pay full price, and then sell to these customers at very high margins (owing to the high trade discount). Thus, different depths of the manufacturer's trade discount may lead to distinctly different pricing and consumer targeting strategies by the retailer.

• The retailer tends to offer a higher (absolute) pass–though of trade discounts when the fraction of regular customers (those who are willing to pay full price) is small, the retailer has high to moderate holding costs, customers have low to moderate holding costs and the manufacturer discounts infrequently. In such circumstances, if the retailer stockpiles, he incurs significant inventory holding costs due to the combination of his own high to moderate holding costs and the long time between trade discount offers, so he has a strong incentive to offer discounts to induce customers to stockpile. Customers are willing to stockpile to a fair extent because their inventory holding costs are not too high. Also, because the fraction of regular customers is small, if the retailer offers discounts, he does not suffer much loss of sales at full price even if these customers switch to buying only at a discount due to the (negative) brand equity effect. From this, it is clear that a combination of favorable factors is needed for a high pass–through rate, and the absence of any one of these factors could lead to a low pass–through rate. For example, if the manufacturer offers trade discounts frequently and the retailer has low inventory holding costs, the retailer may simply choose to stockpile during trade discount opportunities and gradually sell it at his regular price. It is important for manufacturers to consider the complex interactions among factors that affect retailer's discounting choice when designing a trade discount plan.

• Although pass-through rates measured by the traditional metric (retailer's discount divided by the manufacturer's discount) tend to be modest, the incremental pass–through rates (retailer's additional discount above and beyond his normal discount divided by the manufacturer's discount) are commonly zero or negative, and rarely are they greater than 10% in our numerical study. Interestingly, (traditional) pass–through rates tend to be higher for small manufacturer discounts (10% or 20% of the wholesale price) but incremental pass through rates are higher when manufacturer discounts are high. These results raise the question of the choice of metrics to provide economic incentives that are aligned with the manufacturer's goals.

We close this chapter with a discussion of how the findings from our analytical model and numerical study compare with empirical findings in the literature. Although virtually all empirical studies have been conducted in competitive settings with multiple products and multiple retailers, overall, the findings from our analysis and numerical study are very consistent with what the empirical studies indicate. First, several recent studies (Chintagunta 2002, Besanko et al. 2005 and Nijs 2010) report negative pass through rates, i.e., retailers raising the regular price when the manufacturer offers a trade discount. In a study focused on other aspects of retail pricing, Chintagunta (2002) observed a negative correlation between wholesale and retail prices for one of the five products that he studied at one retail chain. The product happened to have the highest price, highest retail margin, and smallest volume in its category. He conjectured that retailers could be raising prices during manufacturer promotions to maximize profits from loyal customers. Besanko et al. (2005) and Nijs (2010) also found negative pass-through in some instances but did not provide details on characteristics of the pertinent products nor explanations for the negative pass-through rates. In our numerical results, we observed problem instances with relatively few customers who are willing to pay the full price for which the retailer's optimal policy is to offer a large discount when the trade discount is small because, in these circumstances, it is more profitable to target the large number of customers who are only willing to purchase a discount. On the other hand, when the manufacturer's trade discount is large, the retailer uses a completely different strategy and offers zero or a small discount, targeting the customers who are willing to pay the full price. In our model, the regular price is fixed and we assume that discounts are non-negative, so we do not have outcomes in which the retailer raises the regular price. But problem instances in which retailer switches strategies as described above have similar characteristics to the situation in which Chintagunta conjectured that a negative pass–through was occurring: relatively few customers are willing to pay the full price, and the per–unit margin considering the trade discount is quite high.

Results from empirical studies also indicate that pass-through rates are high when demand is elastic (Nijs et al. 2010; Meza and Sudhir 2006; Walters 1989), and that the retailer offers a higher pass–through in response to trade deals when the manufacturer discounts infrequently (Walters 1989; Nijs et al. 2010). Our analytical results are consistent with both of these findings. The impact of demand elasticity is not surprising, but our analytical model helps to explain why the frequency of trade discounts matters when both the retailer and customers stockpile. In particular, if trade discounts are infrequent, the retailer tends to stockpile more, but in turn, needs to offer a deeper discount to reduce his inventory holding costs. In addition, Nijs et al. (2010) found the retail pass-through to be higher for products with larger package sizes. Our results also suggest that retailers tend to offer a high pass-through when the product has a high holding cost; thus pass–through rates for bulky products are higher than for small products with a similar unit cost.

The observations above indicate that our analytical model, although it is stylized and ignores competition, is helpful in explaining phenomena observed in practice, including phenomena such as negative pass–through rates for which relatively few compelling explanations have been offered in the literature.

We have analyzed a stylized model with many simplifying assumptions in order to gain insights into the first-order effects of retail and consumer stockpiling and the effects of retailer discounting on customers' willingness to pay on the retailer's optimal response to manufacturer trade discounts. Further research is needed to study these issues in contexts with multiple products, competition between retailers and between manufacturers, other types of promotions, and other realistic factors.

Chapter 4

Conclusions

In this dissertation, we have studied two different pricing issues that manufacturers face when retailers offer periodic discounts and customers stockpile in response. The combination of these two factors can exacerbate the well–known bullwhip effect. Our analysis in Chapter 2 explores one pricing scheme that a manufacturer can use to dampen the bullwhip effect by inducing the retailer to order more frequently and simultaneously improve profitability. Retailer discounting can also lead to a reduction in customers' willingness to pay. Our analysis in Chapter 3 explores the question of how a manufacturer's periodic trade discounts affect a retailer's optimal ordering and pricing (discounting) policy when the retailer stockpiles in response to trade discounts, customers stockpile in response to retail discounts, and retail discounts affect customers' willingness to pay. For the model in each chapter, we provide a full characterization of the retailer's optimal response to the manufacturer's pricing decisions as well as the consumers' response to the retailer's pricing scheme. We also perform associated numerical studies which, together with our analytical results, provide insight into how both manufacturers and retailers should make decisions in these problem settings, and circumstances in which various policies are most effective in increasing profit.

In Chapter 2, we propose and analyze a new Pareto–improving pricing scheme in which the manufacturer subsidizes the retailer's order setup (transportation) cost in exchange for a (possibly) higher wholesale price. The retailer responds by choosing regular and discount prices and his order frequency to maximize his own profit in view of the customers' response. The retailer incurs a fixed cost per shipment for transportation and inventory holding costs, in addition to unit purchase costs. If the retailer offers a discount, he offers it immediately upon the arrival of an order from the manufacturer so as to clear some inventory as soon as possible by inducing customers to stockpile. There are two customer segments that differ in their reservation prices and holding costs (which affect their propensity to stockpile). Customers wish to consume the product at a constant rate and make purchase (including stockpiling) decisions to maximize their utility from consumption less purchase costs and inventory holding costs. As such, they do not consume if the gross cost of the product—the sum of the purchase cost and the cost of holding inventory until the product is consumed—exceeds their reservation price.

We derive a detailed characterization of the retailer's optimal pricing policy for each of three dominant pricing strategies. In the first dominant strategy, the retailer sets a high regular price (equal to the higher of the reservation prices for the two segments) and may offer a discount, but the discount is not large enough to entice customers with the lower reservation price to buy. In the second dominant strategy, the retailer also sets a high regular price but discounts deeply enough to entice customers with the lower reservation price to purchase when the discount is offered. In the third dominant strategy, the retailer sets a low regular price (equal to the lower of the reservation prices for the two segments) and may also offer a discount. For each pricing strategy, there is an associated optimal order cycle duration. We derive conditions that define regions in the two-dimensional space of values of the transportation subsidy and wholesale price in which the retailer prefers each of his dominant pricing strategies.

The manufacturer solves an optimization problem to maximize his own profit within each region taking into account the retailer's response, and selects the pricing scheme that maximizes his own profit across the three regions while ensuring that the retailer is no worse off than under the current pricing scheme. For different manufacturer pricing policies, the retailer may choose a completely different pricing strategy which can lead to a substantially different demand and profit for the manufacturer.

We performed an extensive numerical study that enabled us to identify the types of changes in the retailer's decisions that commonly occur in response to changes in the manufacturer's pricing scheme, and circumstances that make pricing changes profitable for the manufacturer. Below, we discuss some of the underlying intuition that provides a common thread for our findings. It is clear that any transportation subsidy and any increase in the wholesale price (which then increases the retailer's holding cost per unit) will lead the retailer to order more frequently. Even with this increase in order frequency, the retailer's net transportation costs decrease, but the increased order frequency also provides the side–benefit of a reduction in inventory levels. So, a $1 increase in the transportation subsidy may provide the retailer more than a $1 reduction in his overall operating costs. If these savings are not offset by an increase in the wholesale price, the retailer may find it profitable to switch to a different pricing strategy that embraces the customer segment with the lower reservation price. For example, he may switch from a policy with everyday high pricing, under which only the segment with the higher reservation price purchases, to a policy in which the regular price is high but deep discounts are offered periodically that entice customers with the lower reservation price to buy. As another example, the retailer may switch from a pricing policy in which the regular price is high and deep discounts are offered periodically (as described above) to a policy with everyday low pricing under which both segments purchase at a steady rate and never forego consumption because the price is always below their reservation prices. Both of the switches in the retailer's pricing strategy that we have just described lead to an increase in the manufacturer's overall demand, which provides a potential opportunity for an increase in profits, also.

A transportation subsidy may also help the manufacturer to mitigate the bullwhip effect. Although demand is deterministic in our model, in settings with uncertain demand, shorter retailer's order cycles mean that both the retailer and the manufacturer need to forecast demand over a shorter time horizon, which helps to reduce forecast errors and their amplification, thereby mitigating the bullwhip effect.

In Chapter 3, we investigate the retailer's pass-through of the manufacturer's trade discounts. In our model, the manufacturer offers a fixed wholesale price and periodic trade discounts. The retailer optimizes his ordering plan (including stockpiling when the manufacturer offers a discount) and the pattern of discounts to be offered to customers. The retailer orders in batches due to a fixed transportation cost per shipment and incurs holding costs. Therefore, if he offers discounts, he offers them immediately upon arrival of orders from the manufacturer so as to clear some inventory as soon as possible by inducing customers to stockpile. Customers differ in their reservation prices. We account for the fact that retail discounts reduce consumers' reservation prices, which we call the (negative) brand equity effect. To the best of our knowledge, ours is the first model to include the brand equity effect together with retailer and consumer stockpiling in the retailer's decision model for choosing discounting and ordering patterns when a manufacturer offers periodic trade discounts.

For a given frequency and depth of the manufacturer's trade discount, we characterize the retailer's optimal discounting pattern for a given retailer ordering schedule that spans the time between the manufacturer's trade discount offers. The brand equity effect links the retailer's discount decisions in different periods, resulting in a non-concave optimization problem. We are able to surmount this difficulty by obtaining a characterization of the retailer's optimal discounts at times when there is no trade discount in effect, and using this result to develop a nested optimization problem whose solution provides the optimal pattern of retail discounts. We also show that the retailer should stockpile and schedule his next order later than he would otherwise whenever the manufacturer offers a trade discount; this result helps to eliminate dominated ordering schedules. We solve for the retailer's jointly optimal ordering and discounting patterns by enumerating the non-dominated ordering schedules and optimizing the retailer's discounting pattern for each.

We obtained many interesting findings from our analytical model and a numerical study based on it. One of the structural results from our analysis of the model is that there may be threshold effect in the retailer's pass-through. That is, the retailer may pass none of the manufacturer's trade discount on to consumers—owing to the loss of brand equity— unless the trade discount is above certain threshold, at which point the retailer's discount exhibits a jump. We show that the threshold effect is due to the brand equity effect; it does not arise in the absence of this effect.

Another finding is that the retailer's discount does not necessarily increase as the manufacturer's trade discount increases. For example, the retailer may choose to offer a large discount when the manufacturer offers a small trade discount if he faces a small portion of customers who are willing to pay the full price. Here, it may be less important to protect brand equity than it is to reach a larger portion of the customer base by offering a deep discount. Under the same conditions, if the manufacturer's trade discount is high, the retailer may use a completely different pricing strategy in which he offers a small discount to protect brand equity and bolster the number of customers who are willing to pay a high price, and sells to this relatively small proportion of customers at a very high margin, owing to large trade discounts. Indeed, we have found from both our analytical and numerical results that the retailer's discount may be increasing, decreasing, increasing then decreasing, decreasing then increasing, or even fluctuating as the manufacturer's discount increases. The patterns are consequences of complex interactions among factors such as the frequency of manufacturer's trade discount, the retailer's and customers' holding costs, and the portion of customers who are willing to pay full price.

From our numerical study, we have found that the retailer offers a high pass-through if (i) he has a strong incentive to reduce his inventory holding costs, (ii) if a large portion of customers will stockpile fairly large quantities or purchase more in response to discounts. Any factors that contribute to either (i) or (ii) increase the retailer's pass through. The retailer has a strong incentive to reduce his inventory holding costs if his inventory holding cost rate is high and/or if the time between manufacturer trade discount offers is long. These two factors make it expensive for him to stockpile in response to the manufacturer's trade discount and he can reduce these costs by offering a discount to customers. A large portion of customers will stockpile fairly large quantities in response to a retailer's discount if the customers' holding costs are relatively low. Customers with low reservation prices who do not purchase at the regular price will stockpile to satisfy consumption for a greater portion of the time (i.e., they forego consumption for a smaller portion of the time) if the time between manufacturer trade discount offers (and thus also the time between retailer discounts) is shorter. The factors contributing to condition (ii) increase the payoff the retailer receives from any given discount, and thus increase his incentive to offer a discount.

The retailer's pass-through rate is an important measure of the efficiency of manufacturer's trade promotions. By including both a brand equity effect and retailer and consumer stockpiling behavior in the retailer's decision model, we provide very important managerial insights for manufacturers about how retailers may respond to their discounts, which has strong implications for how manufacturers should structure their trade discounts.

In this dissertation, we have explored two ways in which, by offering a subsidy (of the transportation cost) or periodic discounts, a manufacturer may be able to improve his market penetration and profitability. In both cases, the benefits of the subsidy or discount are amplified because the subsidy or discount drives the retailer to make operational decisions (e.g., ordering frequency or stockpiling decisions) as well as pricing decisions differently than he would have in the absence of the manufacturer subsidy or discount. For example, in the model in Chapter 2, the transportation subsidy leads the retailer to order more frequently, thereby reducing his inventory holding costs. This, then, sometimes via a very complicated mechanism, allows him to sell profitably to another segment of customers, which yields benefits to the manufacturer as well. In the model in Chapter 3, the manufacturer's periodic trade discounts generally leads the retailer to stockpile to take advantage of the trade discount. But to avoid the associated increase in inventory holding costs, the retailer may elect to offer deeper discounts to customers, which enables the retailer to reach customers who otherwise would not purchase. Although we have found that the retailer rarely passes on a large portion of the manufacturer's trade discount to customers, the retailer responds to the trade discounts by changing his ordering schedule and discount pattern in such a way that his own profit is maximized, and this sometimes leads to deep retail discounts that generate substantial additional sales, which can lead to increased profits for the manufacturer.

As we have seen in this dissertation, manufacturers' pricing policies affect retailers' pricing decisions as well as operational decisions, and opportunities for profit improvement may be significant if the manufacturer carefully designs his pricing policy with a good understanding of these complex interactions.

Bibliography

[1] D. D. Achabal, S. McIntyre, and S. A. Smith. Maximizing profits from periodic department store promotions. *Journal of Retailling*, 66(4):383–407, 1990.

[2] P. Aggarwal and R. Vaidyanathan. Use it or lose it: Purchase acceleartion effects of time-limited promotions. *Journal of Consumer Behavior*, 2(4):393–403, 2003.

[3] K. L. Ailawadi, K. Gedenk, C. Lutzky, and S. A. Neslin. Decomposition of the sales impact on promotion-induced stockpiling. *Journal of Marketing Research*, 44(3):450–467, 2007.

[4] K. L. Ailawadi, K. Gennik, C. Lutzky, and S. A. Neslin. Decomposition of the sales impact of promotion-induced stockpiling. *Journal of Marketing Research*, 44(3):450–468, 2007.

[5] K. L. Ailawadi and B. A. Harlam. Fingdings–retailer promotion pass-through: a measure, its magnitude, and its determinants. *Marketing Science*, 28(4):782–791, 2009.

[6] K. L. Ailawadi and S. A. Neslin. The effect of promotion on consumption: buying more and consuming it faster. *Journal of Marketing Research*, 35(3):390–398, 1998.

[7] M. K. Armstrong. Retail response to trade promotion: An incremental analysis of forward buying and retail promotion. Unpublished doctoral dissertation, University of Texas at Dallas, Richardson, TX, 1991.

[8] J. L. Assuncao and R. J. Meyer. The rational effect of price promotions sales and consumption. *Management Science*, 39(5):517–535, 1993.

[9] F. M. Beasley. An examination of stockpiling behavior in response to price deals. *Academy of Marketing Studies Journal*, 2(1):23–34, 1998.

[10] D. R. Bell, J. Chiang, and V. Padmanabhan. The decomposition of promotional response: an empirical evalutaion. *Management Science*, 18(4):504–526, 1999.

[11] D. R. Bell, G. Iyer, and V. Padmanabhan. Price competition under stockpiling and flexible consumption. *Journal of Marketing Research*, 39(3):292–303, 2002.

[12] D. Besanko, J.-P. Dube, and S. Gupta. Own-brand and cross-brand retail pass-through. *Marketing Science*, 24(1):123–137, 2005.

[13] R. C. Blattberg, G. D. Eppen, and J. Lieberman. The theoretical and empirical evaluation of price deals for consumer nondurables. *Journal of Marketing*, 45(1):116–129, 1981.

[14] R. A. Briesch, L. Krishnamurthi, T. Mazumdar, and S. Raj. A comparative analysis of reference price models. *The Journal of Consumer Research*, 24(2):202–214, 1997.

[15] B. J. Bronnenberg, V. Mahajan, and W. R. Vanhonacker. The emergence of market structure in new repeat-purchase categories: The interplay of market share and retailer distribution. *Marketing Research*, 37(1):16–31, 2000.

[16] S. Bucovetsky. Price dispersion and stockpiling by consumers. *Review of Economic Studies*, 50(162):443–465, 1983.

[17] G. Cachon. Managing supply chain demand variability with scheduled ordering policies. *Management Science*, 45(6):1091–1108, 1999.

[18] P. Chandon and B. Wansink. When are stockpiled products consumed faster? a convenience-salience framework of postpurchase consumption incidence and quantity. *Journal of Marketing Research*, 39(3):321–335, 2002.

[19] M. Chevalier and R. C. Curhan. Retail promotions as a function of trade promotions: A descriptive analysis. *Sloan Management Review*, Fall:19–32, 1976.

[20] P. K. Chintagunta. Investigating category pricing behavior at a retail chain. *Journal of Marketing Research*, 39(2):141–154, 2002.

[21] R. Croson and K. Donohue. Behavioral causes of the bullwhip effect and the observed value of inventory information. *Management Science*, 52(3):323–336, 2004.

[22] I. S. Currim and L. G. Schneider. A taxonomy of consumer purchase strategies in a promotion intensive environment. *Marketing Science*, 10(2):91–110, 1991.

[23] M. Dada and K. N. Srikanth. Pricing policies for quantity discounts. *Management Science*, 33(10):1247–1252, 1987.

[24] M. G. Dekimpe and D. M. Hanssens. Sustained spending and persistent response: A new look at long-term marketing profitability. *Journal of Marketing Research*, 36(4):397–412, 1999.

[25] M. G. Dekimpe, D. M. Vincent, R. Nijs, and J.-B. E. Steenkamp. Measuring short- and long-run promotional effectiveness on scanner data using persistence modelling. *Applied Stochastic Models in Business and Industry*, 21:409–416, 2005.

[26] R. J. Dolan. Quantity discounts: Managerial issues and research opportunities. *Marketing Science*, 6(1):1–22, 1987.

[27] K. Golabi. Optimal inventory policies when ordering prices are random. *Operations Research*, 33(3):575–588, 1985.

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

[29] S. Gupta. Impact of sales promotions on when, what, and how much to buy. *Journal of Marketing Research*, 25(4):342–355, 1988.

[30] K. Helsen and D. C. Schmittlein. Some characteristics of stockpiling behavior under uncertainty. *Marketing Letters*, 3(1):5–17, 1992.

[31] H. Helson. *Adaptation-level Theory.* New York: Harper and Row Publishers, Inc., 1964.

[32] I. Hendel and A. Nevo. The post-promotion dip puzzle: What do the data have to say? *Quantitative Marketing and Economics*, 1(4):409–424, 2003.

[33] I. Hendel and A. Nevo. Intertemporal substitution and storable products. *Journal of the European Economic Associationn*, 2(2/3):536–547, 2004.

[34] W. Holland and M. S. Sodhi. Quantifying the effects of batch size and order errors on the bullwhip effect using simulation. *International Journal of Logistics: Research and Applications*, 7(3):251–261, 2004.

[35] A. P. Jeuland and C. Narasimhan. Dealing-temporary price custs-by a seller and a buyer discrimination mechanism. *Journal of Business*, 58(3):295–308, 1985.

[36] A. P. Jeuland and S. M. Shugan. Managing channel profits. *Marketing Science*, 2(3): 239–272, 1983.

[37] M. Johansson, P. Cowling, and B. McCarthy. Synchronous supply. *Manufacturing Engineering*, 79(1):36–40, 2000.

[38] G. Kalyanaram and R. S. Winer. Empirical generalizations from reference price research. *Marketing Science*, 14(3):161–169, 1995.

[39] P. K. Kopalle, C. F. Mela, and L. Marsh. The dynamic effect of discounting on sales: empirical analysis and normative pricing. *Marketing Science*, 18(3):317–332, 1999.

[40] P. K. Kopalle, A. G. Rao, and J. L. Assancao. Asymmetric reference price effects and dynamic pricing policies. *Marketing Science*, 15(1):60–85, 1996.

[41] A. Krishna. The impact of dealing patterns on purchase behavior. *Marketing Science*, 13 (4):351–373, 1994.

[42] A. Krishna and G. V. Johar. Consumer perceptions of deals: Biasing effects of varying deal prices. *Journal of Experimental Psychology: Applied*, 2(3):187–206, 1996.

[43] L. Krishnamurthi, T. Mazumdar, and S. P. Raj. Asymmetric response to price in consumer brand choice and purchase quantity decisions. *Journal of Consumer Research*, 19(3):387–400, 1992.

[44] N. Kuman, S. Rajiv, and A. Jeuland. Effectiveness of trade promotions: Analyzing the determinants of retail pass through. *Marketing Science*, 20(4):382–404, 2001.

[45] R. Lal, J. D. C. Little, and J. M. Villas-Boas. A theory of forward buying, merchandising and trade deals. *Marketing Science*, 15(1):21–37, 1996.

[46] R. Lal and R. Staelin. An approach for developing an optimal discount policy. *Management Science*, 30(12):1524–1539, 1984.

[47] J. M. Lattin and R. E. Bucklin. Reference effects of price and promotion on brand choice behavior. *Journal of Marketing Research*, 26(3):299–310, 1989.

[48] H. L. Lee, V. Padmanabhan, and S. Whang. The bullwhip effect in supply chains. *Sloan Management Review*, 38(3):93–102, 1997.

[49] G. E. Mayhew and R. S. Winer. An empirical analysis of internal and external reference prices using scanner data. *Journal of Consumer Research*, 19(1):62–70, 1992.

[50] C. F. Mela, K. Jedidi, and D. Bowman. The long-term impact of promotions on consumer stockpiling behavior. *Journal of Marketing Research*, 35(2):250–262, 1998.

[51] R. J. Meyer and J. Assuncao. The optimality of consumer stockpiling strategies. *Marketing Science*, 9(1):18–41, 1990.

[52] S. Meza and K. Sudhir. Pass-through timing. *Quantitative Marketing and Economics*, 4 (4):351–382, 2006.

[53] G. Miragliotta. Layers and mechanisms: A new taxonomy for the bullwhip effect. *International Journal of Production Economics*, 104(2):365–381, 2006.

[54] K. B. Monroe. Buyers' subjective perceptions of price. *Journal of Marketing Research*, 10 (1):70–80, 1973.

[55] S. Moorthy. A general theory of pass-through in chanels with category management and retail competition. *Marketing Science*, 24(1):110–122, 2005.

[56] V. Nijs, K. Misra, E. T. Anderson, K. Hansen, and L. Krishnamurthi. Channel pass-through of trade promotions. *Marketing Science*, 29(2):250–267, 2010.

[57] V. R. Nijs, M. G. Dekimpe, J.-B. E. M. Steenkamp, and D. M. Hanssens. The category-demand effects of price promotions. *Marketing Science*, 20(1):1–22, 2001.

[58] K. Pauwels. How retailer and competitor decisions drive the long-term effectiveness of manufacturer promotions for fast moving consumer goods. *Journal of Retailing*, 83(3): 297–308, 2006.

[59] K. Pauwels, D. M. Hanssens, and S. Siddarth. The long-term effects of price promotions on category incidence, brand choice, and purchase quantity. *Journal of Marketing Research*, 39(4):421–439, 2002.

[60] A. Potter and S. M. Disney. Bullwhip and batching: An exploration. *International Journal of Production Economics*, 104(2):408–418, 2006.

[61] J. S. Raju. Theoretical models of sale promotions: Contribution, limitations and a future research agenda. *European Journal of Operational Research*, 85(1):1–17, 1995.

[62] C. E. Riddalls and S. Bennett. The optimal control of batched production and its effect on demand amplification. *International Journal of Production Economics*, 72(2):159–168, 2001.

[63] R. Schmalensee. Monopolistic two-part pricing arrangements. *Bell Journal of Economics*, 12(2):445–466, 1981.

[64] A. G. Sogomonian and C. S. Tang. A modeling framework for coordinating promotion and production decisions within a firm. *Management Science*, 39(3):191–203, 1993.

[65] J. D. Sterman. Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Science*, 35(3):321–339, 1989.

[66] K. Sudhir and S. Datta. Pricing in marketing channels. In V. R. Rao, editor, *Handbook of Pricing Research in Marketing*. Edward Elgar Publishing, Cheltenham, U.K., 2009.

[67] R. K. Tyagi. A characterization of retailer response to manufacturer trade deals. *American Marketing Association*, 36(4):510–516, 1999.

[68] R. G. Walters. An empirical investigation into retailer respnose to manufacturer trade promotions. *Journal of Retailing*, 65(2):253–272, 1989.

[69] R. Warburton. An analytical investigation of the bullwhip effect. *Production and Operations Management*, 13(2):150–160, 2004.

[70] Z. K. Weng. Channel coordination and quantity discounts. *Management Science*, 41(9): 1509–1522, 1995.

[71] R. S. Winter. A reference price model of brand choice for frequently purchased products. *Journal of Consumer Research*, 13(2):250–256, 1986.

[72] D. Wright and X. Yuan. Mitigating the bullwhip effect by ordering policies and forecasting methods. *International Journal of Production Economics*, 113(2):587–597, 2008.

Appendix A: Derivation of the Retailer's Optimal Solution for Δ^{UB} and Δ_0 when Consumers Stockpile Extensively

The formulas for Π_1 in (2.1) and Π_3 in (2.3) are based on the assumption that the stockpiling duration of both segments is less than or equal to T. If the retailer increases Δ so that one segment would be willing to stockpile more, that segment will receive the additional discount on all units purchased and the retailer garners no additional inventory savings related to that segment because the segment will not stockpile further. In these cases, we need to modify the objective function and solution. The question is whether the retailer's net benefit (inventory savings less reduction in revenue from the additional discount) due to the *other* segment's additional stockpiling is large enough to compensate. For both Π_1 and Π_3, the retailer's objective depends on whether h_1 or h_2 is smaller; we divide our analysis accordingly.

Analysis of Π_1 for Large Δ

In this case, when Δ reaches $\min\{h_1, h_2\}T$, one segment stockpiles for the entire order cycle.

Case (i): $h_1 \leq h_2$

In this case, $t_2 = \Delta/h_2 < t_1 = T$ because segment 1's stockpiling duration exceeds that of segment 2, which is $t_2 = \Delta/h_2$ ($< T$). Substituting these into the objective, Π_1, specialized for these circumstances, we obtain:

$$\Pi_1(\Delta, T) = (r_2 - w)(\lambda_1 + \lambda_2) - \frac{K + \frac{\lambda_2}{h_2}(1 - \frac{h_r}{2h_2})\Delta^2}{T} - \lambda_1\Delta - \frac{h_r}{2}\lambda_2 T.$$

Subcase (a): $1 - \frac{h_r}{2h_2} > 0$

In this subcase, the specialized Π_1 is decreasing in Δ so we do not want to increase Δ beyond $h_{min}T$; thus $\Delta^* = h_1 T$. Hence the optimal solution is

$$T^* = \sqrt{\frac{K}{h_1\lambda_1 + 0.5h_r\lambda_2 + \frac{\lambda_2}{h_2}(1 - \frac{h_r}{2h_2})h_1^2}} \quad \text{and} \quad \Delta^* = \Delta^{UB} = h_1 T^*. \tag{A-1}$$

The optimal objective is

$$\Pi_1^* = (r_2 - w)(\lambda_1 + \lambda_2) - 2\sqrt{h_1\lambda_1 + 0.5h_r\lambda_2 + \frac{\lambda_2 h_1^2}{h_2}(1 - \frac{h_r}{2h_2})}\sqrt{K'}. \tag{A-2}$$

Subcase (b): $1 - \frac{h_r}{2h_2} < 0$.

In this subcase, the specialized Π_1 is neither concave nor convex, and we find that the unique stationary point is a saddle point by checking the second order conditions at the stationary point. Therefore, the optimal solution must be on the boundary, i.e., at the lower bound $\Delta = h_1 T^*$ or at the upper bound $\Delta = h_2 T^*$. The optimal solution at the lower bound was derived above, and the optimal solution at the upper bound (note that in this case, both t_1 and t_2 are equal to T) is:

$$T = \sqrt{\frac{K}{(\lambda_1+\lambda_2)h_2}} \quad \text{and} \quad \Delta^{UB} = h_2 T.$$

The associated objective value is $\Pi_1^* = (r_2 - w)(\lambda_1 + \lambda_2) - 2\sqrt{h_2(\lambda_1 + \lambda_2)}\sqrt{K'}$. This solution is dominated by $\Pi_1(0)$ in (2.5) (because $2h_2 < h_r$). But under the conditions of this subcase (i.e., $h_1 \leq h_2$ and $1 - \frac{h_r}{2h_2} < 0$), $\Pi_1(0)^*$, is dominated by the solution with $\Delta = h_1 T^*$ (i.e., the lower bound on Δ), whose profit is the same at that given in (A-2). Therefore, the solution with $\Delta = \Delta^{UB}$ is dominated, and the optimal solution is that given in (A-1).

Case (ii): $h_1 > h_2$

In this case, $t_1 = \Delta/h_1$ and $t_2 = T$. The objective, Π_1, specialized to these circumstances, is

$$\Pi_1(\Delta, T) = (r_2 - w)(\lambda_1 + \lambda_2) - \frac{K + \frac{\lambda_1}{h_1}(1 - \frac{h_r}{2h_1})\Delta^2}{T} - \lambda_2\Delta - \frac{h_r}{2}\lambda_1 T.$$

We can obtain solutions using analysis similar to that for Case (i) above.

Subcase (a): $1 - \frac{h_r}{2h_1} > 0$

In this subcase, the optimal solution is

$$T^* = \sqrt{\frac{K}{h_2\lambda_2 + 0.5h_r\lambda_1 + \frac{\lambda_1 h_2^2}{h_1}(1 - \frac{h_r}{2h_1})}} \quad \text{and} \quad \Delta^* = \Delta^{UB} = h_2 T^*.$$

The optimal objective is

$$\Pi_1^* = (r_2 - w)(\lambda_1 + \lambda_2) - 2\sqrt{h_2\lambda_2 + 0.5h_r\lambda_1 + \frac{\lambda_1 h_2^2}{h_1}(1 - \frac{h_r}{2h_1})}\sqrt{K'}.$$

Subcase (b): $1 - \frac{h_r}{2h_1} < 0$

The analysis for this subcase parallels that of Case (i), subcase (b). The stationary point is a saddle point. The solution with Δ at its upper bound is $T^* = \sqrt{\frac{K}{(\lambda_1+\lambda_2)h_1}}, \Delta^* = \Delta^{UB} = h_2 T^*, \Pi_1^* = (r_2 - w)(\lambda_1 + \lambda_2) - 2\sqrt{h_1(\lambda_1 + \lambda_2)}\sqrt{K'}$, and this solution is dominated by $\Pi_1^*(0)$ (because $2h_1 < h_r$), which in turn is dominated by the solution with Δ at its upper bound.

Analysis of Π_3 for Large Δ

In this case, at $\Delta = \min\{h_1/T, h_2/T + r_1 - r_2\}$, one segment's stockpiling duration reaches T.

Case (i): $h_1 \leq h_2$

In this case, we have $t_1 = \frac{\Delta}{h_1} > t_2 = \frac{\Delta-(r_1-r_2)}{h_2}$, so segment 1's stockpiling duration is longer than that of segment 2. Substituting $t_1 = T$ and $t_2 = \frac{\Delta-(r_1-r_2)}{h_2}$ into the objective, we have

$$\Pi_3(\Delta, T) = (r_1 - w)\lambda_1 - \lambda_1\Delta - \frac{K + \frac{\lambda_2}{h_2}[\Delta^2 - (2r_1 - r_2 - w)\Delta + (r_1 - w)(r_1 - r_2)]}{T} \qquad \text{(A-3)}$$

We define $g^{UB}(\Delta) = \frac{\lambda_2}{h_2}[\Delta^2 - (2r_1 - r_2 - w)\Delta + (r_1 - w)(r_1 - r_2)]$. Notice that g^{UB} has the same structural form as $g(\Delta)$ in (2.4). We can rewrite Π_3 as

$$\Pi_3(\Delta, T) = (r_1 - w)\lambda_1 - \lambda_1\Delta - \frac{K + g^{UB}(\Delta)}{T}.$$

Like our assumption that $K + g(\Delta) > 0$, we assume that $K + g^{UB}(\Delta) > 0$ so T should be set as large as possible, i.e., to the upper bound, Δ/h_1. Substituting $T = \Delta/h_1$, the objective can be rewritten as $\Pi_3 = (r_1 - w)\lambda_1 + \frac{2\lambda_2}{h_1}(r_1 - r_2 - w) - (\lambda_1 + h_1\lambda_2/h_2)\Delta - \frac{h_1 K + \frac{\lambda_2 h_1}{h_2}(r_1-w)(r_1-r_2)}{\Delta}$, which is concave in Δ. The optimal unconstrained solution (stationary point) is

$$\Delta^* = \sqrt{\frac{h_1 h_2 K + h_2(r_1 - w)(r_1 - r_2)\lambda_2 h_1}{h_2\lambda_1 + h_1\lambda_2}} \quad \text{and} \quad T^* = \Delta^*/h_1. \tag{A-4}$$

We need to check that the stationary point for Δ^* is in the relevant range, i.e., greater than $r_1 - r_2$. If not, then the optimal constrained solution is $\Delta^* = r_1 - r_2$ and $T^* = \Delta^*/h_1$.

Case (ii): $h_1 > h_2$

In this case, we have $\frac{\Delta}{h_1} \geq \frac{\Delta-(r_1-r_2)}{h_2}$ if $\Delta \leq \hat{\Delta} = \frac{h_1(r_1-r_2)}{h_1-h_2}$ and $\frac{\Delta}{h_1} \leq \frac{\Delta-(r_1-r_2)}{h2}$ otherwise. In other words, segment 1 stockpiles more than segment 2 if Δ is less than a threshold value, $\hat{\Delta}$, and stockpiles less than segment 2 if Δ is larger. Hence we need to optimize Δ in the two subintervals and choose the better solution.

Subcase (a): $\Delta \leq \hat{\Delta}$

We have $t_1 = T$ and $t_2 = \frac{\Delta-(r_1-r_2)}{h_2}$ in this case, hence the specialized Π_3 is exactly the same as (A-3). We need to check whether the stationary point for Δ in (A-4) is less than or equal to $\hat{\Delta}$. If so, the solution is optimal; otherwise, the optimal solution is on the boundary:

$$\Delta^* = \hat{\Delta} \quad \text{and} \quad T^* = \Delta^*/h_1.$$

Subcase (b): $\Delta \geq \hat{\Delta}$

In this case, we have $t_1 = \frac{\Delta}{h_1}$ and $t_2 = T$, so Π_3 specialized to this case is

$$\Pi_3(\Delta, T) = (r_1 - w)(\lambda_1 + \lambda_2) - \frac{K + \frac{\lambda_1}{h_1}(1 - \frac{h_r}{2h_1})\Delta^2}{T} - \frac{h_r}{2}\lambda_1 T - \lambda_2\Delta. \tag{A-5}$$

If $1 - \frac{h_r}{2h_1} > 0$, then the objective is decreasing in Δ so we should set Δ at the lower bound, $h_2T + r_1 - r_2$. With this substitution, the objective can be rewritten as

$$\Pi_3(\Delta, T) = \Theta - [\lambda_2 h_2 + \frac{\lambda_1 h_2^2}{h_1}(1 - \frac{h_r}{2h_1}) + \frac{h_r\lambda_1}{2}]T - \frac{K + \frac{\lambda_1}{h_1}(1 - \frac{h_r}{2h_1})(r_1 - r_2)^2}{T},$$

where $\Theta = (r_1 - w)(\lambda_1 + \lambda_2) - \lambda_2(r_1 - r_2) - \frac{2\lambda_1}{h_1}h_2(1 - \frac{h_r}{2h_1})(r_1 - r_2)$. The objective function is concave in T and the unconstrained optimal solution is

$$T^* = \sqrt{\frac{\frac{\lambda_1}{h_1}(1-\frac{h_r}{2h_1})(r_1-r_2)^2+K}{\lambda_2 h_2+\frac{\lambda_1 h_2^2}{h_1}(1-\frac{h_r}{2h_1})+h_r\lambda_1/2}} \quad \text{and} \quad \Delta^* = h_2T^* + (r_1 - r_2).$$

It is easy to see that the Δ^* derived above must be greater than $r_1 - r_2$. On the other hand, if the stationary point for Δ^* is less than $\hat{\Delta}$, then the optimal solution is

$$\Delta^* = \hat{\Delta} \quad \text{and} \quad T^* = \frac{\hat{\Delta}-(r_1-r_2)}{h_2}.$$

If $1 - \frac{h_r}{2h_1} < 0$, the objective is neither monotonic nor concave in Δ, and the stationary point is a saddle point. Hence we need to check the boundary solutions $\Delta = h_1T, h_2T$ and $\hat{\Delta}$. (It is easy to show that $\hat{\Delta} \geq r_1 - r_2$, so $r_1 - r_2$ is a redundant boundary in this case.) The solution with Δ at the lower bound (i.e., h_2T) was derived above. The objective function with Δ at its upper bound (i.e., h_1T) is

$$\Pi_3(\Delta, T) = (r_1 - w)(\lambda_1 + \lambda_2) - (\lambda_1 + \lambda_2)h_1T - \frac{K}{T}.$$

The unconstrained optimal solution is

$$T^* = \sqrt{\tfrac{K}{(\lambda_1+\lambda_2)h_1}} \quad \text{and} \quad \Delta^* = h_1 T^*.$$

If the Δ^* derived above is less than $\hat{\Delta}$, then it can be eliminated from consideration.

By substituting $\hat{\Delta}$ for Δ in (A-5) and optimizing the resulting concave function with respect to T, we obtain the solution at $\Delta = \hat{\Delta}$ as

$$\Delta = \hat{\Delta}, T^* = \sqrt{\frac{K + (\lambda_1/h_1)(1 - \frac{h_r}{2h_1})\hat{\Delta}^2}{0.5 h_r \lambda_1}}.$$

We have now derived solutions with Δ at the three boundaries. The boundary solution with the largest objective value is the optimal solution to this subproblem.

Appendix B: Proof of Lemma 1

Lemma 1: If $\beta < 0$, then there exists a $\Delta > r_1 - r_2$ such that $g(\Delta) < \beta * \Delta^2$.

Proof: We can re-express $g(\Delta)$ as:

$$(\frac{\lambda_1}{h_1} - \frac{h_r \lambda_1}{2h_1^2})\Delta^2 + \frac{\lambda_2}{h_2}[\Delta^2 - (2r_1 - r_2 - w)\Delta + (r_1 - w)(r_1 - r_2)].$$

The first term is simply $\beta * \Delta^2$ so $g(\Delta) < \beta * \Delta^2$ if the expression in square brackets above is negative. Let $\Delta = r_1 - r_2 + \epsilon$. Then, with a little algebra, we can show that the expression is negative for all $\epsilon < r_2 - w$. Thus, there exists a $\Delta > r_1 - r_2$ such that $g(\Delta) < \beta * \Delta$. □

Appendix C: Proof of Lemma 2

Lemma 2: Under the conditions of Case I, specifically $h_r > 2h_1$, $g(\Delta_0) < 0$.

Proof: Substituting for Δ_0 in $g(\Delta_0)$, after some algebra, we obtain:

$$g(\Delta_0) = -\frac{[(\lambda_2/h_2)(2r_1 - r_2 - w)]^2}{4(\sum \frac{\lambda_i}{h_i} - \frac{h_r \lambda_1}{2h_1^2})} + (r_1 - w)(r_1 - r_2)\lambda_2/h_2.$$

The condition $h_r > 2h_1$ implies that the denominator of the first term is less than λ_2/h_2, so the first term is less than $-(\lambda_2/h_2)(2r_1 - r_2 - w)^2/4$. With a few straightforward algebraic steps, we can show that $g(\Delta_0) < -(w - r_2)^2/4 < 0$. □

Appendix D: Derivation of Results in Tables 2.2 and 2.3 in Section 2.4.1

We first derive the optimal discount for each of the retailer's profit functions under the assumptions of Section 2.4.1, i.e., that neither segment's stockpiling duration is more than T. In Appendix A, we address situations where one of the segments is willing to stockpile for a duration exceeding T.

Π_1 (see (2.1)) is decreasing in Δ if $\alpha > 0$, in which case Δ^* is the lower limit of the feasible interval (i.e., zero). The condition $\alpha < 0$ implies that the retailer prefers to increase Δ so long as additional discounting induces *both* segments to stockpile more. We know that it is suboptimal for the retailer to choose Δ such that both segments would be willing to stockpile to satisfy their consumption for *more* than T. When Δ reaches $\min\{h_1, h_2\}T$, one of the two segments stockpiles for the entire replenishment interval. This defines Δ^{UB} under the assumptions of Section 2.4.1.

Π_2 (see (2.2)) is increasing in Δ if $\beta < 0$ and decreasing in Δ otherwise. Thus, Δ^* is, accordingly, the lower limit (i.e., zero) or the upper limit (i.e., $r_1 - r_2$) of the feasible interval.

We now consider Π_3 (see (2.4)). The only term in Π_3 that depends on Δ is g. The stationary point of g is $\Delta_0 = \frac{\lambda_2(2r_1-r_2-w)}{2h_2(\sum \lambda_i/h_i - h_r\lambda_1/2h_1^2)}$. Notice that the numerator of Δ_0 is positive so the sign of Δ_0 depends upon the sign of the denominator, which is the same as the sign of γ. Recall that g appears as $-g(\Delta)$ in Π_3. If $\gamma > 0$, then $\Delta_0 > 0$ and $-g$ is concave. Consequently, Δ_0, the stationary point, is the optimal discount if it is feasible (i.e., $\Delta_0 > r_1 - r_2$). On the other hand, if $\gamma < 0$, then it is straightforward to show that $-g$ is convex and increasing for all $\Delta > 0$. (The stationary point is negative.)

Further analysis allows us to show dominance of certain policies over others: For Cases I and II, Π_2 achieves its maximum at the upper limit of its feasible region, $r_1 - r_2$. It is easy to show that $\Pi_2(r_1 - r_2) = \Pi_3(r_1 - r_2)$. It can also be shown (see proof of Lemma 1 in Appendix B) that if $\beta < 0$, there exists $\Delta > r_1 - r_2$ such that $g(\Delta) < 0$ and $g(\Delta) < \beta * \Delta^2$. Furthermore, Π_3 is increasing for Δ up to $\Delta_0 > r_1 - r_2$. Therefore $\Pi_3(\Delta_0)$ dominates $\Pi_2(r_1 - r_2)$ if $\beta < 0$.

For Case V, $\beta < 0$ implies that Π_2 is increasing in Δ, and $\gamma < 0$ implies that Π_3 is also increasing in Δ. Noting that $\Pi_2(r_1 - r_2) = \Pi_3(r_1 - r_2)$, we can conclude that $\Pi_3(\Delta_{ub})$ dominates $\Pi_2(r_1 - r_2)$.

Appendix E: Lemma 3

Lemma 3: *If Π^L is monotonically increasing in the interval $[a, b]$, convex in $[a, \Delta_0]$ and concave in $[\Delta_0, b]$, then for any two points $\Delta_1, \Delta_2 \in [a, b]$ such that $\frac{\partial \Pi^L}{\partial \Delta^L}(\Delta_1) = \frac{\partial \Pi^L}{\partial \Delta^L}(\Delta_2)$, we have $\Pi^L(\frac{\Delta_1+\Delta_2}{2}) = 0.5\Pi^L(\Delta_1) + 0.5\Pi^L(\Delta_2)$.*

Proof. There exist at most two distinct positive solutions, say Δ_1 and Δ_2, to the first order condition $\frac{\partial \Pi^L}{\partial \Delta^L} = 3a^L{\Delta^L}^2 + 2b^L\Delta^L + c^L = \mu$, and the sum of Δ_1 and Δ_2 is $\frac{-2b^L}{3a^L} = \frac{4}{3}b^L$. Given that Δ_0 is the inflection point of Π^L, we know that $\Delta_0 = \frac{-b^L}{3a^L} = \frac{2}{3}b^L$. Therefore, $\frac{\Delta_1+\Delta_2}{2} = \Delta_0$. What remains to be shown is that $\Pi^L(\Delta_0) = 0.5\Pi^L(\Delta_1) + 0.5\Pi^L(\Delta_2)$.

As shown in Figure 4.1, the slope of the segment connecting $\Pi^L(\Delta_1)$ and $\Pi^L(\Delta_2)$, which we define as α, is $\frac{\Pi^L(\Delta_2)-\Pi^L(\Delta_1)}{\Delta_2-\Delta_1}$. Hence the segment connecting $\Pi^L(\Delta_1)$ and $\Pi^L(\Delta_2)$ can be written as $g(\Delta^L) = \Pi^L(\Delta_1) + \alpha(\Delta^L - \Delta_1)$, and we have $g(\Delta_0) = 0.5\Pi^L(\Delta_1) + 0.5\Pi^L(\Delta_2)$.

To show that $\Pi^L(\Delta_0) = g(\Delta_0)$ is equivalent to show $\Pi^L(\Delta_0) = \Pi^L(\Delta_1) + \alpha(\Delta_0 - \Delta_1) = \Pi^L(\Delta_1) + \frac{\Pi^L(\Delta_2)-\Pi^L(\Delta_1)}{\Delta_2-\Delta_1}(\Delta_0 - \Delta_1) = \Pi^L(\Delta_1) + 0.5[\Pi^L(\Delta_2) - \Pi^L(\Delta_1)]$. Applying the explicit form of Π^L, the above equation can be written (for both the cases of $m \le \overline{R} - p$ and $m > \overline{R} - p$) as $0.25(\Delta_1 + \Delta_2) + b^L(\Delta_1 - \Delta_2)^2 = (\Delta_1)^3 + (\Delta_2)^3$. Given $\Delta_1 + \Delta_2 = \frac{4}{3}b^L$, the above equation can be written as $\frac{16}{27}(b^L)^2 + b^L(\Delta_1 - \Delta_2)^2 = \frac{4}{3}b^L[(\Delta_1)^2 - \Delta_1\Delta_2 + (\Delta_2)^2]$. The above equation can be further simplified to be $\frac{16}{9}(b^L)^2 = (\Delta_1 + \Delta_2)^2$. The left hand side is equal to the right hand side. □

Appendix F: Proof of Proposition 1

Let $(\Delta^{H^*}, \Delta_i^{L^*})$ denote the optimal solution to the retailer's problem. To prove the main result, we first need to establish that $\frac{\partial \Pi^L}{\partial \Delta^L}(\Delta_i^{L^*}) \geq 0$ for all $\Delta_i^{L^*} > 0$, $i \in \{1, \ldots, N\}$.

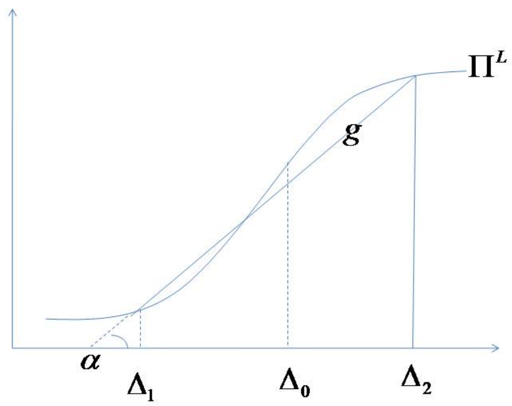

Figure 4.1: Π^L

Let m^* denote the optimal value of m (i.e., m^* satisfies $\Delta^{H^*} + \sum_{i=1}^N \Delta_i^{L^*} = m^*T$). Then Δ^{H^*} and $\Delta_i^{L^*}$ can be expressed as functions of m^* in the nested problem. Suppose that for some $i \in \{1, ..., N\}$, we have $\frac{\partial \Pi^L}{\Delta^L}(\Delta_i^{L^*}) < 0$ and $\Delta_i^{L^*} > 0$. Then there exists $\tilde{m} = m^* - \epsilon/T$ and $\tilde{\Delta}_i^L = \Delta_i^{L^*} - \epsilon$, for some small $\epsilon > 0$ such that that $\Pi^L(\Delta^{H^*}, \Delta_1^{L^*}, \ldots, \Delta_N^{L^*}) < \Pi^L(\Delta^{H^*}, \Delta_1^{L^*}, \ldots, \Delta_{i-1}^{L^*}, \tilde{\Delta_i^L}, \Delta_{i+1}^{L^*}, \ldots, \Delta_N^{L^*})$. This is true because decreasing $\Delta_i^{L^*}$ will improve the profit in the i^{th} low cycle because $\frac{\partial \Pi^L}{\Delta^L}(\Delta_i^{L^*}) < 0$ and a reduction in m will improve the profit in the other cycles because of higher brand equity. This contradicts the optimality of m^*. Therefore, we must have $\frac{\partial \Pi^L}{\Delta^L}(\Delta_i^{L^*}) \geq 0$.

Now, we prove the main result, namely that the positive discounts in low cycles must be equal to each other in the optimal solution. Let λ^H be the Lagrange multiplier associated with the constraint $\Delta^H \geq 0$ and λ_i^L be the Lagrange multiplier associated with the constraint $\Delta_i^L \geq 0$, and μ be the Lagrange multiplier associated with the constraint $\Delta^H + \sum_{i=1}^n \Delta_i^L = mT$. Then the Karush–Kuhn–Tucker conditions for the retailer's problem (holding m fixed) are as follows:

$$-\frac{\partial \Pi^H}{\partial \Delta^H} = \lambda^H - \mu \tag{A-6}$$

$$-\frac{\partial \Pi^L}{\partial \Delta_i^L} = \lambda_i^L - \mu, \ i = 1, \ldots, n \tag{A-7}$$

$$\Delta^H + \sum_{i=1}^n \Delta_i^L = mT \tag{A-8}$$

$$\lambda^H \Delta^H = 0, \quad \lambda_i^L \Delta_i^L = 0, \ i = 1, \ldots, n \tag{A-9}$$

$$\lambda^H \geq 0, \quad \lambda_i^L \geq 0, \ i = 1, \ldots, n \tag{A-10}$$

The conditions in (A-7) and (A-9) imply that for any $i \in 1, ..., n$, we have either ($\Delta_i^{L^*} = 0$ and $\lambda_i^L > 0$) or ($\Delta_i^{L^*} > 0$, $\lambda_i^L = 0$, and $\frac{\partial \Pi^L}{\partial \Delta_i^L} = \mu$). Therefore, we have $\frac{\partial \Pi^L}{\partial \Delta_i^L}(\Delta_i^{L^*}) = \mu \geq 0$ for all positive $\Delta_i^{L^*}$. There are at most two positive solutions that satisfy $\frac{\partial \Pi^L}{\partial \Delta_i^L} = \mu$ due to the cubic form of Π^L. As shown in Figure 3.1, Π^L is either strictly concave or initially convex and then concave in the region of $\Delta^L \geq 0$. If $b^L \leq 0$, then Π^L is strictly concave and there exists one positive solution, which implies that any positive Δ_i^{L*}s are equal to each other. If $b^L > 0$, then Π^L is initially convex, then concave increasing then concave decreasing, so there are two distinct positive solutions to the first order condition $\frac{\partial \Pi^L}{\partial \Delta_i^L}(\Delta_i^{L*}) = \mu$. Denote the two solutions by Δ_1 and Δ_2. By Lemma 3, the solution ($\Delta_i^{L*} = \Delta_1$ and $\Delta_i^{L*} = \Delta_2$) is dominated by ($\Delta_i^{L*} = \Delta_j^{L*} = (\Delta_1 + \Delta_2)/2$). Therefore, we must have $\Delta_i^{L*} = \Delta_j^{L*}$ for any $\Delta_i^{L*}, \Delta_j^{L*} > 0, i, j \in 1, ..., N$. □

Appendix G: Characterization of the Retailer's Dominance Map for Case III in Section 2.4.2

The condition $g(\Delta_0) < 0$ is equivalent to the condition

$$\frac{\lambda_2}{h_2}(r_1 - r_2)^{-1}(\frac{\lambda_1}{h_1} - \frac{h_r\lambda_1}{2h_1^2})^{-1} > 4(r_1 - w)(r_2 - w)^{-2}. \tag{A-11}$$

When $\beta > 0$ (which applies to Cases III and IV), the left hand side of (A-11) is positive. The right hand side of (A-11) is convex and increasing in the applicable range of w ($< r_2$). Let w_2 represent the value of w that satisfies (A-11) as an equality. Then Π_3^* is dominant for $w < w_2$ and Π_2^* is dominant for larger values of w. Thus, for $w > w_2$ the relevant comparison is between Π_1^* and Π_2^*, whereas for $w < w_2$ the relevant comparison is between Π_1^* and Π_3^*.

We also need to ensure that for $w < w_2$, we have $\Delta_0 > r_1 - r_2$; otherwise we would be evaluating Π_3 at an infeasible value of Δ. We address this issue by determining whether all ws satisfying (A-11), i.e., all $w < w_2$, also satisfy the condition $\Delta_0 > r_1 - r_2$. With some algebra, the latter condition can be expressed as:

$$\frac{\lambda_2}{h_2}(r_1 - r_2)^{-1}(\frac{\lambda_1}{h_1} - \frac{h_r\lambda_1}{2h_1^2})^{-1} > 2(r_2 - w)^{-1}. \tag{A-12}$$

The left hand side of this inequality is the same as the left hand side of (A-11). Because $r_1 > r_2$, the right hand side of (A-12) is less than the right hand side of (A-11), so any w satisfying (A-11) also satisfies (A-12). Thus, the condition $\Delta_0 > r_1 - r_2$ imposes no additional constraints.

Expressions for Π_1^* and Π_3^* were derived in Section 2.4.2 (cf. (2.6) and (2.7)). We now derive Π_2^*. Maximizing $\Pi_2(0, T) = (r_1 - w)\lambda_1 - \frac{K}{T} - \frac{\lambda_1 h_r}{2}T$ with respect to T gives $T^* = \sqrt{\frac{2K}{\lambda_1 h_r}}$, so

$$\Pi_2^* = (r_1 - w)\lambda_1 - \sqrt{2h_r\lambda_1}\sqrt{K}. \tag{A-13}$$

We first consider the case where Π_2^* dominates Π_3^* and then the case where the reverse holds.

Subcase (a): Π_2^* dominant over Π_3^*

Here, Π_3^* can be excluded, so we need to compare Π_1^* and Π_2^* to determine the regions of (w, K) in which the retailer prefers each. It is easy to see that the two function are convex decreasing in K for any w and the former is declining at a faster rate as K increases. Thus Π_2^* is dominant if $(r_1 - w)\lambda_1 > (r_2 - w)(\lambda_1 + \lambda_2)$ or equivalently, $w > \frac{r_2(\lambda_1+\lambda_2)-r_1\lambda_1}{\lambda_2} = w_1$. Note that this is the same threshold on w as was defined in Case I. For smaller values of w, however, the dominance regions vis-a-vis Π_1^* and Π_2^* (ignoring Π_3^*) are defined by a monotonic switching curve, i.e., the (w, K) pairs satisfying $\Pi_1^* = \Pi_2^*$, which can be expressed as:

$$w\lambda_2 = r_2(\lambda_1 + \lambda_2) - r_1\lambda_1 - \sqrt{2h_r(\lambda_1 + \lambda_2)}\sqrt{K} + \sqrt{2h_r\lambda_1}\sqrt{K}. \tag{A-14}$$

Π_1^* is dominant to the southwest of this curve and Π_2^* dominates to the northeast.

Subcase (b): Π_3^* dominant over Π_2^*

When $w < w_2$ at the given K, then Π_2^* can be excluded from consideration. The (w, K) pairs satisfying $\Pi_1^* = \Pi_3^*$ can be expressed as

$$w\lambda_2 = r_2(\lambda_1 + \lambda_2) - r_1\lambda_1 - \sqrt{2h_r(\lambda_1 + \lambda_2)}\sqrt{K} + \sqrt{2h_r\lambda_1}\sqrt{K + g(\Delta_0)}. \tag{A-15}$$

Equation (A-15) is defined only for $K' > -g(\Delta_0) > 0$. (In Section 4.2.1, we discussed what happens when $K' + g(\Delta_0) < 0$.) Because $g(\Delta_0)$ is a quadratic, there are two values of w' that satisfy (A-15) for each K' but only the smaller solutionis valid; the larger w' is greater than r_2. Π_1^* is dominant to the southwest of the switching curve and Π_3^* is dominant to the northeast.

When $w = w_2$, $g(\Delta_0) = 0$ so the two switching curves are identical at this point.

Appendix H: Sketches of Solution Procedures for the Manufacturer's Problem

Dominance Map 1, Option 1

Taking the derivative of Π_1^m in (2.7) with respect to K', we obtain

$$\frac{\partial \Pi_1^m}{\partial K'} = \frac{K\sqrt{(\lambda_1 + \lambda_2)h_r}}{2\sqrt{2}} K'^{-1.5} + \left[\frac{\sqrt{(\lambda_1 + \lambda_2)h_r}}{2\sqrt{2}} - \frac{h_m(\lambda_1 + \lambda_2)^{1.5}}{2P\sqrt{2h_r}}\right] K'^{-0.5}. \tag{A-16}$$

The first term is positive, so a sufficient condition for Π_1^m to be increasing is that the expression in parentheses is non-negative, or $h_r/h_m \geq (\lambda_1 + \lambda_2)/P$. Thus, if the retailer's holding cost is high in comparison to that of the manufacturer, the manufacturer will choose K' as the largest feasible value. Otherwise, setting (A-16) to zero yields $K'^* = K(\frac{h_m(\lambda_1+\lambda_2)}{h_r P} - 1)^{-1}$. For any fixed K', the objective is increasing in w', however, so the solution lies on one of the two constraints. Recall that Δ_0 is a linear function of w', so g is a quadratic function of w'. Thus, Lagrangian methods may be required to solve the problem if the IC constraint is binding. On the other hand, if it is not binding, the optimal solution can be obtained in closed form by using the relationship between w' and K' in the participation constraint.

Dominance Map 1, Option 3 The expression for $f(w')$ is:

$$\lambda_2 c_1[1 - h_m\lambda_2 c_1/2P]w'^2 + c_2\lambda_2(1 - h_m\lambda_2 c_1/P)w' - h_m\lambda_2^2 c_2^2/2P,$$

$$\text{where} \quad c_1 = -\frac{\lambda_2/h_2^2}{2(\frac{\lambda_1}{h_1} - \frac{h_r\lambda_1}{2h_1^2})} \quad \text{and} \quad c_2 = \frac{(\lambda_2/h_2^2)(2r_1 - r_2)}{2(\sum \frac{\lambda_i}{h_i} - \frac{h_r\lambda_1}{2h_1^2})} - \frac{r_1 - r_2}{h_2}.$$

Because the objective function, participation constraint and IC constraint all involve $\sqrt{K' + g(\Delta_0')}$, algebraic solution approaches are difficult, but numerical solution methods are relatively straightforward. Despite their complicated form, we have found the functions to be relatively well-behaved, although it is not possible to prove unimodality in general.

Dominance Map 2, Option 1

When $w' > w_2$, the participation and IC constraints have different gradients with respect to K' (i.e., one of the functions is steeper for all K'), so the two functions can cross at most once. The problem can be solved easily by first computing the crossing point, then solving two problems, each with the applicable binding constraint. One also needs to impose constraints on w' to ensure that the solution lies on the correct side of the crossing point and that $w' > w_2$.

When $w' < w_2$, The solution approach is similar to that given in the discussion of Dominance Map 1, except that one additionally needs to consider the upper bound on w.

Dominance Map 2, Option 2

The manufacturer's objective is increasing in w' for fixed K', so for any K', the manufacturer increases w' as much as possible subject to the retailer's participation and IC constraints. The IC constraint is not binding in this case because for any fixed K', there exists a sufficiently large w' that the retailer will prefer Π_2^*. The solution thus lies on the retailer's participation constraint and with some algebra, the solution can be obtained in closed form.

Appendix I: Solution for the Retailer's Problem When $N = 0$

The special case of $N = 0$ is easy to solve and does not require a nested optimization procedure. The retailer's objective is simply to optimize Π^H, and we have $m = \Delta^H/T$. The objective function is

$$\Pi^H = a_0^H(\Delta^H)^3 + b_0^H(\Delta^H)^2 + c_0^H\Delta^H + \gamma_0^H$$

where the coefficients are defined as follows. If $m = \Delta^H/T < R - p$, we have:

$$\begin{aligned}
a_0^H &= (-0.5 + 1/T)\frac{\Gamma}{h\overline{R}}, \\
b_0^H &= [0.5(p - w + \Delta^M) - (\overline{R} - p) - 0.5h_rT^H/T]\frac{\Gamma}{h\overline{R}}, \\
c_0^H &= [0.5h_rT^H(\overline{R} - p) - (p - w + \Delta^M - 0.5h_rT^H)hT^H/T]\frac{\Gamma}{h\overline{R}}, \\
\gamma_0^H &= (p - w + \Delta^M - 0.5h_rT^H)hT^H(\overline{R} - p)\frac{\Gamma}{h\overline{R}} - K.
\end{aligned}$$

On the other hand, if $m = \Delta^H/T \geq (\overline{R} - p)$ we have:

$$\begin{aligned}
a_0^H &= (-0.5 + 1/T - 1/(2T^2))\frac{\Gamma}{h\overline{R}}, \\
b_0^H &= [0.5(p - w + \Delta^M) - (\overline{R} - p) - (p - w + \Delta^M)(1/(2T^2) - 1/T) + (\overline{R} - p)/T]\frac{\Gamma}{h\overline{R}}, \\
c_0^H &= (\overline{R} - p)[(p - w + \Delta^M)(1/1/(2T)) - 0.5(\overline{R} - p)]\frac{\Gamma}{h\overline{R}}, \\
\gamma_0^H &= 0.5(p - w + \Delta^M)(\overline{R} - p)^2\frac{\Gamma}{h\overline{R}} - K.
\end{aligned}$$

We need to solve two problems, one constrained to the region $\Delta^H < (R - p)T$ and the other constrained to the region $\Delta^H \geq R - p$. The solution that returns the higher objective value is the optimal solution. The solution procedures for the constrained problem of $\Delta^H < (R - p)T$ and that of $\Delta^H \geq (R - p)T$ are entirely the same, and the following derivation of the optimal solution applies two both of the two constrained problems.

The solution depends upon the sign of a_0^H. We consider the two cases below. It is easy to show that in both cases, the objective function is a cubic function of Δ^H, but the form of the function for $\Delta^H \geq 0$ depends upon the other coefficients. As such, the optimal solution may be zero, one of the stationary points, or the upper limit on Δ^l (i.e., $(R - p)T$). Below, we describe the form of the cubic function in the region $\Delta^H \geq 0$ for the various combinations of coefficients and the corresponding optimal solutions.

- Subcase (a): $a_0^H < 0$

  If $c_0^H > 0$, then Π^H is increasing and then decreasing in Δ^H, hence $\Delta^{H*} = \min\{(R-p)T, \frac{-b_0^H - \sqrt{b_0^{H\,2} - 3a_0^H c_0^H}}{3a_0^H}\}$

  If $b_0^H < 0$ and $c_0^H < 0$, then Π^H is decreasing in Δ^H, hence $\Delta^{H*} = 0$.

  If $b_0^H > 0$ and $c_0^H < 0$, then Π^H is decreasing then increasing and then decreasing again in Δ^H, hence $\Delta^{H*} = 0$ or $\Delta^{H*} = \min\{(R-p)T, \frac{-b_0^H - \sqrt{b_0^{H\,2} - 3a_0^H c_0^H}}{3a_0^H}\}$.

- Subcase (b): $a_0^H \geq 0$

  If $b_0^H > 0$ and $c_0^H > 0$, then Π^H is increasing in Δ^H, hence $\Delta^{H*} = (R-p)T$.

  If $b_0^H < 0$ and $c_0^H > 0$, then Π^H is increasing then decreasing and then increasing in Δ^H, hence $\Delta^{H*} = \frac{-b_0^H + \sqrt{b_0^{H\,2} + 3a_0^H c_0^H}}{3a_0^H}$ or $\Delta^{H*} = (R-p)T$.

  If $c_0^H < 0$, then Π^H is decreasing then increasing in Δ^H, so $\Delta^{H*} = 0$ or $(R-p)T$. □

Appendix J: Proof of Proposition 2

Let $(\Delta^{H*}, \Delta_1^{L*}, \ldots, \Delta_N^{L*})$ denote the optimal discounts offered by the retailer, and T^{H*} and T^{L*} denote the optimal durations of the high and low cycles, respectively. We assume $T^{H*} \geq T^{L*}$.

We will show that if the discount offered in the low cycle, Δ_i^{L*}, is greater than the discount offered in the high cycle, Δ^{H*}, then the retailer can improve his profit by switching Δ^{H*} and Δ_i^{L*}, and this contradicts the optimality of $(\Delta^{H*}, \Delta_1^{L*}, \ldots, \Delta_N^{L*})$.

As noted earlier, we prove the result by contradiction. Suppose that $\Delta^{H*} < \Delta_i^{L*}$ for some $i \in 1, \ldots, N$. Then

$$\begin{aligned} \Upsilon &\triangleq \Pi(\Delta^{H*}, \Delta_1^{L*}, ..., \Delta_N^{L*}) - \Pi(\Delta_i^{L*}, \Delta_1^{L*}, ..., \Delta^{H*}, ...\Delta_N^{L*}) \\ &= \Pi^H(\Delta^{H*}) + \Pi^L(\Delta_i^{L*}) - \Pi^H(\Delta_i^{L*}) - \Pi^L(\Delta^{H*}). \end{aligned}$$

If $m \leq \overline{R} - p$, then

$$\Upsilon = 0.5\Delta^M((\Delta^{H*})^2 - (\Delta_i^{L*})^2) + 0.5h_r(R-p-m)[T^H(\Delta^{H*} - \Delta_i^{L*}) + T^L(\Delta_i^{L*} - \Delta^{H*})]).$$

and if $m \leq \overline{R} - p$, then

$$\Upsilon = 0.5\Delta^M((\Delta^{H*})^2 - (\Delta_i^{L*})^2) + \Delta^M(\Delta^{H*} - \Delta_i^{L*}).$$

Hence

$$\Pi(\Delta^{H*}, \Delta_1^{L*}, ..., \Delta_N^{L*}) - \Pi(\Delta_i^{L*}, \Delta_1^{L*}, ..., \Delta^{H*}, ...\Delta_N^{L*}) < 0$$

because $\Delta^{H*} < \Delta_i^{L*}$ and $T^{H*} \geq T^{L*}$. This contradicts the optimality of $(\Delta^{H*}, \Delta_1^{L*}, \ldots, \Delta_N^{L*})$. Hence we must have $\Delta^{H*} \geq \Delta_i^{L*}$ for all $i \in \{1, \ldots, N\}$ if $T^{H*} \geq T^{L*}$. □

Appendix K: Optimizing Δ^L for a Given n

In this appendix, we derive the optimal Δs for a fixed n. If there are n low cycles in which the retailer offers a positive discount, then we have $m = \frac{\Delta^H + n\Delta^L}{T}$, where we have omitted the subscript on Δ_i^L because, by Proposition 1, all of the positive Δ^L values are equal, so we need not distinguish them. (Recall, however, that some of the Δ_i^L values are zero.) To convert the retailer's problem to a single-variable problem, we re-express Δ^H as $mT - n\Delta^L$, and the retailer's objective becomes a cubic function of Δ^L. We can infer the functional form from the first derivative of the retailer's objective with respect to Δ^L, which is:

$$\frac{\partial \Pi}{\partial \Delta^L} = a(\Delta^L)^2 + b\Delta^L + c$$

where

$$\begin{aligned} a &= 1.5(n^2-1)\frac{n}{h\overline{R}T}, \text{ and} && \text{(A-17)} \\ b &= [-2(\overline{R}-p-m)^+(n+1) + (p-w)(n+1) + (\Delta^M - 3mT)n]\frac{n}{h\overline{R}T}. && \text{(A-18)} \end{aligned}$$

If $m < \overline{R} - p$, then

$$c = (\overline{R}-p-m)^+[2mT - 0.5h_r(T^H - T^L)] + mT(1.5mT - (p - w + \Delta^M))\frac{n}{h\overline{R}T} \quad \text{(A-19)}$$

and if $m > \overline{R} - p$, then

$$\begin{aligned} c &= (\overline{R}-p-m)^+[2mT - (p-w+\Delta^M) + 0.5(\overline{R}-p-m) + (p-w)/n - 0.5(\overline{R}-p-m)/n] \\ &+ mT(1.5mT - (p-w+\Delta^M))\frac{n}{h\overline{R}T}. && \text{(A-20)} \end{aligned}$$

For a fixed m, the problem is to allocate a total discount, m, among the high cycle and n low cycles, where the discounts in the low cycle are equal to one another. The optimal solution depends upon the signs of b and c, which define which of the four possible forms of Π (as shown in Figure 4.2) applies. We assume $T^H \geq T^L$ in this section, which is commonly true in practice because retailers tend to stockpile hence high cycles tend to be longer than low cycles. We allow T^H and T^L to be in any order in the Numerical Study in Section 3.5.

- $b > 0, c > 0$

  If $b > 0$ and $c > 0$, Π is convex increasing in Δ^L. The solution is $\Delta^H = 0$ and $\Delta^L = mT$. We conclude that this solution is suboptimal based on the results in Proposition 2.

- $b < 0, c > 0$

  As shown in the upper right plot in Figure 4.2, the two local optima are (i) the stationary point $\Delta^L = \frac{-b+\sqrt{b^2-4ac}}{2a}$, if it exists, and (ii) the upper limit of Δ^L (i.e., $\Delta^L = \frac{mT}{n}, \Delta^H = 0$). The upper limit is suboptimal based on Proposition 2. If $b^2 - 4ac \geq 0$, the stationary point exists and $\Delta^{L*} = \frac{-b+\sqrt{b^2-4ac}}{2a}$; if $b^2 - 4ac < 0$, the objective is increasing in Δ^L so the solution $\Delta^L = mT/n$ is suboptimal .

- $-\infty < b < \infty$, $c \leq 0$

  When $c \leq 0$, the only difference between the cases with $b \leq 0$ and $b > 0$ is that when $b > 0$, Π^L is convex for small Δ^L. But in both cases, the retailer's profit function is decreasing then increasing and hence the optimal solution is either at the lower limit (i.e., $\Delta^L = 0$), or the upper limit ($\Delta^L = mT/n$). However, by Proposition 2, the latter solution is suboptimal so the solution is $(\Delta^{H*} = mT, \Delta^{L*} = 0)$ in this case.

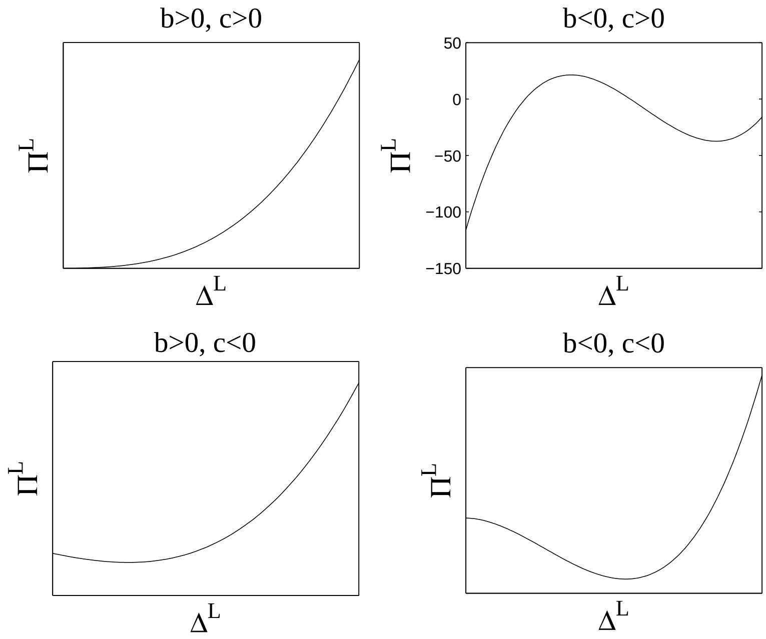

Figure 4.2: Π^L For Different Signs of b, c

Therefore, the optimal solution to the inner problem is either the boundary solution ($\Delta^{H*} = mT, \Delta^{L*} = 0$) or the stationary point ($\Delta^{L*} = \frac{-b+\sqrt{b^2-4ac}}{2a}, \Delta^{H*} = mT - n\Delta^{L*}$). We solve the outer problem utilizing these two solutions to the inner problem.

Appendix L: Outer Optimization Problem: Optimizing m

We derive the solution for m^* for the two potentially optimal solutions: (1) $\Delta^{H*} = mT$ and $\Delta^{T*} = 0$ and (2) $\Delta^{L*} = \frac{-b+\sqrt{b^2-4ac}}{2a}$ where a, b and c are defined in (A-17) through (A-20), and $\Delta^{H*} = mT - n\Delta^{L*}$.

Potential Solution 1: $\Delta^H = mT, \Delta^L = 0$ with $\alpha_1 < 0$

Substituting for Δ^H and Δ^L in the objective function, the retailer's objective can be rewritten as $\Pi(m) = \alpha_1 m^3 + \alpha_2 m^2 + \alpha_3 m + \alpha_4$ where
(Case 1:) if $m \leq \overline{R} - p$, then

$$
\begin{aligned}
\alpha_1 &= 0.5T^2(2-T)\frac{\Gamma}{h\overline{R}}, \\
\alpha_2 &= [(0.5(p - w + \Delta^M) - (\overline{R} - p))T^2 - 0.5h_r T^H T]\frac{\Gamma}{h\overline{R}}, \\
\alpha_3 &= [0.5h_r(\overline{R} - p)T^H T - (p - w + \Delta^M - 0.5h_r T^H)hT^H - n(p - w - 0.5h_r T^L)hT^L]\frac{\Gamma}{h\overline{R}}, \\
\alpha_4 &= (\overline{R} - p)[(p - w + \Delta^M - 0.5h_r T^H)hT^H + n(p - w - 0.5h_r T^L)hT^L]\frac{\Gamma}{h\overline{R}} - (n+1)K
\end{aligned}
$$

(Case 2:) if $m > R - p$, then

$$
\begin{aligned}
\alpha_1 &= -0.5T(T-1)^2\frac{\Gamma}{h\overline{R}}, \\
\alpha_2 &= [(0.5(p-w+\Delta^M) - (R-p))T^2 + ((\overline{R}-p) - (p-w+\Delta^M))T \\
&\quad +0.5((p-w)(n+1)+\Delta^M)]\frac{\Gamma}{h\overline{R}}, \\
\alpha_3 &= [((\overline{R}-p)(p-w+\Delta^M) - 0.5(\overline{R}-p)^2)T - ((p-w)(n+1)+\Delta^M)(\overline{R}-p)]\frac{\Gamma}{h\overline{R}}, \\
\alpha_4 &= 0.5(\overline{R}-p)^2[(p-w)(n+1)+\Delta^M)]\frac{\Gamma}{h\overline{R}} - (n+1)K
\end{aligned}
$$

If $m \leq \overline{R} - p$, then α_1 is negative for $T > 2$, and if $m > \overline{R} - p$, then α_1 is negative for $T > 1$. Note that if $T < 2$, then there cannot be a low cycle and the special case with no low cycles applies. (See Appendix I for an analysis of this special case.) Furthermore, if $T = 2$, then the high and low cycles are forced to have the same duration, which is quite restrictive. In the analysis that follows, we initially focus on cases with $T > 2$, and then return to the case of $T = 2$. The analysis depends upon the signs of α_2 and α_3, and the corresponding shape of the retailer's objective is shown in Figure 3.2; we divide our analysis accordingly. In all cases, the same analysis applies to both $m \leq \overline{R} - p$ and $m > \overline{R} - p$.

(a) $\alpha_2 \leq 0, \alpha_3 \leq 0$

$\Pi(m)$ is decreasing in m for $m \geq 0$, and hence $m^* = 0$.

(b) $\alpha_2 > 0$, $\alpha_3 < 0$

$\Pi(m)$ is decreasing then increasing and then decreasing again as m increases from 0. There are two local optima: $m^* = 0$ and the larger of the two stationary points, $m^* = \frac{\alpha_2 + \sqrt{\alpha_2{}^2 - 3\alpha_1\alpha_3}}{-3\alpha_1}$. This stationary point exists if and only if $\alpha_2^2 + 2T^2(T-2)\alpha_3 \geq 0$. If $\alpha_2^2 + 2T^2(T-2)\alpha_3 < 0$, then $\Pi(m)$ is decreasing in m, so $m^* = 0$ is the optimal solution. Therefore, the optimal solution is $m^* = \min\{\overline{R} - p, \frac{\alpha_2 + \sqrt{\alpha_2{}^2 - 3\alpha_1\alpha_3}}{-3\alpha_1}\}$ if the stationary solution exists and $m^* = 0$ otherwise.

(c) $-\infty < \alpha_2 < \infty$ and $\alpha_3 > 0$

$\Pi(m)$ is increasing and then decreasing in m for $m \geq 0$, so the unconstrained optimal solution is at the stationary point, $\hat{m} = \frac{\alpha_2 + \sqrt{\alpha_2{}^2 - 3\alpha_1\alpha_3}}{-3\alpha_1}$. (Note that the stationary point exists because the expression under the square root is positive.) Therefore, in Case 1, $m^* = \min\{\overline{R} - p, \hat{m}\}$ and in Case 2, $m^* = \hat{m}$.

We now provide further economic interpretation of the coefficients in the objective function. Assuming that $\alpha_1 < 0$, the other key parameters that affect characteristics of the solution are α_2 and α_3. If $\alpha_2 > 0$ at a given value of Δ^M, then, roughly speaking, the retailer gains more from selling to discount customers than he loses due to the reduced margin on units sold to regular customers If $\alpha_3 > 0$ at a given value of Δ^M, then the retailer's savings in inventory costs from discounting is greater than the reduction in profit due to the fact that the (negative) brand equity effect reduces the number of regular customers. If $\alpha_2 > 0$and $\alpha_3 < 0$ (or the reverse), then the retailer faces competing forces when optimizing his discounts. If both α_2 and α_3 are negative, then the retailer has little incentive to offer discounts at the given value of Δ^M.

Potential Solution 2: $\Delta^{L^*} = \frac{-b-\sqrt{b^2-4ac}}{2a}, \Delta^H = mT - n\Delta^{L^*}$

Substituting for Δ^{L^*} and Δ^H in $\Pi(\Delta^L)$ gives the objective expressed as a function of m: $\Pi = \frac{a}{3}\Delta^{L^{*3}} + \frac{b}{2}\Delta^{L^{*2}} + c\Delta^{L^*} + \theta$, where $\Delta^{L^*} = \frac{-b-\sqrt{b^2-4ac}}{2a}$, $\theta = a^H(mT)^3 + b^H(mT)^2 +$

$c^H(mT) + \gamma^H + nr^L$, a, b, c are functions of m and are defined in (A-17), (A-18), (A-19), and (A-20) $a^H, b^H, c^H, \gamma^H, \gamma^L$ are also functions of m and defined in Section 3.4.

However, the optimal m cannot be expressed analytically. We provide optimality conditions below; numerical methods can be used to find optimal m.

The first order condition $\frac{\partial \Pi(\Delta^L)}{\partial \Delta^L} = 0$ can be written as

$$\begin{aligned} & a\Delta^L(m)\frac{\partial \Delta^L(m)}{\partial m} + 0.5\Delta^L(m)^2[2(n+1) - 3nT] + b\Delta^L \frac{\partial \Delta^L(m)}{\partial m} \\ & +[-2mT - 0.5h_r(T^H - T^L) + 2T(R-p-m)]\Delta^L(m) + c\frac{\partial \Delta^L(m)}{\partial m} = 0 \qquad \text{(A-21)} \end{aligned}$$

where

$$\begin{aligned} \Delta^L(m) &= \frac{-b - \sqrt{b^2 - 4ac}}{2a} \\ \frac{\partial \Delta^L(m)}{\partial m} &= -\frac{1}{2a}\{2(n+1) - 3nT - 2b[2(n+1) - 3nT] - 4a\{-2mT - 0.5(T^H - T^L) \\ &+ \ 2T(R-p-m) + T[1.5mT - (p - w + \Delta^M) + 1.5mT^2]\}/[2\sqrt{b^2 - 4ac}]\} \end{aligned}$$

The optimal value of m is either at the stationary point mentioned above (i.e., the solution to (A-21)) or at the boundary value, 0.

Finally, we turn to the special case of $\alpha_1 = 0$ (which implies that $T = 2$) under the assumption that $m \leq \overline{R} - p$. If $\alpha_1 = 0$ and $m < \overline{R} - p$, the solution procedure for $\alpha_1 < 0$ can be used. The retailer's objective function is $\Pi(m) = \alpha_1 m^3 + \alpha_2 m^2 + \alpha_3 m + \alpha_4$ and expressions for the αs are those listed under Potential Solution 1 at the beginning of this appendix. Because $\alpha_1 = 0$, Π is quadratic in this case, so the solution depends on the signs of α_2 and α_3 and is either at a stationary point or at the boundary of 0. Details follow.

- $\alpha_2 > 0, \alpha_3 > 0$:
  Π is concave increasing for $m \geq 0$, and the optimal solution is $m^* = \overline{R} - p$.
- $\alpha_2 > 0, \alpha_3 < 0$:
  Π is concave decreasing and then increasing, hence m^* is either 0 or $\overline{R} - p$.
- $\alpha_2 < 0, \alpha_3 > 0$:
  Π is concave increasing and then decreasing, and m^* is at the stationary point: $\frac{-\alpha_3 - \sqrt{\alpha_3{}^2 - 4\alpha_2\alpha_3}}{2\alpha_2}$.
- $\alpha_2 < 0, \alpha_3 < 0$:
  Π is concave decreasing for $m \geq 0$, hence $m^* = 0$.

The economic interpretations of α_2 and α_3 are the same as those given earlier, so if both are positive, large discounts are attractive and if both are negative, discounts are undesirable for the retailer. When these coefficients have different signs, there may be a tradeoff, in which case an intermediate value may be optimal.