Essays on Supply Chain Contracting and Retail Pricing

by

Thunyarat Amornpetchkul

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in The University of Michigan
2014

Doctoral Committee:

Associate Professor Hyun-Soo Ahn, Chair
Professor Izak Duenyas
Assistant Professor Ozge Sahin
Associate Professor Mark P. Van Oyen
Assistant Professor Xun Wu

To my parents and all of my teachers

ACKNOWLEDGEMENTS

To me, this dissertation is not only part of the requirements for graduation, but it is also a keepsake that will always remind me of all the things I learned, all the work I did, and all the support I received during my years as a PhD student.

My special thanks and appreciation go to the two advisors, Professor Hyun-Soo Ahn and Professor Ozge Sahin. Our student-advisors relationship may have started from mutual research interests. But as it developed, I think the three of us formed a unique combination of researchers who can be both productive and fun. The main products coming out of our research meetings are research papers, included in this dissertation. But if we were to record all of the other discussions that have taken place during those meetings, I think we would have had a byproduct as a fashion or lifestyle book. Thank you for your guidance, which helped me start my PhD journey. Thank you for your kindness and patience, which gave me enough room to learn along the way. Thank you for your trust and encouragement, which pushed me to go further. Thank you for everything you did for me, which landed me at the finish line.

I also would like to thank the other committee members: Professor Izak Duenyas, Professor Xun (Brian) Wu, and Professor Mark Van Oyen. I truly appreciate their time and valuable suggestions which helped improve various aspects of this dissertation. I am especially grateful to Professor Izak Duenyas for his contribution and advice on the first essay, as well as his support as the department chair.

In addition to the dissertation committee, I would like to extend my gratitude towards the faculty of the Department of Operations and Management Science (now

Technology and Operations) as a whole. Thank you for believing in me that I could meet the department's high expectations, and for giving me all the support I needed to thrive here. In particular, I would like to thank the PhD program coordinators during the past five years: Professor Roman Kapuscinski, Professor Hyun-Soo Ahn, Professor Damian Beil, and Professor Amitabh Sinha, for helping guide me through different phases in the program. Outside of the department, I also received wonderful support on administrative matters from the Doctoral Studies Program staff: Brian Jones, Roberta Perry, and Kelsey Zill. They have always been very kind and responsive in addressing my questions and needs.

Another group of people who made my PhD experience unexpectedly enjoyable is the "OMS family," as we PhD students like to call ourselves. They are indeed like a second family to me. Although the members have changed over time, as we had graduated and newly admitted students, one thing that never changes is that we always hope for the best for one another. One person to whom I would like to express my deep gratitude for our friendship is Anyan Qi, who joined the PhD program in the same year as me. His outstanding intelligence helped me learn and inspired me to set higher goals for myself. His affability and optimism created our pleasant work environment and kept me sane during difficult times. I cannot imagine having a better person to go through these challenging years with – thank you, Anyan.

Last but not least, I am always thankful to my family and loved ones. Their love and heartfelt encouragement helped build the confidence that carried me through the PhD journey. My achievement today owes much to their continuous support that I will never forget.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

ABSTRACT

Essays on Supply Chain Contracting and Retail Pricing

by

Thunyarat (Bam) Amornpetchkul

An important operational decision that a seller has to make is how to price his product under different situations. This dissertation addresses three unique pricing problems, commonly faced by a seller in a supply chain, in a series of three essays.

The first essay considers a supplier's problem of choosing which type of contracts to offer to a retailer whose demand forecasts can be improved over time. It is shown that there exist mechanisms which enable the supplier to always benefit from the retailer's improved demand forecasts. Such a mechanism consists of an initial contract, offered to the retailer before she obtains improved forecasts, and a later contract (contingent on the initial contract), offered to the retailer after she obtains improved forecasts.

The second essay investigates a retailer's problem of choosing which form of price promotions to offer to consumers, some of which are more inclined to increase spending when satisfied with the value of the deals. Two types of promotions are considered: i) all-unit discount, where a price reduction applies to every unit of a purchase that meets the minimum requirement, and ii) fixed-amount discount, where the final amount that a consumer has to pay is reduced by a predetermined discount amount if the purchase meets the minimum requirement. It is shown that both discount schemes

can induce consumers to overspend. However, depending on consumer valuation of the product, one scheme can be more profitable to the retailer than the other.

The third essay discusses a dual-channel retailer's problem of choosing a price differentiating policy (charging different prices for the same product sold at different channels) and/or inventory transshipping policy (transferring inventory between the channels) to balance available inventory and demand arriving at each channel. It is shown that the two mechanisms have different implications on sales volume. Which mechanism is more effective depends on the retailer's initial inventory position. Furthermore, when implemented concurrently, the benefit from price differentiation and inventory transshipment mechanisms may either substitute or complement each other.

CHAPTER 1

Introduction

A fundamental question for any sellers in a supply chain is what pricing mechanism to use to generate most profits from selling their products. The answer to this question heavily depends on the nature of the businesses as well as the characteristics of the buyers. For example, a supplier selling to a retailer who has superior information about the end demand would benefit from a mechanism that promotes demand information sharing. A retailer selling to customers who enjoy receiving discounts would find it profitable to offer a price promotion that induces larger purchases. For a retailer who operates in more than one channel, it is important to use a pricing mechanism that helps balance available inventory and demand at each channel in order to maximize the overall profit.

This dissertation explores seller's problems across two different areas of a supply chain: upstream (a supplier selling to a retailer) and downstream (a retailer selling to customers). More precisely, the dissertation consists of three essays; one on Supply Chain Contracting, and the other two on Retail Pricing. Each essay investigates operational problems arising from interactions between the respective supply chain parties as a seller or a buyer. Despite different focuses, all essays consider realistic business situations where the seller and the buyer make decisions based on their own benefits, and the buyer's behavior may be influenced by her perspectives towards the pricing mechanism offered by the seller.

The first essay titled "Mechanisms to Induce Buyer Forecasting: Do Suppliers Always Benefit from Better Forecasting?" explores the effects of improved demand information on the supplier's and the retailer's profitability under different types of supply chain contracts. More specifically, three types of contracts that a supplier (seller) can offer to a retailer (buyer) are considered: 1) a contract that is offered before the buyer can obtain improved forecasts, 2) a contract that is offered after the buyer has obtained improved forecasts, and 3) a contingent (dynamic) contract where an initial contract is offered to the buyer before she obtains improved forecasts, followed by a later contract (contingent on the initial contract) offered after improved forecasts have been obtained. In a scenario where the supplier is certain that the buyer can obtain more accurate forecasts over time, the contingent contract is shown to be the most profitable mechanism for the supplier. The contingent contract also guarantees the supplier an increasingly larger profit as the buyer's forecast accuracy increases. In a different scenario where the supplier is uncertain whether the buyer can improve forecasts over time, the essay discusses how the supplier can modify the contingent contract to screen the buyer on both her demand and forecasting capability information. Under such a contract, the supplier's profit increases with the probability that the buyer is capable of improving forecast accuracy. In contrast to the existing literature, the results from this essay show that there exist mechanisms which enable the supplier to always benefit from better demand information.

The second essay, "Conditional Promotions and Consumer Overspending," discusses the implications of sales promotions on consumer spending. In particular, when a deal comes with an eligibility requirement in the form of a minimum purchase quantity or a minimum spending, it may lead some consumers to end up buying more than what they need just to qualify for the discount offer. This essay investigates the effects of conditional promotions (e.g., buy 2 or more get 30% off, spend \$50 or more get \$15 off) on consumer purchase decisions and the retailer's profitability. Two popular types of conditional promotions are considered: i) all-unit discount, where a price reduction applies to every unit of a purchase that meets the minimum requirement, and ii) fixed-amount discount, where the final amount that a consumer has to pay is reduced by a predetermined discount amount if the consumer's purchase meets the minimum requirement. The results from this essay show that both discount schemes can induce consumers to overspend. However, consumer overspending benefits the retailer only when there is a sufficiently large proportion of highly deal-prone or high-valuation consumers in the market. Additionally, depending on the nature of products, one discount scheme can be more profitable to the retailer than the other. The all-unit discount outperforms the fixed-amount discount when consumers are not willing to pay the regular price for the product; while, the fixed-amount discount is more profitable than the all-unit discount when there exist consumers who would make a purchase even without a discount. These findings suggest that adopting an appropriate type of conditional discounts can effectively improve the retailer's profit over what obtained through selling at the regular price or a conventional price markdown.

The third essay, "Dynamic Pricing or Dynamic Logistics?" aims to understand how the pricing mechanism and inventory transshipping mechanism can help improve the retailer's profit in a dual-channel environment. This study considers a dynamic pricing problem of a retailer who sells a product through two channels (e.g., online and physical store), where inventory is kept at two separate locations, dedicated for demand arriving at each channel. To balance inventory and demand at each channel, the retailer may employ a price differentiation policy and/or an inventory transshipment policy. A price differentiation policy helps manage demand by allowing the retailer to charge different prices for the same product sold at different channels in each period. On the other hand, an inventory transshipment policy acts on the inventory side by allowing the retailer to transfer inventory between the channels when needed. This essay characterizes the retailer's optimal pricing and transshipping policy, and compares the effectiveness of the two mechanisms in improving profits. The findings show that the optimal price differentiation policy in the current period always results in a larger expected sales volume, compared to the optimal uniform pricing policy. On the other hand, the optimal transshipment decision may result in a larger or smaller expected sales. While price differentiation provides a larger profit improvement than transshipment does in many situations, transshipment is shown more effective when the retailer holds significantly less inventory at the high-margin channel. Furthermore, when implemented concurrently, the benefit from price differentiation and inventory transshipment mechanisms may either substitute or complement each other. The two mechanisms can substitute each other when the retailer's objective is to correct his inventory position. However, when the retailer prefers to maintain the same balance of inventory at the channels, the two mechanisms work together, complementarily.

The rest of this dissertation is organized as follows. Chapter 2, 3, and 4 discuss the first, second, and third essay, respectively. An overall conclusion of the dissertation is provided in Chapter 5.

CHAPTER 2

Mechanisms to Induce Buyer Forecasting: Do Suppliers Always Benefit from Better Forecasting?

2.1 Introduction

In this chapter, we consider a supplier selling goods to a buyer under information asymmetry and multiple forecast updates before the selling season. We assume that the buyer, due to her proximity to the markets in which she is selling, may have more information about demand than the supplier. Furthermore, as the selling period approaches, the buyer may have the capability to obtain even better (more accurate) forecasts. We focus on investigating when the buyer would have the incentive to obtain better forecasts, and what kinds of contract offerings would allow the supplier to benefit from the better information obtained by the buyer over the procurement season. We are interested in how temporal changes in forecast accuracy affect whether the supplier benefits from the buyer obtaining improved forecasts. Previous literature has obtained contradictory results, showing that it is possible for the supplier's profits to decrease when buyers obtain improved demand forecasts. We note however that these results were obtained under the assumption that the supplier and the buyer utilize static contracts, where contract ordering takes place only once. In this essay, we consider another type of contract which allows multiple ordering opportunities, and show that such mechanism can guarantee the supplier's benefit from the buyer's improved demand information. More specifically, three unique contributions of this essay are: 1) we consider dynamic (contingent) contracts and show how they can be utilized in conjunction with forecast updates in favor of the supplier. We show that if dynamic contracts are used effectively, then the supplier can in fact always benefit from temporal improvements in the buyer's forecast accuracy (in contrast to the static case) so long as the buyer is capable of obtaining forecast updates. We also show how dynamic contracts can be easily adapted to benefit the supplier even when the buyer may refuse to obtain forecast updates. 2) We derive results that are robust under many possible business situations (e.g., endogenous/exogenous retail price with/without salvage values). And, 3) we provide analytical results regarding the effects of the supplier's uncertainty about the buyer's forecasting capability on the supplier's and the buyer's profit. In particular, we show that even in presence of such uncertainty, the supplier can design a sophisticated screening contract which allows him to benefit from more accurate demand information.

The value of a buyer's demand forecast on supply chain profits has drawn a lot of attention recently. It is intuitive to expect that both supplier and buyer benefit from better demand information. However, under information asymmetry, and certain type of contract structures, it may not be true that both parties benefit from improved demand information. For example, *Taylor* (2006) showed that the supplier may prefer to contract with the buyer before more accurate demand information is received. Most of the other OM papers on this topic to date have focused on static contracts and single forecast update scenarios. However, in this essay, we model an evolving information asymmetry between a buyer and a supplier due to a second forecast update by the buyer and introduce dynamic or contingent contracts. We show that if the supplier has enough power to offer take-it-or-leave-it contingent contracts, and if the buyer has capability to obtain better forecasts, then contingent contracts would always result in higher profits for the supplier than static contracts. Furthermore, utilizing dynamic contracts, the supplier can always take advantage of the buyer's improved demand forecasting to increase his profits. We consider a simple two-period model similar to those considered in other papers (e.g. *Taylor* 2006). We assume that in period 1, the buyer and the supplier have some priors on demand. We capture the initial information asymmetry between the two parties by assuming that the buyer may have more detailed prior information due to her proximity to the market, previous selling experience, etc. Furthermore, the buyer may or may not have the capability to obtain a better second forecast of demand in period 2. The supplier can produce in both periods, but faces a higher production cost if producing in period 2 (This reflects the higher capacity cost due to expedited production or transportation costs.). In such situations, most of the contracts that have been considered in the literature are either "early contracting," where the buyer and supplier sign a contract in period 1, or "late contracting," where the contract takes place only after the buyer has obtained the more refined forecast. If we limit ourselves to only these kinds of contracts, then consistent with previous literature, there exist situations where both parties prefer to contract with less accurate demand information. However, we show that the supplier can offer a contingent contract, where he offers the buyer a menu of choices in period 1, and also a menu of choices in period 2 (which is a function of what was chosen in period 1). In this case, we show that this contract always provides the supplier with higher profits than either type of static contracts; hence, the supplier always benefits from the forecast refinement. Although the contingent contract is not always the most profitable for the buyer, there exist situations where the buyer also prefers it and the contingent contract is a win-win solution for the supply chain.

As a simple example that describes the setting of this chapter, consider the famous Sport Obermeyer Ltd case (*Hammond and Raman*, 1994) taught in most MBA programs. In the case, Sport Obermeyer first has an initial forecast, then has most of its demand uncertainty resolved at the Las Vegas trade show where it displays its ski jackets for the new season and receives orders. However, to obtain better forecasts, Sport Obermeyer institutes an early-write program where it invites some of its largest and most representative buyers to an all-expenses paid ski vacation in Aspen a few months before the Las Vegas trade show and gauges the buyer's reaction to the products, receives some early orders, and uses the reactions and the early orders to update its forecasts for the different ski jacket models. A key take-away of this case study, as it is taught in many business schools, is to show the importance of obtaining better demand forecasts before the final demand is revealed. Realizing the importance of more accurate demand information, many manufacturers and retailers update their demand forecasts multiple times in a procurement season as in the Sport Obermeyer case. However, today many companies selling goods in the U.S. use fairly large contract manufacturers or supply chain integrators in Asia to get their products manufactured. Increasingly, these suppliers have become much larger and more powerful in their respective supply chains. Therefore, in certain product categories, especially if the product requires advanced know-how, it is very difficult for a small manufacturer to produce its products without using one of these large contract manufacturers. As these contract manufacturers become larger and more powerful, they are able to offer take-it-or-leave-it contracts to relatively smaller buyers. In an article on aligning incentives in supply chains, *Narayanan and Raman* (2004) write "Companies should explore contract-based solutions before they turn to other approaches, because contracts are quick and easy to implement." As the contract manufacturer increasingly gets more power to set contractual terms, a reasonable question to ask is whether a buyer would be willing to obtain better forecasts and share these with the contract manufacturer. Consider a small start-up high tech company who would probably have to contract with much more powerful contract manufacturers or supply chain integrators or a small start-up apparel manufacturer who would have to contract with Li&Fung to get its products manufactured. Is it still true that obtaining more detailed forecasts will benefit such a manufacturer facing a much more powerful supplier as was the case 20 years ago?

A novel aspect of our research is that we also consider the situation where the buyer's capability to obtain more detailed forecasts may be unknown to the supplier. Thus, our analysis is divided into two cases: 1) where all parties know that the buyer is capable of obtaining more accurate forecasts, and 2) where the supplier is uncertain of the buyer's capability. Even a very powerful supplier that can offer take-it-or-leave-it contracts may not be able to force all buyers to obtain more accurate forecasts. For example, a buyer may claim that her staff does not have the technical sophistication, the resources, or the market leads necessary to obtain more accurate forecasts than what is available in period 1. If the supplier knows that the buyer in fact does have such capabilities, then any refusal to obtain more accurate forecasts will lead the supplier to update his beliefs about the demand that the buyer is facing. However, the supplier may be truly uncertain about the buyer's forecasting capabilities. For example, even though Wal-Mart is very well regarded for its precision in matching supply to demand, it struggled in estimating demand when entering the markets in China, Brazil, and Indonesia. When even Wal-Mart struggles in forecasting in these countries, a supplier facing a buyer that claims obtaining better forecasts is not possible may have some uncertainty about the buyer's forecasting capability. Therefore, it is interesting to explore how such a supplier can offer contracts to a buyer by screening them both for forecasting capability as well as demand type.

Our study aims to answer the following research questions:

1. Which type of contracts is most profitable for the buyer and supplier?

2. How does the buyer decide (if she is capable) whether or not to obtain more accurate forecasts? How do the types of contracts offered by the supplier affect this decision?

3. How does the supplier's knowledge of whether the buyer is capable of obtaining better forecasts affect the kind of contracts he offers to the buyer?

These questions differentiate our work from most of the supply chain coordination literature in that our emphasis is not on coordinating contracts, but rather, which contract is most profitable to which party, and whether (and when) multiple forecasts benefit the buyer or supplier. We note that the answer to question 2, which asks if a buyer would ever suffer (or benefit) from a more accurate forecast, also depends greatly on the supplier's knowledge of the buyer's capability. If the supplier knows that the buyer is capable of obtaining more accurate forecasts, an announcement that the buyer *chooses* not to obtain forecasts can lead the supplier to update his beliefs about the buyer's demand expectation. We take this into account and address whether a buyer can ever decide not to obtain forecasts (because obtaining forecasts can result in profit reduction) so long as the supplier knows the buyer has the capability to obtain forecasts. Additionally, since the supplier may not be certain whether the buyer indeed has the capability to obtain more detailed forecasts, we also address how the supplier should revise his contract offerings taking into account his priors on the buyer's forecast capability. Thus, our main research focus is not only to see whether the supplier and the buyer can benefit from contracting dynamically, but also (and more importantly) to determine "when" or under "which circumstances" the dynamic contract is implementable (both parties agree to contract), and when it is not. This is why we analyze the buyer's preferences for contracts which leads to the question of whether the buyer can refuse to obtain more accurate forecasts. This in turn leads us to analyze how the supplier would interpret this refusal when he is sure the supplier is capable of obtaining forecast updates and when he is not.

The rest of the chapter is organized as follows. In Section 2, we review the literature on contracting with information asymmetry and forecast updating. Section 3 introduces the model framework, and discusses the three contract choices we analyze. In Section 4, we study which of the three contract types (early static, late static, or dynamic) the buyer and supplier prefer. We also address the question of whether a buyer can refuse to obtain better forecasts if this refusal has signal value to the supplier in Section 5. In Section 6, we address the case where the supplier is uncertain about the buyer's accurate forecast capability (or cost) and show how the supplier can write a two-dimensional screening contract (on buyer's second forecast capability and demand type) to screen the buyer. We conclude with discussion and future research directions.

2.2 Literature Review

In this essay, we study the nonlinear optimal static and contingent contracts that can be signed before or after the buyer obtains more accurate demand forecast when the information is asymmetric in the supply chain. We review three areas of research that are related to the present work. Methodologically, this essay draws results from Incentive Theory, a branch of Economics studying strategic interaction between two parties under asymmetric information. Incentive Theory deals with both static and dynamic screening problems. Its focus has mainly been on deriving the optimal static screening contract for a principal who wants to optimally elicit information from a privately informed agent, also known as an adverse selection problem. For more information on static adverse selection problems see *Laffont and Martimort* (2002). Multi-period models with dynamic information structures are less understood. *Fudenberg et al.* (1990) is one of the first papers to study a dynamic principal-agent model with an underlying stochastic process. *Bolton and Dewatripont* (2005) provides a good summary of the literature on dynamic principal agent models. In this essay, we consider both static and dynamic (contingent) contracts in a single procurement season. There are a number of papers in the operations management literature that study the dynamic procurement contracts in a principal agent framework. *Plam-*

beck and Zenios (2000) and *Zhang and Zenios* (2008) study dynamic principal agent models and show that the models can be solved using dynamic programming. *Lobel and Xiao* (2013) study the manufacturer's problem of designing a long-term dynamic supply contract, and show that the optimal contract takes a simple form: a menu of wholesale prices and associated upfront payments. While these papers assume that the principal and the agent contract repeatedly over multiple procurement seasons, we assume that they contract only once but the contract terms may require repeated (dynamic) interaction in a single procurement season. We are interested in modeling the multiple forecast updates in a procurement season and identify situations where the dynamic contracts are implementable.

The second related area is on the effect of the accuracy of the demand forecasts on supply chain, supplier, and buyer profits. The issue of buyer's demand forecast accuracy on supply chain profits has drawn increasing attention. It is natural to think that both supplier and buyer benefit from better forecasts. However, recently, *Taylor* (2006), *Taylor and Xiao* (2010), and *Miyaoka and Hausman* (2008) show that more accurate or precise forecasts are not always profitable to the supplier and the retailer. *Taylor* (2006) examines the impact of information asymmetry, forecast accuracy, and retailer sales effort on the manufacturer's sale timing decision. He characterizes the sales timing preference as a function of the production cost. *Miyaoka and Hausman* (2008) consider the effects of having the wholesale price determined by different parties and at different times. They present scenarios where the supplier and the retailer are hurt or rewarded by the improved forecasts. One fundamental difference between the present work and the earlier literature is that we investigate when it benefits the supplier for the buyer to obtain multiple forecast updates in a procurement season; while, the existing literature mostly focuses on the refinement of a single demand forecast, and whether increased accuracy of this one demand forecast benefits the supplier or the supply chain under almost exclusively static contracts.

Additionally, we investigate the contract structures that promote or inhibit such forecast updates such as dynamic (contingent) contracts that allows the supplier to screen the buyer multiple times as she updates her forecast. This allows us to provide managerial insights, which are different from what have been shown in the literature, that temporal increases in forecast accuracy in fact can always benefit the supplier if an appropriate mechanism is utilized.

Others who examine different aspects of information asymmetry and forecast sharing in supply chains are *Cachon and Lariviere* (2001), *Özer and Wei* (2006), and *Taylor and Xiao* (2009). *Cachon and Lariviere* (2001) focus on information asymmetry and study forecast sharing between a manufacturer and a supplier. In their model, the retailer offers the contract and channel coordination is achievable only if she dictates the capacity decision. Similarly, *Özer and Wei* (2006) study forecast sharing but assume that the supplier offers the contract. They consider capacity reservation and advance purchase contracts to assure credible forecast sharing. *Taylor and Xiao* (2009) study incentives to induce buyer forecasting with rebates and returns contracts if the forecast update is costly. They design contracts that induce the buyer to forecast and compare these with the contracts that do not induce forecasting. These papers assume a single forecast update and no uncertainty on the buyer's forecasting capability. Another relevant work to ours is *Lariviere* (2002). He considers a supplier selling to a retailer who may be capable (incur a cheap forecasting cost) or incapable (incur an expensive forecasting cost) of forecasting demand, similar to our model in Section 6. To induce the capable retailer to forecast and share improved demand information, the supplier employs either price-based returns mechanisms (buy backs) or quantity-based returns mechanisms (quantity flexibility contracts). His paper considers a single-period and single-forecast model, and focuses on comparing the performance of the two restricted return mechanisms mostly relying on a numerical study. On the other hand, we focus on investigating the effects of uncertainty in the buyer's forecasting capability and the buyer's forecast accuracy on the supplier's and the buyer's profit using a general non-linear contract. Solving a two-dimensional screening problem, we analytically show that the supplier benefits from the increased forecast accuracy and increased probability of facing a capable buyer while the buyer's profit decreases as the supplier's prior about the capability probability increases. Interestingly, when the capability of the buyer is uncertain and the supplier screens both dimensions, as the forecast accuracy in period 2 improves, buyer's profit stays constant. For a general multidimensional screening problem, see *Rochet and Choné* (1998). While the contract constraints in their multidimensional screening problem are similar to what we consider in Section 6, they only consider a single-period problem and their model does not involve demand forecasts.

The third related area is the optimal contract structure and timing of orders when the demand information evolves over time. *Ferguson* (2003), and *Ferguson et al.* (2005), study a buyer that produces and assembles components using parts procured from the supplier. Similar to our model *Ferguson et al.* (2005) assumes that the demand uncertainty is partially resolved before the buyer makes its production decision. The buyer can commit early (before the forecast update) or later (after the forecast update). They consider a wholesale price contract with a single type of buyer and single production opportunity. *Iyer and Bergen* (1997) study how the retailer's and the manufacturer's profits change when the retailer orders before or after a demand forecast update. *Gurnani and Tang* (1999) study a two-period model where the buyer updates his demand forecast in period 2 and can place orders in both periods. Assuming the unit cost in the second period is uncertain and could be higher or lower than the unit cost in the first period, they provide conditions under which the buyer may prefer to delay her order. Similar to these papers, *Brown and Lee* (1997), *Donohue* (2000), *Huang et al.* (2005), *Barnes-Schuster et al.* (2002), *Seifert et al.* (2004), and *Erhun et al.* (2008) study multiple ordering opportunities where a delayed commitment can be either purchased upfront as an option or purchased later at a higher per-unit cost for symmetric information scenarios. A common modeling assumption of all of these papers is that the supplier fully knows the buyer's demand information and therefore he does not act strategically. *Courty and Hao* (2000) study a screening contract where consumers know at the time of contracting only the distribution of their valuations, but subsequently learn their actual valuations. The seller offers a menu of refund contracts, specifying an advanced payment and a refund that can be claimed after the consumer's valuation is realized. Under such a contract, the consumer is sequentially screened, as in our contingent contract. However, the context and the model of their paper are significantly different as they focus on valuation uncertainty with a single update while we consider demand forecast accuracy in a supply chain management problem. Finally *Oh and Özer* (2012) consider a problem of a supplier selling to a manufacturer when both parties can obtain asymmetric demand forecast for the same product. The supplier decides when to build capacity, how much capacity to build, whether to offer a menu of contracts to elicit private forecast information from the manufacturer, and if so, what contract to offer. They provide a capacity reservation contract which can be close to optimal. While they study how the contract terms are affected by demand forecast and costs, while we focus on comparing the performance of different types of contracts, mechanisms to induce retailer to obtain higher forecasts accuracy and investigating the effects of increased forecast accuracy on the supplier's and the buyer's profit.

2.3 Model and Preliminary Results

2.3.1 Model

We consider a supply chain composed of a single supplier (he) and a single buyer (she). At the beginning of the season, both the supplier and the buyer have priors on the buyer's demand distribution but do not know the realization. For simplicity, we will restrict our analysis to the case where the buyer is expected to have either high (H) or low (L) demand, with priors p_H and p_L respectively. We model the information asymmetry by assuming that based on experience with the market, past sales, etc., the buyer can privately observe information about her demand type (high or low) in period 1. The buyer who receives a high (low) demand signal is called high (low) type buyer. In period 2, the buyer can update her demand forecast to be more accurate. The supplier, on the other hand, only has priors on the buyer's demand type at all times.

Below, we provide further details of the buyer's demand forecast evolution, the buyer's revenue, and the supplier's choices of contract types to offer to the buyer.

Demand Forecast Evolution

In period 1, the buyer observes a demand signal S_1, which is type $i \in \{L, H\}$ with probability p_i^1. The accuracy of the period 1 signal is denoted by θ_1, such that the buyer's actual demand type coincides with the signal of S_1 with probability θ_1. We assume $\theta_1 \in [\max(p_L, p_H), 1)$ so that the observed signals provide additional information regarding the buyer's demand type. In period 2, the buyer observes another demand signal S_2, which is of type $j \in \{L, H\}$ with probability p_{ij}. The period 2 signal is accurate with probability θ_2, where $\theta_2 \geq \theta_1$ to reflect the improvement of demand forecast accuracy over time. We assume that the more accurate information overwrites the less accurate one. That is, after the buyer observes the period 2 signal, her actual demand type will match the period 2 signal with probability θ_2, and the period 1 signal becomes irrelevant.[1] Finally, at the end of the second period, the buyer will observe her actual demand type $\xi \in \{L, H\}$, and realize the actual demand. If $\xi = k$, then her demand realization will be $D_k = \mu_k + \epsilon, k \in \{L, H\}$,

[1] We note that our model is similar to that adopted by Taylor (2006) except for the fact that in the current model, the buyer can obtain a second signal which is more accurate, whereas in Taylor's model, there is only one signal before demand is realized.

where μ_k is the mean of actual demand type k, and ϵ is a zero-mean random variable, whose cumulative distribution function (cdf) F is continuous and differentiable over $[-\delta, \delta]$. This variable ϵ represents the idiosyncratic risk that affects both demand types, referred to as "market uncertainty."

Table 1 summarizes the notation used in this chapter.

Table 2.1: Notation used in Chapter 2

Notation	Math. Definition	Value when $i = L$	Description
S_1	S_1	L if $S_1 = i$	Period 1 signal of demand type
S_2	S_2	L if $S_2 = i$	Period 2 signal of demand type
ξ	ξ	L if $\xi = i$	Actual demand type
D_i	D_i	D_L	Demand of type i
μ_i	μ_i	μ_L	Mean of demand type i
θ_1	$P(\xi = i \mid S_1 = i)$	θ_1	Accuracy of period 1 forecast
θ_2	$P(\xi = i \mid S_2 = i)$	θ_2	Accuracy of period 2 forecast
ϵ	ϵ	ϵ	Market uncertainty
δ	δ	δ	Parameter controlling the support of the market uncertainty
p_i	$P(\xi = i)$	p_L	Unconditional probability of having demand type i
p_i^1	$P(S_1 = i)$	$\frac{\theta_1 + p_L - 1}{2\theta_1 - 1}$	Probability of observing signal of demand type i in period 1
p_i^2	$P(S_2 = i)$	$\frac{\theta_2 + p_L - 1}{2\theta_2 - 1}$	Probability of observing signal of demand type i in period 2
p_{ij}	$P(S_2 = j \mid S_1 = i)$	$p_{LL} = p_{HH} = \frac{\theta_1 + \theta_2 - 1}{2\theta_2 - 1}$, $p_{LH} = p_{HL} = 1 - p_{LL}$	Probability of observing signal type j in period 2, given that the signal observed in period 1 is type i
p_{ij}^1	$P(\xi = j \mid S_1 = i)$	$p_{LL}^1 = p_{HH}^1 = \theta_1$, $p_{LH}^1 = p_{HL}^1 = 1 - \theta_1$	Probability of having demand type j, given that the signal observed in period 1 is type i
p_{ij}^2	$P(\xi = j \mid S_2 = i)$	$p_{LL}^2 = p_{HH}^2 = \theta_2$, $p_{LH}^2 = p_{HL}^2 = 1 - \theta_2$	Probability of having demand type j, given that the signal observed in period 2 is type i

Buyer's Revenue

We define $\Gamma(D, q)$ as the buyer's revenue from selling q units in a market with demand $D \in \{D_L, D_H\}$. Let $\Gamma'(D, q) := \frac{d\Gamma(D,q)}{dq}$ and $\Gamma''(D, q) = \frac{d^2\Gamma(D,q)}{dq^2}$.

Assumption 2.1. *: $\Gamma(D, q)$ satisfies the following properties.*

1. *$\Gamma(D_i, q) \geq \Gamma(D_j, q)$ if $D_i \succcurlyeq D_j$ (where $\succcurlyeq$ indicates stochastic ordering).*

2. $\Gamma(D, q_1) \geq \Gamma(D, q_2)$ *if* $q_1 \geq q_2$.

3. $\Gamma(D_i, q_1) - \Gamma(D_i, q_2) \geq \Gamma(D_j, q_1) - \Gamma(D_j, q_2)$ *if* $D_i \succcurlyeq D_j$ *and* $q_1 \geq q_2$.

4. $\Gamma''(D_i, q) \leq 0$ *and* $-\Gamma''(D_i, q)$ *is unimodal in* q *for all* i.

These four properties are satisfied by many revenue functions commonly used in the contracting literature. Property 1 to 3 characterize natural behavior that the revenue should increase in demand and the quantity that the buyer has available for sales. Property 4 helps guarantee the unimodality of the supplier's profit in contract quantities. We will discuss two of the most standard revenue models that satisfy these properties.

Exogenous price with salvage value: If the market is highly competitive and the buyer has limited pricing power, the retail price r is exogenous to the system. Let $s, 0 \leq s < r$, be the salvage value that the buyer can obtain for each unsold unit. Then, the buyer's revenue $\Gamma(D, q)$ is given by $rE\min(D, q) + sE(q - D)^+$. This revenue satisfies Properties 1-3. As long as the density of the market uncertainty ϵ is unimodal (e.g., Normal, Uniform, Exponential), Property 4 is satisfied as well. In this model, the retail price and salvage value are public information, known to both the supplier and the buyer prior to their contracting. The buyer observes her demand signals, then chooses a contract providing a quantity q and charging a transfer payment t, which maximizes her expected profit of $rE\min(D, q) + sE(q - D)^+ - t$.

Endogenous price: If the buyer has pricing power, then we need to define a demand response function. Suppose the demand curve of type $\xi \in \{L, H\}$ is linear in retail price r, and is given by $D(r, \xi) = a + \mu_\xi + \epsilon - br$, similar to *Taylor* (2006). We assume $\mu_L < \mu_H$, and hence, $D(r, L) \preccurlyeq D(r, H)$. The buyer sets the optimal retail price. Without loss of generality, we assume $a = 0$, and normalize b to 1. Then, for a buyer type ξ with q units for sale, the optimal retail price is $\min(q, \frac{\mu_\xi+\epsilon}{2})$, and the resulting revenue is given by $\Gamma(D_\xi, q) := (\mu_\xi + \epsilon - \min(q, \frac{\mu_\xi+\epsilon}{2}))\min(q, \frac{\mu_\xi+\epsilon}{2})$. It is easy to check that $\Gamma(D_\xi, q)$ satisfies all four properties. In this model, prior to contracting, the buyer's demand curve as a function of demand type is known to both the supplier and the buyer, and the supplier knows that the buyer will set the retail price that maximizes her revenue. The buyer chooses a contract from the menu based on her observed demand signals. After the total order quantity q is delivered and the actual demand type ξ and market uncertainty ϵ are realized at the end of period 2, the buyer sets the corresponding optimal retail price $\min(q, \frac{\mu_\xi+\epsilon}{2})$.

Types of Contracts

We assume that the supplier is powerful enough to offer the buyer a menu of take-it-or-leave-it contracts. If a traditional one-time contract is to be offered, the supplier has options to offer the contract in period 1, before the buyer obtains a more accurate demand forecast (*early static contract*), or in period 2, after an improved demand forecast has been received (*late static contract*). In this essay, we also consider another possibility where the supplier can offer a menu of contracts that span both periods. The first menu is offered in period 1, and the second menu contingent on the first contract is offered in period 2 (*dynamic contract*).

The supplier has to produce at least the quantity contracted with the buyer. He can produce in period 1 and/or period 2 but the deliveries occur at the end of period 2. The production cost in period $t \in \{1, 2\}$ is c_t, where $0 < c_1 \le c_2$. Notice that while producing in period 1 is less expensive, it exposes the supplier to overproduction or underproduction risks if the buyer has the option to order in the second period [2].

Dynamic Contract: The supplier offers the following menu of contracts in period 1:

$$(q_H, t_H)\Big\{(q_{HH}, t_{HH}), (q_{HL}, t_{HL})\Big\}, \qquad (q_L, t_L)\Big\{(q_{LH}, t_{LH}), (q_{LL}, t_{LL})\Big\}$$

If the buyer chooses (q_i, t_i) in period 1, she pays t_i for the initial order quantity q_i. After choosing (q_i, t_i), she can re-order from the menu $\{(q_{iH}, t_{iH}), (q_{iL}, t_{iL})\}$ in

[2] We assume that the inventory holding cost is negligible without loss of generality.

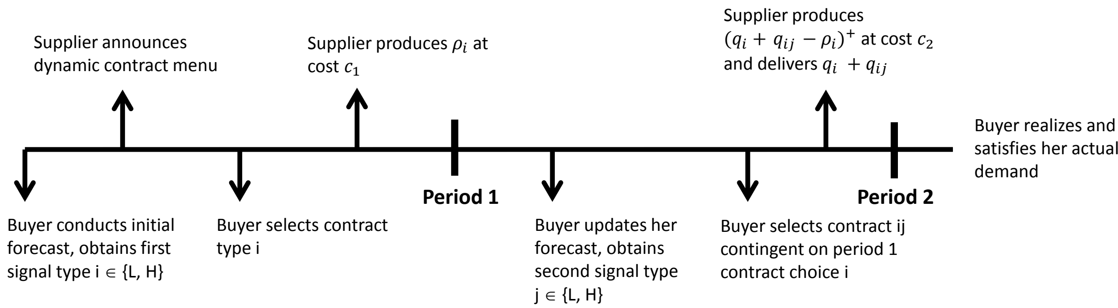

Figure 2.1: Sequence of events: Dynamic Contract

period 2. She pays t_{ij} for the additional order quantity q_{ij}. The total order $q_i + q_{ij}$ is delivered at the end of period 2. Notice that the contract (q_i, t_i) is meant for the buyer who observes a signal i in period 1, and the contingent contract (q_{ij}, t_{ij}) is meant for the buyer who subsequently observes a signal j in period 2.

The supplier decides how much to produce upfront in period 1 after the buyer makes the initial selection from the period 1 menu of contracts. We define ρ_i as the supplier's decision variable of the production quantity in period 1, given that the buyer chooses the type i contract from the period 1 menu. The benefit of producing in period 1 is the cheaper unit production cost. However, delaying part of production to period 2 allows the supplier to produce after learning exactly how much the buyer will order in total, and hence, reduces the risk of over- or underproduction. The sequence of events with dynamic contract is displayed in Figure 1.

The supplier's optimization problem under the dynamic contract is given by:

$$\max_{\mathbf{q},\mathbf{t},\rho} \left\{ \sum_{i\in\{L,H\}} p_i^1(-c_1\rho_i + t_i) + \sum_{i\in\{L,H\}} p_i^1 \sum_{j\in\{L,H\}} p_{ij}[t_{ij} - c_2(q_i + q_{ij} - \rho_i)^+] \right\} \tag{2.1}$$

$s.t.$ Period 1 Participation Constraints

$$\sum_{j\in\{L,H\}} p_{ij} \sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i + q_{ij}) - t_{ij}] \geq t_i,\ i \in \{L,H\}$$

Period 1 Incentive Constraints

$$\sum_{j\in\{L,H\}} p_{ij} \sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i + q_{ij}) - t_{ij}] - t_i \geq$$
$$\sum_{j\in\{L,H\}} p_{ij} \sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_{-i} + q_{(-i)l}) - t_{(-i)l}] - t_{-i},\ i \in \{L,H\},\ l \in \{L,H\}$$

Period 2 Participation Constraints

$$\sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i + q_{ij}) - t_{ij}] \geq \sum_{k\in\{L,H\}} p_{jk}^2\Gamma(D_k, q_i),\ i \in \{L,H\},\ j \in \{L,H\}$$

Period 2 Incentive Constraints

$$\sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i + q_{ij}) - t_{ij}] \geq \sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i + q_{i(-j)}) - t_{i(-j)}],$$
$$i \in \{L,H\},\ j \in \{L,H\}$$

Nonnegativity Constraints

$$\rho_i, q_i, q_{ij}, t_i, t_{ij} \geq 0\ i \in \{L,H\},\ j \in \{L,H\}$$

The first term in the objective function includes the initial payment and period 1 production cost $c_1\rho_i$. The second term accounts for the period 2 payment and the remaining production cost $c_2(q_i + q_{ij} - \rho_i)^+$ for the total quantity ordered by the buyer. The first constraint is the participation constraint that guarantees the type-i buyer's expected profit from the whole horizon is non-negative in period 1. The second constraint is the incentive compatibility constraint, which ensures that the type-i buyer selects the contract designed for her type in period 1. Similarly, the third and fourth constraints are the participation and incentive compatibility constraints in period 2. They guarantee non-negative expected profits from participating in the

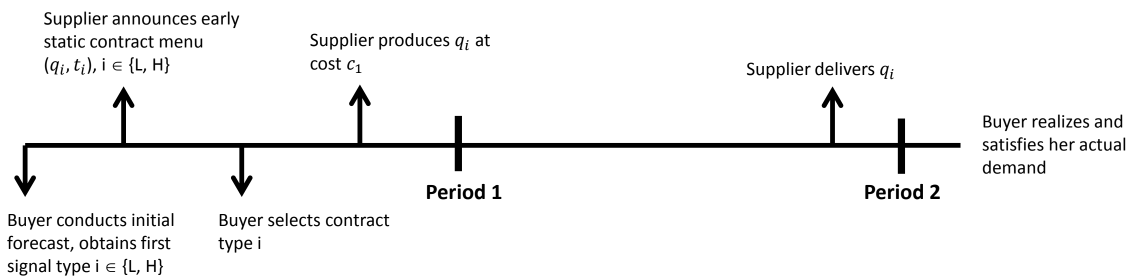

Figure 2.2: Sequence of events: Early Static Contract

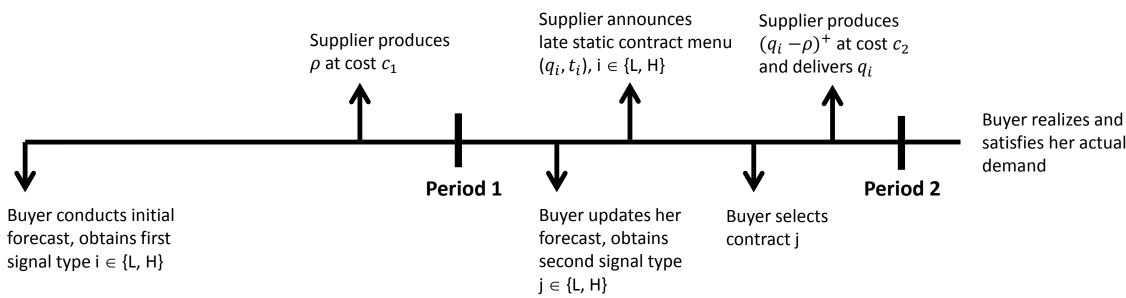

Figure 2.3: Sequence of events: Late Static Contract

period 2 contracts, and maximum expected profits from committing to the contract corresponding to the buyer's second demand signal type.

Static Contracts: The early static and late static contracts are special cases of dynamic contract. More precisely, under an early static contract, the supplier announces a menu of contracts $\{(q_H, t_H), (q_L, t_L)\}$ in period 1 to screen the buyer's period 1 signal type. Hence, it can be viewed as a dynamic contract with constraints $q_{ij} = 0$ and $t_{ij} = 0,\ i, j \in \{L, H\}$ in period 2. Under a late static contract, the supplier offers a menu of contracts $\{(q_H, t_H), (q_L, t_L)\}$ in period 2 to screen the buyer's period 2 signal type. Hence, it is equivalent to a dynamic contract with constraints $q_i = 0$ and $t_i = 0, i \in \{L, H\}$ in period 1. The sequence of events with early and late static contract are given in Figure 2 and 3, respectively. In Appendix A, we provide the supplier's optimization problems and solutions of the early static and late static contracts.

2.3.2 Preliminary Results

Propositions 2.1 through 2.3 characterize the structure of an optimal dynamic contract (The proofs are provided in Appendix A). We will use these properties in the next section when we discuss which contract structure (dynamic, early static or late static) is most beneficial for the buyer or seller under different conditions.

There are multiple contracts that result in the same expected profit for the buyer and the supplier. Proposition 2.1 shows that in one form of the optimal dynamic contracts, all contract quantities are deferred to the second period contracts ($q_i = 0, i \in \{L, H\}$ in period 1). In this contract, the supplier charges t_i in period 1 as an option price, which gives the buyer the right to order $q_i + q_{iH}$ or $q_i + q_{iL}$ in period 2, and pay the additional fee t_{ij} if necessary. The buyer will have the total order, $q_i + q_{ij}$, delivered by the end of period 2. This contract structure is similar to that of a capacity reservation contract commonly used in practice.

Proposition 2.1. *For an optimal dynamic contract with contract quantities* $\{q_i, q_{ij}\}$, $i, j \in \{L, H\}$, *there exists an equivalent dynamic contract with* $q'_i = 0$ *and* $q'_{ij} = q_i + q_{ij}$, $i, j \in \{L, H\}$.

Similarly, we can show that the supplier can transfer the payments of the period 2 low-type contracts to period 1 ($t_{iL} = 0, i \in \{L, H\}$ in period 2) without losing optimality. Proposition 2.2 states that there exists an optimal dynamic contract such that if the second forecast indicates the demand is low (i.e. the buyer observes HL or LL), then the buyer is not charged another fee in the second period. Only when the buyer needs additional units to meet expected high demand, she has to pay an extra fee in the second period.

Proposition 2.2. *For an optimal dynamic contract with transfer payments* $\{t_i, t_{ij}\}$, $i, j \in \{L, H\}$, *there exists an equivalent dynamic contract with* $t'_i = t_i + t_{iL}$, $t'_{iH} = t_{iH} - t_{iL}$, *and* $t'_{iL} = 0$, $i \in \{L, H\}$.

By Proposition 2.1 and 2.2, we can construct an equivalent dynamic contract starting from any other optimal dynamic contract in the following form: $(0, t_H)\{(q_{HH}, t_{HH}), (q_{HL}, 0)\}$, $(0, t_L)\{(q_{LH}, t_{LH}), (q_{LL}, 0)\}$. Proposition 2.3 characterizes the supplier's optimal production policy and the structure of an optimal dynamic contract in this simplified form. The structure of the optimal early static and late static contract are characterized in Appendix A.

Proposition 2.3. *If the buyer selects type i contract in period 1, then the supplier's period 1 optimal production is*

$$\rho_i^*(\boldsymbol{q}) = \begin{cases} q_{iH} & \text{if } \frac{c_1}{c_2} \leq p_{iH} \\ q_{iL} & \text{if } \frac{c_1}{c_2} > p_{iH} \end{cases}$$

The optimal contract is not unique. Under one optimal contract, the payments to the supplier by the buyer are given by

$$\begin{aligned} t_L &= (1-\theta_1)\Gamma(D_H, q_{LL}) + \theta_1 \Gamma(D_L, q_{LL}) \\ t_{LH} &= \theta_2[\Gamma(D_H, q_{LH}) - \Gamma(D_H, q_{LL})] + (1-\theta_2)[\Gamma(D_L, q_{LH}) - \Gamma(D_L, q_{LL})] \\ t_H &= \theta_1 \Gamma(D_H, q_{HL}) + (1-\theta_1)\Gamma(D_L, q_{HL}) - (2\theta_1 - 1)[\Gamma(D_H, q_{LL}) - \Gamma(D_L, q_{LL})] \\ t_{HH} &= \theta_2[\Gamma(D_H, q_{HH}) - \Gamma(D_H, q_{HL})] + (1-\theta_2)[\Gamma(D_L, q_{HH}) - \Gamma(D_L, q_{HL})] \\ t_{LL} &= t_{HL} = 0. \end{aligned}$$

The optimal dynamic contract quantities can be obtained by solving the following *equations:*

$$-[\frac{(2\theta_1-1)}{p_L^1}+p_{LH}\theta_2-\theta_1][\Gamma'(D_H,q_{LL})-\Gamma'(D_L,q_{LL})]+p_{LL}\Gamma'(D_L,q_{LL})$$
$$=(c_1-p_{LH}c_2)^+$$
$$\theta_2\Gamma'(D_H,q_{LH})+(1-\theta_2)\Gamma'(D_L,q_{LH})=\min\left(c_2,\frac{c_1}{p_{LH}}\right)$$
$$[\frac{(1-\theta_2)(\theta_2-\theta_1)}{(2\theta_2-1)}]\Gamma'(D_H,q_{HL})+[\frac{\theta_2(\theta_2-\theta_1)}{(2\theta_2-1)}]\Gamma'(D_L,q_{HL})=(c_1-p_{HH}c_2)^+$$
$$\theta_2\Gamma'(D_H,q_{HH})+(1-\theta_2)\Gamma'(D_L,q_{HH})=\min\left(c_2,\frac{c_1}{p_{HH}}\right).$$

If $q_{LL}>q_{LH}$*, it is optimal to bunch the quantity for the low-type's period 2 contracts and offer* $q_{LL}=q_{LH}=\bar{q}_L$ *which satisfies* $(1-(1+p_H^1)\theta_1)\Gamma'(D_H,\bar{q}_L)+((1+p_H^1)\theta_1-p_H^1)\Gamma'(D_L,\bar{q}_L)=c_1p_L^1$.
If $q_{HL}>q_{HH}$*, it is optimal to bunch the quantity for the high-type's period 2 contracts and offer* $q_{HL}=q_{HH}=\bar{q}_H$ *which satisfies* $p_H^1\theta_1\Gamma'(D_H,\bar{q}_H)+p_H^1(1-\theta_1)\Gamma'(D_L,\bar{q}_H)=c_1p_H^1$.

An optimal dynamic contract can always be fully characterized as long as the buyer's revenue function $\Gamma(D,q)$ is known and satisfies the properties in Assumption 1. We also note that the task of solving for the optimal transfer payments and contract quantities for a dynamic contract has essentially the same difficulty level as designing a conventional static contract. The major difference between offering a static and a dynamic contract is that a static contract distinguishes only between the two types of the buyer (H and L in period 1 or period 2); while, a dynamic contract screens for four different types of the buyer (HH, HL, LH, LL), based on all the possible combinations of period 1 and period 2 signals observed by the buyer.

2.4 Contract Preferences and Effects of Forecast Accuracy

We now address the question of which contracts are most profitable for the supplier or the buyer under which conditions. First, we assume that the buyer will always obtain a more accurate forecast in period 2. As noted earlier, early static and late static contract are special cases of dynamic contract. Hence, it is straightforward to see that from the supplier's point of view, dynamic contract weakly dominates early static and late static contracts. The advantage of dynamic contract over early static contract comes from the supplier's ability to screen not only the initial demand estimate, but also the improved demand information, under the dynamic contract. This allows the supplier to potentially sell more to the buyer who observes the high-type signal in the second period. In comparison with late static contract, the superiority of dynamic contract comes from its structure that enables the supplier to screen the initial demand forecast. By learning about the buyer's type upfront in period 1, the supplier can make a better production decision, resulting in cheaper production costs under a dynamic contract.

Given that the supplier would always prefer to use the dynamic contract (compared to early static or late static contracts), an interesting question is whether the buyer's receiving better forecasts over time is beneficial to the supplier. There are two ways to address this question. First, we note that the buyer receives a second signal with accuracy $\theta_2 > \theta_1$ in period 2. It is straightforward to see that the supplier would always prefer that the buyer receive this second signal. That is, if the buyer did not receive this second (more accurate signal) or if the buyer received a second signal but the accuracy of this second signal was identical to the first period, the supplier would definitely be worse off. Indeed, the better the accuracy of the signal that the buyer receives in period 2, the higher profits the supplier can receive so long as he uses the dynamic contract.

Theorem 2.1. *The supplier's profit under an optimal dynamic contract monotonically increases with the buyer's second-period forecast accuracy,* θ_2.

We would like to point out how Theorem 2.1 complements the analysis of *Taylor* (2006). In that paper, Taylor (using a model similar to ours but with a single period analysis) analyzed a situation where the buyer received only one signal and showed that increasing the accuracy of that signal does not necessarily increase the profits of the supplier. In our setting, if the supplier used a late static contract, increasing the accuracy of θ_2 would not always increase the profits of the supplier, similar to Taylor's result. Likewise, if the buyer uses an early static contract, the only relevant forecast signal is the first one and an increase in the accuracy of this signal does not always increase the profits of the supplier either. However, a contrasting and interesting result we show here is that as long as the supplier uses the dynamic contract we specify above, a second (improved) forecast always benefits the supplier, and the more improved the forecast is, the greater the benefit to the supplier. This is because with a dynamic contract, the supplier screens both the initial and more accurate demand signals, allowing him to effectively extract most of the potential gain from the reduction in the mismatch between the buyer's ordered quantity and actual demand without having to pay high rents to the buyer. An example of the supplier's profit under the three contract types as the forecast accuracy improves is shown in Figure 4.

It is also worth pointing out that the dynamic contract can be more profitable to the supplier than the early static contract even when the buyer contracting under the early static contract has an initial forecast that is more accurate than the improved forecast of the buyer contracting under the dynamic contract. That is, even when the demand information the supplier obtains from a dynamic contract is inferior to that obtained from an early static contract, the supplier's profit can still be higher under the dynamic contract. Proposition 2.4 provides sufficient conditions for this scenario.

When the demand is more likely to be low ($p_L \geq 0.5$), if the buyer does not learn additional information from a demand forecast ($\theta_1 = \max\{p_L, p_H\} = p_L$), then the buyer will always buy the low-type contract under the early static contract ($p_L^1 = 1$). If the buyer obtains a more informative demand forecast ($\theta_1 > p_L$), then with a positive probability, the buyer will observe a high signal and select the high-type contract. However, when the additional cost to sell to the high-type buyer (production cost $c_1(q_H - q_L)$ and high-type rent $(2p_L - 1)[\Gamma(D_H, q_L) - \Gamma(D_L, q_L)]$) is greater than the expected gain from having a high demand ($p_L[\Gamma(D_H, q_H) - \Gamma(D_H, q_L)] + p_H[\Gamma(D_L, q_H) - \Gamma(D_L, q_L)]$), the supplier finds it less profitable to sell to the high-type. In this case, the supplier's profit under an early static contract is decreasing in the forecast accuracy θ_1 when the accuracy is small.[3] Hence, the supplier can earn larger profits from a dynamic contract even when the buyer's accuracy is lower.

Proposition 2.4. *If demand is more likely to be low ($p_L \geq 0.5$) and if the optimal early static contract when $\theta_1 = p_L$ is such that $c_1(q_H - q_L) + (2p_L - 1)[\Gamma(D_H, q_L) - \Gamma(D_L, q_L)] > p_L[\Gamma(D_H, q_H) - \Gamma(D_H, q_L)] + p_H[\Gamma(D_L, q_H) - \Gamma(D_L, q_L)]$, then there exists $\bar{\theta} \in (0, 1]$ such that the supplier's profit under the optimal early static contract with a buyer whose period 1 accuracy is $\theta < \bar{\theta}$ is less than the supplier's profit under an optimal dynamic contract with a buyer whose period 2 accuracy is $\theta_2 < \theta$.*

While the supplier can always benefit from more accurate forecasts with a dynamic contract, it is required that the buyer obtains a more accurate forecast in the second period for a dynamic contract to be viable. This actually raises two interesting questions: 1) Is the buyer willing to obtain better forecasts?; i.e., do better forecasts also always benefit the buyer assuming the buyer has the capability to obtain them?, and 2) What if the supplier is uncertain about the buyer's capability to obtain more accurate forecasts in the first place? To address these two questions, we first develop an understanding of which contracts are most profitable for each type of the buyer.

[3] Notice that this result is analogous to Taylor and Xiao (2010).

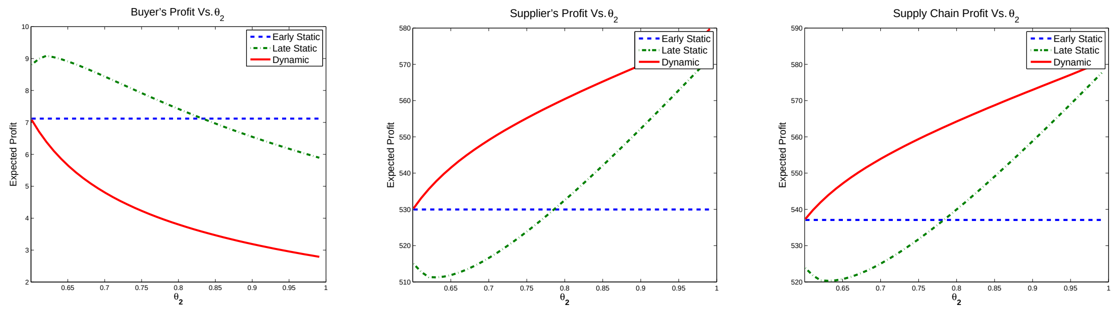

Figure 2.4: Preferences of the high-type buyer, supplier, and supply chain: $r = 1, c_1/c_2 = 0.67, c_1 = 0.1, p_L = 0.46, \mu_H = 900, \mu_L = 400,\ \delta = 50, \theta_1 = 0.6)$

Proposition 2.5 presents the contract preferences of the low-type buyer (who observes a low demand signal in period 1).

Proposition 2.5. *The low-type buyer always prefers late static contract to early static and dynamic contract. She is indifferent between early static and dynamic contracts.*

Under early static and dynamic contract, the low-type buyer commits to the low-type contract in period 1, before she obtains a more accurate demand forecast. Hence, she is screened as the lowest type, and is awarded zero expected profits since the low-type participation constraints in period 1 under both early static and dynamic contract are binding at optimality. The supplier sets the quantity and transfer payment such that the low-type buyer makes a positive profit only if her actual demand turns out to be high-type; she loses money ex-post otherwise. If a late static contract is offered, however, the low-type buyer has a chance to observe an improved demand signal in period 2, before she commits to a contract. With a positive probability, her second signal can be high-type and she can receive a positive expected profit from the high-type contract. Otherwise, if her second signal is low-type, she receives a zero expected profit. Thus, her ex-ante expected profit is positive under a late static contract. Since the low-type buyer prefers to contract late and is indifferent between early static and dynamic contract, she would always agree to update her forecasts even when the supplier offers her a dynamic contract.

The situation for the high-type buyer is different. In a screening contract, the profit to the high-type buyer comes from the information rent that the supplier has to offer in order to prevent the high-type from deviating to lower type contracts. Hence, the high-type buyer always makes positive profits under all three types of contracts. However, it is not immediate under which situations, the high-type buyer will prefer which contract type to the others. Proposition 2.6 shows that the high-type buyer always prefers to contract early rather than dynamically. Furthermore, under the dynamic contract, the high-type buyer's expected profit is hurt even more as her second period information accuracy improves. This is because under early static contract, the buyer only reveals her less accurate demand information, leaving sufficient amount of uncertainty which results in higher rents. On the other hand, the buyer reveals both her initial and improved demand information under dynamic contract, leaving little rents to her. The additional demand information revealed in the second period always benefits the supplier rather than the buyer because it decreases the uncertainty about the buyer's type.

Proposition 2.6.

1. *The high-type buyer's profit under the early static contract is at least as high as her profit under the dynamic contract.*

2. *The high-type buyer's profit under the dynamic contract monotonically decreases with the second-period forecast accuracy.*

Since the early static contract is more profitable to the high-type buyer than the dynamic contract, if the buyer expects the supplier to offer her a dynamic contract, she would opt out from conducting a more accurate forecast in order to be offered an early static contract instead. However, if both the supplier's and the buyer's profit are higher with the late static than with the early static contract, then the supplier can benefit from offering the late static contract upfront (in a sense, take the dynamic contract off the table), so that the buyer would agree to obtain a more accurate forecast. In this case, both parties can still benefit from more accurate demand information. An example of such situation where the late static contract is more preferable to both parties than the early static contract is shown in Figure 4, when the period 2 accuracy, θ_2, is between 0.78 and 0.83. Notice that the supplier prefers the late static contract only when the forecast accuracy is sufficiently increased in period 2 ($\theta_2 > 0.78$). This is because the significant improvement in the accuracy of demand forecasts makes it worth waiting to contract late even though the supplier has to incur higher production costs. For the high-type buyer, she prefers late static to early static contract when the period 2 accuracy is sufficiently low ($\theta_2 < 0.83$). This is because a moderate accuracy of her signal leaves enough uncertainty about her demand type, and the supplier, after waiting to contract in period 2, would be willing to offer her a higher rent in exchange for the more accurate and only demand information. This situation is particularly prevalent when the period 2 production cost is not much more expensive than the period 1 production cost (signified by a large $\frac{c_1}{c_2}$). It is worth noting that the supply chain can also benefit from more accurate demand forecast through dynamic and late static contract, especially when the increase in accuracy is substantial. This is because when the period 2 accuracy is high, the value of more accurate demand information outweighs the increase in production costs in the later period.

This section has shown that while the supplier always benefits most from, and therefore, favors the dynamic contracts, the buyer may actually prefer the early static contract and may claim that she will not obtain a more accurate forecast as the high type buyer's profit is monotonically decreasing in period 2 accuracy. But if the supplier knows that the buyer is capable of obtaining a more accurate forecast, how does the supplier interpret this refusal to obtain a better forecast? We answer this question in the next section.

2.5 Capable Buyer and Bayesian Updating Supplier

In the previous section, we showed that although there are some situations where both parties benefit from more accurate forecasts, there are others where the buyer does not. Since the buyer is the one obtaining the better forecast, the question is whether the buyer can simply announce that she is not obtaining a second forecast, when she knows that having the better forecast will put her under a less profitable contract type. The problem is that if the supplier knows that the buyer is capable of obtaining improved forecasts, then a decision by the buyer not to obtain them can lead the supplier to update his beliefs about the buyer's demand type, which will affect the type of contracts he will offer to the buyer.

We study the Perfect Bayesian Equilibrium (PBE) of the two-stage game played by the supplier and the buyer. In period 1, the buyer observes her own period 1 demand signal, and then announces whether she intends to obtain the second demand signal or not. If she does, the supplier offers a dynamic contract menu (since it is always most profitable for him). If the buyer does not update, the supplier offers an early static contract menu. Note that if the Bayesian updating leads the supplier to be certain about the buyer's type, then the supplier can offer a first-best contract for the buyer's type.

In this Bayesian game, the supplier has an initial belief about the buyer's period 1 signal type described by the probabilities $p_i^1,\ i \in \{L, H\}$. Let $\{U, N\}$ be the set of the buyer's possible strategies where U is to update, and N is to not update. The buyer with signal type i chooses U with probability σ_i^U, and N with probability $\sigma_i^N = 1 - \sigma_i^U$. After the buyer announces her strategy $S \in \{U, N\}$, the supplier updates his belief about the type distribution to $\tau_i(S),\ i \in \{L, H\}$ as follows: $\tau_i(S) = \frac{p_i^1 \sigma_i^S}{\sum_{k \in \{L,H\}} p_k^1 \sigma_i^S}$.

The standard result for this Bayesian game is that there is no separating PBE.

There are two pure strategy pooling PBE: one where both types of buyers announce they will obtain better demand forecasts, and the other where both types of buyers announce they will not obtain better demand forecasts [4]. This result is driven by the fact that the lowest type buyer is indifferent between obtaining and not obtaining a more accurate forecast. If one assumes that the lowest type buyer will obtain a better forecast when she is indifferent, then the equilibrium is the one where all types of the buyer obtain more accurate forecasts. However, if one assumes that when indifferent, the lowest type buyer will actually not obtain better forecasts, then the equilibrium is the one where all types of the buyer choose not to obtain better forecasts, which makes the supplier unable to benefit from more accurate demand information.

To rule out the no-updating PBE, the supplier can in fact utilize a side payment to induce the lowest type buyer to update demand forecasts. More precisely, the supplier can announce upfront that all buyers who obtain a better forecast will be given a small side payment. In this case, the lowest type buyer will have an incentive to obtain a more accurate forecast, leading to a unique all-update equilibrium. This shows that even when the buyer has a bargaining power to refuse updating her demand forecasts, the supplier can offer a dynamic contract with a side payment to guarantee the buyer's willingness to obtain improved forecasts. The only exception where such a dynamic contract may not be preferable to the supplier is when there is a cost involved with updating the forecasts. In this case, the supplier's optimal strategy depends on the update cost, as discussed in Theorem 2.2.

Theorem 2.2. *There exists a threshold $\tilde{K}$ such that if the forecast update costs $K \leq \tilde{K}$, the unique PBE with a side payment is where both types of the buyer obtain more accurate demand forecasts. If $K > \tilde{K}$, the supplier offers an early static contract.*

The cost of obtaining better demand forecast has essentially no impact on the

[4] This result extends naturally to the case where there are $n > 2$ types of buyers, where the types are ordered according to the period 1 demand signal.

buyer's behavior since the cost is transferred to the supplier through the optimal contract design. It is profitable for the supplier to offer a side payment to incentivize both types of the buyer to obtain more accurate demand information at the equilibrium as long as the update cost is reasonable. If the forecast update is so costly that the benefit to the supplier is less than the cost of obtaining better forecasts, then the supplier chooses to offer an early static contract that does not require a second forecast.

An important managerial implication here is that while the buyer may not always benefit from more accurate forecasts, so long as the supplier offers even a small side payment for obtaining better forecasts, he can induce the Bayesian equilibrium where all types of the buyer agree to obtain more accurate forecasts. This is, of course, only possible when the supplier is certain about the buyer's capability of conducting more accurate forecasts in the second period, which may not always be the case in practice. Hence, we discuss the supplier's best contract options in presence of an uncertainty about the buyer's forecasting capability in the next section.

2.6 Screening the Forecasting Capability

When the supplier is truly uncertain as to how accurate of a demand forecast the buyer can obtain, we propose that the supplier can improve the screening nature of his contracts by screening the buyer not only on demand level, but also on forecasting capability. Suppose that the supplier has a prior on the capability of the buyer to obtain improved forecasts. The supplier can design a capability and type screening menu of contracts such that the capable (high accuracy θ_2) buyer prefers dynamic contract, and the incapable (low accuracy θ_1) buyer prefers early static contract.

The model in this section also applies to a situation where the forecast updates are costly, but the supplier is uncertain about the buyer's forecasting costs, as modeled by Lariviere (2002). That is, the buyer may incur a small update cost K_L or a large update cost K_H. If the buyer's update cost is small ($K_L \leq \tilde{K}$), then it is worthwhile for the supplier to induce the buyer to obtain improved forecasts using a dynamic contract. On the other hand, if the buyer's update cost is large ($K_H > \tilde{K}$), then the supplier prefers to contract with the buyer under an early static contract. In this case, the large update cost will never be incurred (since the buyer is offered an early static contract). Then, the buyer with a small update cost is equivalent to the capable buyer, and the buyer with a large update cost is equivalent to the incapable buyer in this model.

The capability to obtain a better demand forecast is private information to the buyer. However, we assume the supplier estimates that with probability ϕ, the buyer is capable of obtaining a more accurate demand forecast, and with probability $1-\phi$, she is not capable. The capable buyer is simply the buyer described in previous sections, who receives a better forecast in period 2. On the other hand, the incapable buyer observes a signal whether she is low-type or high-type in period 1, but is incapable of making that forecast more accurate in period 2. Since the incapable buyer does not receive a second signal, if she purchases a dynamic contract, she uses a strategy such that with probability υ_{iH}, she chooses the iH contract in period 2, and with probability $\upsilon_{iL} = 1 - \upsilon_{iH}$, she chooses the iL contract, where $i \in \{L, H\}$ is her signal type in period 1. Notice that if a buyer chooses the early static contract, then she will always receive the same expected profit, whether she is capable or not, since the contract decision is only made in period 1 and the more accurate forecasting is irrelevant.

The supplier offers a menu of contracts that screens both the type and forecasting capability of the buyer. He offers a menu of the early static contract, $(q_i, t_i), i \in \{L, H\}$, and the dynamic contract, (q_i^D, t_i^D), $\Big((q_{iH}, t_{iH}), (q_{iL}, t_{iL})\Big), i \in \{L, H\}$. The supplier's optimization problem is as follows

$$\max_{\mathbf{q},\mathbf{t},\rho} \Bigg\{(1-\phi)\sum_{i\in\{L,H\}} p_i^1(-c_1 q_i + t_i) + \phi\Big[\sum_{i\in\{L,H\}} p_i^1(-c_1\rho_i + t_i^D) +$$

$$\sum_{i\in\{L,H\}} p_i^1 \sum_{j\in\{L,H\}} p_{ij}[t_{ij} - c_2(q_i^D + q_{ij} - \rho_i)^+]\Big]\Bigg\}$$

$s.t.$ Period 1 Participation Constraints

$$\sum_{j\in\{L,H\}} p_{ij}^1[\Gamma(D_j, q_i) - t_i] \geq 0,\ i \in \{L,H\}$$

$$\sum_{j\in\{L,H\}} p_{ij} \sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i^D + q_{ij}) - t_{ij}] \geq t_i^D,\ i \in \{L,H\}$$

Period 1 Incentive Constraints

$$\sum_{j\in\{L,H\}} p_{ij}^1[\Gamma(D_j, q_i) - t_i] \geq \sum_{j\in\{L,H\}} p_{ij}^1[\Gamma(D_j, q_{-i}) - t_{-i}],\ i \in \{L,H\}$$

$$\sum_{j\in\{L,H\}} p_{ij} \sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i^D + q_{ij}) - t_{ij}] - t_i^D \geq$$

$$\sum_{j\in\{L,H\}} p_{ij} \sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_{-i}^D + q_{(-i)l}) - t_{(-i)l}] - t_{-i}^D,\ i \in \{L,H\},\ l \in \{L,H\}$$

Period 2 Participation Constraints

$$\sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i^D + q_{ij}) - t_{ij}] \geq \sum_{k\in\{L,H\}} p_{jk}^2\Gamma(D_k, q_i^D),\ i \in \{L,H\},\ j \in \{L,H\}$$

Period 2 Incentive Constraints

$$\sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i^D + q_{ij}) - t_{ij}] \geq \sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i^D + q_{i(-j)}) - t_{i(-j)}],$$

Forecasting Capability Incentive Constraints

$$\sum_{j\in\{L,H\}} p_{ij}^1[\Gamma(D_j, q_i) - t_i] \geq \sum_{l\in\{L,H\}}\sum_{j\in\{L,H\}} \upsilon_{il} p_{ij} \sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i^D + q_{il}) - t_{il}] - t_i^D, i \in \{L,H\}$$

$$\sum_{j\in\{L,H\}} p_{ij} \sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i^D + q_{ij}) - t_{ij}] - t_i^D \geq \sum_{j\in\{L,H\}} p_{ij}^1[\Gamma(D_j, q_i) - t_i], i \in \{L,H\}$$

Nonnegativity Constraints

$$q_i, t_i, \rho_i, q_i^D, q_{ij}, t_i^D, t_{ij} \geq 0\ i \in \{L,H\},\ j \in \{L,H\}$$

The optimization problem includes the participation and incentive compatibility constraints for both the early static and dynamic contract. The participation and incentive compatibility constraints of the early static contract ensures that the type $i \in \{L,H\}$ buyer who does not have the forecasting capability picks the early static contract meant for her type. Similarly, dynamic contract constraints ensures that type

$i \in \{L, H\}$ buyer who has the forecasting capability picks the dynamic contract meant for her type. The forecasting capability incentive constraints prevent the incapable buyer from deviating to the dynamic contract by ensuring that, regardless of the mixed strategy chosen by the buyer (any value of v_{ik}), the incapable buyer will do worse by choosing a dynamic contract. Similarly, the second forecasting capability incentive constraint ensures that the capable buyer will prefer to obtain a revised forecast and will not deviate to the early static contract.

In this setting, we have the following results regarding the supplier's and the buyer's optimal strategies and profit.

Theorem 2.3. *Under an optimal two-dimensional screening contract:*

1. *The optimal strategy for the incapable type i buyer if she were to contract dynamically is to choose the iL contract in period 2 with probability 1, i.e. $v_{iL} = 1, i \in \{L, H\}$.*

2. *Both capable and incapable buyer of the same type receive the same expected profit, where the profit is zero for the low-type buyer and is positive for the high-type buyer.*

3. *The supplier's profit monotonically increases with the capability probability ϕ and the period 2 forecast accuracy θ_2.*

4. *The high-type buyer's profit is monotonically decreasing in the capability probability ϕ and is independent of period 2 forecast accuracy θ_2.*

Part 1 and 2 of the theorem are driven by the binding constraints in an optimal two-dimensional screening contract. More precisely, when the buyer is incapable of obtaining a more accurate demand signal to help her make decision on which contract to choose in the second period, she minimizes her risk of losing money from ordering too much by always relying on the low-type contract. This is because an optimal screening contract is designed such that the high-type buyer's expected profit from deviating to the low-type contract is the same as the expected profit from choosing the high-type contract; however, the low-type buyer could lose money from deviating to the high-type contract. The forecasting capability incentive constraints are also binding at optimality, resulting in the same expected profit to the capable and incapable buyer of the same type. Like under other screening contracts, the low-type expected profit is zero while the high-type expected profit is positive due to the information rents.

Part 3 and 4 of the theorem describe the effects of the capability probability as well as the second forecast accuracy on the supplier's and the high-type buyer's expected profit. In Section 4, it has been shown that the supplier always prefers dynamic contract; while, the high-type buyer always prefers early static contract to dynamic contract. Hence, as the capability probability increases, implying it is more likely that dynamic contract will take place, the high-type buyer's expected profit is reduced but the supplier's profit is increased. Figure 5a and b display the expected profits of the high-type buyer and supplier for different values of capability probability. (Note also that an increase in the market uncertainty δ benefits the buyer and hurts the supplier as expected.) If instead the capability probability is fixed, but the capable buyer's second forecast accuracy is improved, we show that the supplier can continue to take advantage of the buyer's better demand information. This result is consistent with Theorem 2.1, where we show the supplier is better off contracting with a more accurate buyer using a dynamic contract. What is less expected here is that a change in the capable buyer's improved forecast accuracy has no effect on the buyer's expected profit. This is because the second forecast accuracy only influences the performance of the capable buyer, but not the incapable buyer. Since the supplier holds the same belief about the buyer's forecasting capability and since it is optimal to offer both capable and incapable buyer the same profit, the supplier gives the same optimal level of rent to the buyer, and keeps the remaining increase in supply chain profit due to better forecast accuracy to himself.

The results in this section show that even in the presence of the uncertainty about the buyer's forecast capability, there still exist a mechanism which allows the supplier to always benefit from more accurate demand information. An implication here is that more accurate information is always potentially beneficial to the supplier as long as he designs a contract wisely.

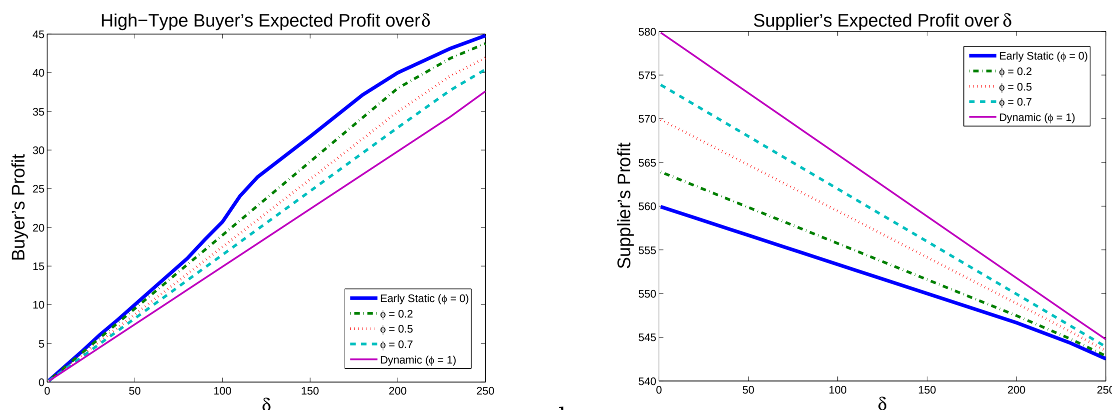

a. b.

Figure 2.5: High-type buyer's and supplier's profits: $r = 1, \frac{c_1}{c_2} = 0.4, c_1 = 0.2, p_L = 0.5, \mu_H = 1000, \mu_L = 600, \theta_1 = 0.6, \theta_2 = 0.8$

2.7 Conclusion

In this chapter, we consider a stylized model of a multi-period procurement game between a powerful supplier and a buyer under demand information asymmetry. We investigate whether and when the parties benefit from more accurate demand forecasts obtained by the buyer under different contract structures. A key finding of our study is that a supplier who knows that a buyer is capable and is going to obtain more accurate forecasts will always benefit from offering a dynamic contract, which screens both the buyer's initial and improved forecasts. We show however that there are cases where the buyer prefers the early static contract the most, and therefore, is disincentivized to obtain more accurate forecasts in the second period. But if the supplier knows the buyer is capable of obtaining a forecast update, we argue that the buyer can not simply refuse to obtain better forecasts as this would lead the supplier to update his beliefs about the buyer's demand type and end up offering a first-best contract to extract even more profit from the buyer. In fact, we show that as long as the supplier is certain that the buyer is capable of obtaining more accurate demand forecasts, he can always induce the buyer to update her demand forecasts by offering a small side payment. An interesting situation is when the supplier may be uncertain about the buyer's capability to obtain more accurate forecasts. In this case, the buyer may claim an inability to obtain more accurate forecasts especially when better forecasts are likely to hurt her profits. But in response, we show that the supplier can offer a more sophisticated contract to screen the buyer both on her forecast updating capability and demand type.

An important conclusion of this chapter is that suppliers can always benefit from demand forecasts obtained by buyers that become more accurate over time. However, to achieve this benefit, the supplier may sometimes be required to design more sophisticated contracts than what are common in practice. Our findings complement what has been found in the existing contract literature which considers mostly simple static contracts and reports that the supplier's profit can be hurt by improved demand forecasts from the buyer. This essay indicates a need for more sophisticated dynamic contracts. We note that similar contracts with contingent clauses have recently been observed in practice. However, the standard dynamic contracts may not be enough when the supplier is facing a buyer who can claim an inability to obtain more accurate forecasts. In this case, the supplier may have to develop dual-screening contracts, which have been much less common in practice. Our finding may indicate that there is a certain point in the supply chain contract's complexity level where the benefit obtained by the contract is outweighed by its complexity. Further research is needed on simpler mechanisms that achieve most if not all of the more sophisticated mechanisms developed here.

CHAPTER 3

Conditional Promotions and Consumer Overspending

3.1 Introduction

Price promotion is a key component in today's retailing. In 2010, the total promotion spending in the United States reached $337 billions, out of which 85% was spent on price discounts (*Borrell Associates*, 2010). Retailers employ price promotions to stimulate sales in their stores, which not only generate larger revenues, but also reduce costs by accelerating the disposal of excess inventory (*Blattberg et al.*, 1981). The most common forms of price promotions include price markdowns (e.g., 20% off from regular price), bundling (e.g., buy shirt and pants together and save $20), and conditional discounts (e.g., buy 2 or more and get 20% off).

Among these, conditional discounts are increasingly more popular in practice. Their use in consumer packaged goods increased by 10.3% in 2010 from 2009. They accounted for 26% of total coupon distribution, and was as high as 33% among grocery products (*NCH Marketing Services*, 2011). Retailers typically offer conditional discounts in the form of either percent-off or dollars-off. For example, Travelodge offered 15% off its regular room rate if customers stayed 2 nights or longer. Bath & Body Works offered $10 off a purchase of $30 or more as a Mother's Day sale.

A distinct feature of conditional discounts is that the deal is always coupled with an eligibility requirement, mostly in the form of a minimum purchase quantity or a minimum spending, that customers need to fulfill in order to receive the discount. This strategy can benefit retailers in several ways: it allows finer price discrimination contingent on purchase quantity or spending, and potentially induces buyers with higher consumption level or reservation price to consume more than what they would without a discount offer.

Most existing literature on price promotions identifies price elasticity as the main reason driving buyers to increase purchase quantity in response to a price reduction (*Jeuland and Narasimhan*, 1985; *Bell et al.*, 1999). However, this argument falls short of explaining why some consumers go the extra mile and buy a significantly larger quantity than what they actually need just to qualify for the deals. In fact, stories about impulsive, deal-driven shopping behavior regularly appear in news articles, popular press magazines, and personal blogs (*Klaft*, Sep. 3, 2011; *Tuttle*, Jul. 23, 2010; *Sherman*, Feb. 5, 2009; *Fontinelle*, Aug. 17, 2011; *Tsai*, 2007).

Such shopping behavior is referred to in the literature as "deal-proneness," which is the propensity to purchase products when they are offered on a "deal" basis (*Hackleman and Duker*, 1980). Existing studies on this subject explain that deal-proneness arises as consumers gain psychological benefits from paying a lower-than-expected price for a product (*Schindler*, 1989, 1998; *Laroche et al.*, 2001; *DelVecchio*, 2005). This enhanced positive feeling associated with the transaction is referred to as "transaction utility" (*Thaler*, 1985). Consumers who receive transaction utility when completing a deal are labeled "deal-prone." Consumers who are not sensitive to such cognitive benefits, and value a discount offer based solely on its monetary value, are called "value-conscious." (*Lichtenstein et al.*, 1990) In order to investigate how deal-proneness influences purchase decisions, we define *overspending* as a situation where a deal-prone consumer increases her purchase quantity purely due to the transaction utility received from completing a deal.

The heterogeneity in consumer deal-proneness together with the heterogeneity in consumer willingness to pay result in multiple consumer segments who make different purchase decisions when facing a discount offer. Thus, when designing a price promotion, retailers must account for the impact of the deal on different groups of consumers. This is a challenging problem as it is not clear how terms of promotional offers influence purchase decisions of heterogeneous consumers, and which discount policy should be adopted for a given set of market parameters. This chapter aims to study the effects of conditional discounts on consumer behavior and seller's profitability, focusing on two most common forms of conditional discounts: i) all-unit discount, where the price reduction (in percentage or dollars off) applies to every unit if the eligibility condition is met, and ii) fixed-amount discount, in which a fixed discount (e.g., \$10 off) is applied to the total expense that satisfies the condition. Note that price markdown is a special case of the all-unit discount with no minimum purchase requirement, and mixed bundling with a limit of one deal per transaction is a special case of the fixed-amount discount.

Our main research questions are: 1) How do different types of consumers respond to conditional discounts? Do conditional discounts induce consumers to overspend? 2) When should a retailer offer conditional discounts?, and 3) What are the market conditions that favor a fixed-amount discount to an all-unit discount or vice versa? To answer these questions, we consider a model of a single seller facing heterogeneous consumers whose marginal consumption surplus decreases in the quantity they consume. Consumers are heterogeneous in two dimensions: their cognitive attitude towards a deal (deal-prone or value-conscious) and valuation (how much they value consumption of the product). Facing these consumers, the seller's problem is to decide whether to offer a discount, and if so, which discount type (all-unit or fixed-amount) and what specific terms of discount to offer in order to maximize his expected profit.

To our knowledge, this essay is the first to analytically investigate consumers' heterogeneity in deal-proneness and their response towards conditional discounts. We show that a conditional discount can induce deal-prone consumers to overspend. Furthermore, we derive the seller's optimal discount strategies and show that a conditional discount is strictly more profitable than selling at the regular price when either consumer willingness to pay for the product is low, or there exist deal-prone consumers in the market. If a sufficiently large proportion of consumers are deal-prone, we show that a conditional discount also dominates a conventional price markdown. Hence, conditional discount is an effective tool to increase the seller's profitability. Interestingly, we find that the optimal terms of discounts (discount depth and minimum purchase requirement) and consumer purchase quantities induced by the optimal discounts are not always monotone in the magnitude of transaction utility. Finally, we identify market conditions under which the all-unit or the fixed-amount discount is more profitable than the other. We show however that regardless of the type of discounts used, it is not always optimal to induce consumers to overspend. Consumer overspending benefits the seller only when there is a sufficient proportion of highly deal-prone consumers in the market.

The rest of the chapter is organized as follows. In Section 3.2, we review relevant literature on sales promotion and consumer deal-proneness. In Section 3.3, we describe the framework, introduce the model, and define the two types of conditional discounts considered in this chapter. Section 3.4 provides analyses of the consumer's problem and discusses consumer purchase behavior under different types of discounts. Section 3.5 addresses the seller's problem and discusses optimal discount policies that maximize the seller's profits. To gain more insights about the effects of each dimension of consumer heterogeneity, we also derive the structure of the optimal discount policies and the induced outcomes under two special cases. Section 3.6 presents numerical study showing that the profit improvement the seller can achieve with the use of conditional discounts is significant. Finally, we discuss extensions of the original model in Section 3.7 and conclude the chapter with discussion of the main results and future research directions in Section 3.8.

3.2 Literature Review

We will review three major streams of literature on sales promotion and consumer deal-proneness. The first stream of work relevant to ours is on price promotions. One of the most common forms of price promotion is a simple price markdown. Early studies on this subject explained sales as a mechanism that increases seller's profit from certain groups of buyers (e.g., informed buyers in *Varian* 1980; low-demand and low-holding cost buyers in *Jeuland and Narasimhan* 1985; brand switchers in *Bell et al.* 1999). Other works analyzed the profitability of periodic price reduction policies and proposed optimization models, taking into account consumer's promotion response and retailer's inventory level (*Lazear*, 1986; *Achabal et al.*, 1990; *Smith and Achabal*, 1998). Under these price markdown mechanisms, the deal is not contingent on the purchase quantity. On the other hand, the current essay considers conditional discounts, where the deal is realized only when buyers purchase at least the minimum required quantity. This fundamental difference in pricing schemes allows us to investigate another potential benefit of promotions in enticing consumers to increase their purchase quantity to complete the deal, which is not observable under price markdowns. A more characteristically similar price promotion mechanism to conditional discount is quantity discount. A number of papers on quantity discount analyzed a cost-minimizing or profit-maximizing problem of a supplier selling to a buyer, where price discounts are given on large order sizes (*Monahan*, 1984; *Lee and Rosenblatt*, 1986; *Dada and Srikanth*, 1987; *Corbett and de Groote*, 2000). Another group of quantity discount literature compared different types of quantity discounts (e.g., all-unit vs. package pricing in *Wilcox et al.* 1987; all-unit vs. incremental discount in *Weng* 1995, and *Munson and Rosenblatt* 1998). However, the comparison between all-unit and fixed-amount discount has not been addressed. Furthermore, the quantity discount literature has been largely focusing on the upstream supply chain parties, rather than a retailer selling to end consumers like in the current essay.

The closest price promotion mechanisms to the conditional discounts, both in terms of structure and usage, are certain forms of bundling. The bundling literature traditionally considered only pure bundling, pure unbundling, and mixed bundling (*Schmalensee*, 1984; *Chuang and Sirbu*, 1999; *Herrmann et al.*, 1997). Other papers broadened their scope to consider bundling schemes where many more combinations of products are offered to consumers (*Hanson and Martin*, 1990; *Venkatesh and Mahajan*, 1993; *Armstrong*, 1996). However, the challenges in pricing all different bundles reduce the attractiveness of bundling when selling a large number of products. This gave rise to the study of customized bundling, where bundles are priced based on quantity rather than specific components, the closest mechanism to conditional discounts. *Hitt and Chen* (2005) and *Wu et al.* (2008) analyzed customized bundling problems and compared the profitability of customized bundling, pure bundling, and individual selling. *Hui et al.* (2008) studied the problem of choosing the optimal number of sizes of bundles to offer under a customized bundling scheme. They concluded that, due to the cognitive cost consumers experience when evaluating many bundle options, a small number of versions are sufficient to capture the majority of the potential profit from versioning. An important difference between customized bundling and conditional discount is that under customized bundling, there are often more than one price breakpoints and discount rates. That is, consumers may receive a deeper discount if they purchase an even larger quantity. On the other hand, under a conditional discount, the discount rate or the discount amount is unified as long as the purchase quantity meets the single minimum requirement. This makes conditional discount simpler to implement and more straightforward to communicate to consumers, which explains why conditional discounts are widely used in retailing.

The second stream of literature is on consumer response to price promotions. A handful of papers in the framing literature compared consumer response to price promotions framed in percentage terms and in dollar terms. For example, *Bitta et al.* (1981), *Chen et al.* (1998), *Hardesty and Bearden* (2003), and *Gendall et al.* (2006) compared the consumer perception of the value of price discounts presented in percentages and dollars. Their experimental results showed that discounts on high-priced items framed in dollars were perceived more significant than the same discounts framed in percentage. These papers have different focus and methodology from ours. They experimentally studied consumer response to different presentations of the exact same price promotion; while, in the current essay, we analytically compare between structurally different price discount policies. Other papers in this stream studied the effects of promotions which require a minimum spending or multiple-unit purchases on consumer purchase behavior (*Wansink et al.*, 1998; *Lee and Ariely*, 2006; *Kivetz et al.*, 2006; *Foubert and Gijsbrechts*, 2007). However, they did not compare the two types of conditional discounts considered in the current essay, and did not study the seller's profit-maximizing promotional strategies.

Finally, the third stream of relevant literature is on consumer deal-proneness. The existence and characteristics of deal-prone consumers have been extensively studied over the past few decades (*Hackleman and Duker*, 1980; *Schindler*, 1989, 1998; *DelVecchio*, 2005; *Kukar-Kinney et al.*, 2012). On a related subject, *Thaler* (1985, 1999), and *Lichtenstein et al.* (1990) employed the theory of mental accounting and reference price to explain the consumer propensity to purchase products when they are offered on a "deal" basis as driven by transaction utility, which depends on the perceived value of the deals. *Schindler* (1992) and *Heath et al.* (1995) provided experimental results supporting that consumers are more likely to purchase at a deal when they are informed that the price reduction is significantly lower than the regular price. These studies, however, did not analytically model consumer deal-proneness when multiple units are purchased, and did not investigate how deal-prone consumers respond to different types of deals. There are a limited number of papers which studied the purchase behavior of deal-prone consumers under different promotion types. *Lichtenstein et al.* (1997) identified a consumer segment that is deal-prone across various types of promotions. *Laroche et al.* (2003) studied deal-prone consumer perception and purchase intention when offered coupons and two-for-one promotions. However, similar to the second stream of literature, the papers in this area are mostly empirical and experimental, aiming to understand consumer behavior rather than the seller's profitability. which is contrastingly different from our analytical approach.

Overall, the main difference that distinguishes our work from the existing literature is that we are the first to compare the profitability and the impact on consumer purchase behavior of all-unit and fixed-amount discount, two of the most commonly used price promotions in retailing. Furthermore, we analytically model consumer deal-proneness and investigate its implications on consumer spending under conditional discounts.

3.3 Model

We consider a seller of one product facing heterogeneous consumers. If the seller does not offer a discount, the product is sold at a retail price of p per unit. To reflect a retail price regulation commonly imposed by a manufacturer, we assume that the retail price p is exogenous to the seller, as is the case for a manufacturer's suggested retail price (MSRP) in practice. However, the seller can adjust the price at which the product is sold using a price promotion (discount).

While there are many different forms of discounts used in retailing, this essay focuses on *conditional discount*, which refers to a discount that is applied only when a consumer satisfies the purchase condition (e.g. minimum purchase quantity, minimum spending). In particular, we examine two widely used forms of conditional discounts: all-unit discount (A) and fixed-amount discount (F). These two forms of discounts capture all types of the price promotions identified as most commonly employed according to *Lichtenstein et al.* (1997) (e.g. buy-one-get-one-free, sales, coupons, cents-off).[1]

3.3.1 Types of Conditional Promotions

All-Unit Discount (A)

In an all-unit discount, a discount is applied to all purchased units if a customer's purchase meets a minimum eligibility requirement (e.g. buy 2 or more and get 25% off). To represent the terms of an all-unit discount, let $r \in [0, 1)$ denote the promotion depth, or the "percent-off," and let K denote the minimum purchase quantity required to obtain the discount.[2] For tractability purpose, we assume that K is a continuous parameter. We acknowledge that in practice, all-unit discounts are generally offered for products sold in discrete quantities. However, our use of continuous quantity provides a good approximation of the discrete quantity model without sacrificing the insights.

Let $D^A = (r, K)$ represent an all-unit discount. Then, the purchase price for q units of product under an all-unit discount D^A is given by

$$P(q, D^A) = \begin{cases} pq & \text{if } 0 \leq q < K \\ p(1-r)q & \text{if } q \geq K \end{cases}$$

Notice that a standard price markdown is a special case of all-unit discounts with K equals to the smallest sellable unit of the product (e.g. 1 shirt, half-order, 1 oz.),

[1] Incremental discount is another well-known form of conditional discounts, primarily used by suppliers or manufacturers. However, given its small presence in retailing, it is not considered in *Lichtenstein et al.* (1997). Hence, we do not consider incremental discount in this chapter.

[2] Notice that an all-unit discount with an eligibility requirement in a form of a minimum spending can be represented in the same way since the retail price is fixed. For example, a promotion "spend \$50 or more and get 25% off" for a product priced at \$25 each has $K = 2$.

so that any purchase is qualified for the discount.

Fixed-Amount Discount (F)

In a fixed-amount discount, the final amount that a consumer has to pay is reduced by a predetermined discount amount if the consumer's purchase meets a minimum eligibility requirement (e.g. buy 2 or more get \$25 off). To represent the terms of a fixed-amount discount, let $m \geq 0$ be the "dollars-off," which is the dollar discount amount to be subtracted from the total price of an eligible purchase; let K be the minimum purchase quantity to qualify for the discount.

Let $D^F = (m, K)$ denote a fixed-amount discount. Then, the purchase price for q units of product under a fixed-amount discount D^F is given by

$$P(q, D^F) = \begin{cases} pq & \text{if } 0 \leq q < K \\ pq - m & \text{if } q \geq K \end{cases}$$

Notice that under a fixed-amount discount, the discount amount that a customer receives for an eligible purchase does not go up with the total purchase quantity. On the other hand, under an all-unit discount, the dollar discount amount is larger for an eligible purchase of a larger quantity since the discount is applied to all purchased units. Note also that no discount is a special case of conditional discounts when the discount is zero ($r = 0$ for an all-unit discount, or $m = 0$ for a fixed-amount discount).

Next, we discuss different types of consumers and their corresponding utility when purchasing the product under a conditional discount.

3.3.2 Consumer's Types and Utility

We assume that consumers are heterogenous in two dimensions: valuation of the product, and deal-proneness. A consumer may have a high ("high-type" h) or low valuation ("low-type" l). A high-type consumer is willing to pay a higher price and consume a larger quantity of the product, compared to a low-type consumer. We denote the proportion of high-type consumers in the market by γ. Consumers may also differ in their responses to deals. That is, some consumers are more inclined to purchase when a deal is present. To reflect this, we assume that there are two types of consumers based on their cognitive behavior towards deals: value-conscious (v) and deal-prone (d). A value-conscious consumer only draws utility from purchasing and consuming the product (acquisition utility). On the other hand, a deal-prone consumer draws additional utility when purchasing the product at a sufficiently large discount (transaction utility). The proportion of deal-prone consumers in the market is denoted by β.

Notice that the two attributes of consumers: valuation and deal-proneness, are assessed based on different aspects of consumer purchase behavior. That is, a consumer's valuation is based solely on her liking of the product; whereas, the consumer's deal-proneness is based on her cognitive response towards a pricing scheme. Hence, we assume that the two attributes are independent. This gives rise to four different consumer segments: high-type deal-prone (hd), high-type value-conscious (hv), low-type deal-prone (ld), and low-type value-conscious (lv), with a proportion of $\gamma\beta$, $\gamma(1-\beta)$, $(1-\gamma)\beta$, and $(1-\gamma)(1-\beta)$, respectively.

A consumer's net utility from making a purchase under a conditional discount is a sum of the acquisition utility and transaction utility, based on the *acquisition-transaction utility theory* proposed by *Thaler* (1985). We define acquisition and transaction utility below.

Acquisition Utility

Following the standard practice, we define acquisition utility as a consumer's valuation less the purchase price. Let $V_i(q), i \in \{h, l\}$ denotes a type-i consumer's valuation (willingness to pay) for q units of the product. We assume that a type-i consumer receives a utility of s_i from consuming each additional unit of the product. However, there is a limit, θ_i, above which consuming additional units will not increase the consumer's utility. This reflects the fact that after a certain quantity, the consumer gains no additional surplus from consuming more. That is, a type-i consumer accrues utility from the first θ_i units at the rate of s_i per unit. Beyond this limit, the marginal utility gained from consuming an additional unit becomes zero. We assume $s_l < s_h$ and $\theta_l \leq \theta_h$ to represent that a high-type consumer has a higher willingness to pay and a larger demand for the product, compared to a low-type consumer. Thus, $V_i(q)$ is given by the following equation and illustrated in Figure 3.1.[3]

$$V_i(q) = \begin{cases} s_i q & \text{if } 0 \leq q \leq \theta_i \\ s_i \theta_i & \text{if } q > \theta_i \end{cases}$$

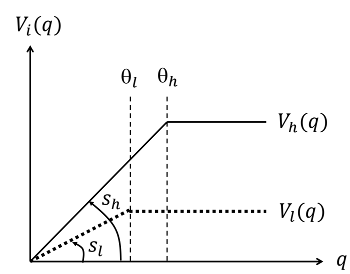

Figure 3.1: Consumer's valuation function

Let $P(q, D^k)$ denotes the purchase price of q units of the product under a conditional discount $D^k, k \in \{A, F\}$, as defined previously. Then, a type-i consumer's acquisition utility from purchasing q units under a conditional discount D^k is given by

$$A_i(q, D^k) := V_i(q) - P(q, D^k).$$

Transaction Utility

In addition to acquisition utility, deal-prone consumers may draw transaction utility from buying the product at a discount. Following empirical evidence that consumers judge the merits of a deal by its promotion depth and dollar savings (*DelVecchio*, 2005; *DelVecchio et al.*, 2007), we assume that deal-prone consumers

[3]We also consider a model with linearly decreasing marginal valuation, which results in a concave utility function, in Section 3.7.2.

draw transaction utility whenever they make a purchase at what they perceive as a "good deal" (measured by a promotion depth or dollar savings). More specifically, for a deal-prone consumer to receive transaction utility, the value of the savings must be greater than a threshold. In an all-unit discount, we assume that a deal-prone consumer receives transaction utility of t if the promotion depth (r) is greater than or equal to a threshold R. If she does not buy enough to qualify for the discount or if the discount does not meet the threshold, she receives zero transaction utility. Likewise, in a fixed-amount discount, a deal-prone consumer receives transaction utility of t if and only if the dollars-off (m) is at least as large as a threshold M. Note however that if a consumer is value-conscious, she does not draw transaction utility from any discount.

Let $T_j(q, D^k), j \in \{v, d\}, k \in \{A, F\}$ denote the transaction utility of a type-j consumer (d for deal-prone and v for value-conscious) when she purchases q units of the product at a conditional discount D^k. Then, the transaction utility $T_j(q, D^A)$ for an all-unit discount $D^A = (r, K)$ is given by:

$$T_d(q, D^A) = \begin{cases} 0 & \text{if } q < K \text{ or } r < R \\ t & \text{if } q \geq K \text{ and } r \geq R \end{cases} \qquad T_v(q, D^A) = 0 \tag{3.1}$$

Similarly, the transaction utility $T_j(q, D^F)$ for a fixed-amount discount $D^F = (m, K)$ is given by equation (3.1) with r and R replaced by m and M, respectively.

We will sometimes refer to t as "cognitive surplus," and R and M as "deal-prone threshold." A large value of t reflects that the deal-prone consumer receives a large additional cognitive gain when completing a good deal. A small value of R and M represents when the deal-prone consumer perceives almost every deal as worthy, and is therefore easily induced to commit to a deal. The value of t, R, and M are product-specific, and may depend on several factors. For example, a deal-prone consumer may have a small t, R, and M for a low price tag product (e.g., bags of chips, yogurt cups) but a larger t, R, and M for a high price tag product (e.g., shoes, shirts, hotel rooms).

Based on the definition of acquisition utility and transaction utility given above, the net utility of a type-ij consumer, $i \in \{l, h\}, j \in \{v, d\}$, who purchases q units of the product at a conditional discount $D^k, k \in \{A, F\}$, is as follows:

$$\begin{aligned} U_{iv}(q, D^k) &= A_i(q, D^k) \\ U_{id}(q, D^k) &= A_i(q, D^k) + T_d(q, D^k) \end{aligned} \tag{3.2}$$

In the next section, we will analyze how different types of consumers respond to deals in the form of all-unit and fixed-amount discounts. In particular, we are interested in comparing the effectiveness of these two types of discounts in boosting the consumer's purchase quantity.

3.4 Consumer's Problem

We first examine how each type of consumers behaves when a conditional discount is offered. More specifically, we are interested in characterizing how the valuation and deal-proneness of a given consumer (defined by type ij) influence her purchase decision whether to buy the product, and if so, how much.

We consider a type-ij consumer's problem of choosing the purchase quantity ($q \geq 0$) that maximizes her utility.[4] When facing a conditional discount $D^k, k \in \{A, F\}$, the consumer's utility from purchasing a quantity q is given by $U_{ij}(q, D^k)$, as in equation (3.2). Notice that acquisition utility $A_i(q, D^k)$ depends on the consumer's valuation type $i \in \{l, h\}$; transaction utility $T_j(q, D^k)$ depends on the consumer's deal-prone type $j \in \{v, d\}$. Both acquisition utility and transaction utility also depend on the discount type $k \in \{A, F\}$. Below, we discuss the consumer's problem of choosing an optimal purchase quantity under an all-unit discount and a fixed-amount discount.

[4] We assume that if the utility from purchasing two different quantities are the same, the consumer always chooses to purchase the larger quantity due to the lower perceived per-unit price.

3.4.1 All-Unit Discount

In an all-unit discount $D^A = (r, K)$, the utility that a type-ij consumer draws when purchasing q units of the product is given by:

$$U_{ij}(q, D^A) = \begin{cases} s_i q - pq & \text{if } 0 \le q < \min\{\theta_i, K\} \\ s_i \theta_i - pq & \text{if } \theta_i \le q < K \\ s_i q - p(1-r)q + T_j(q, D^A) & \text{if } K \le q < \theta_i \\ s_i \theta_i - p(1-r)q + T_j(q, D^A) & \text{if } q \ge \max\{\theta_i, K\} \end{cases}$$

The first two expressions correspond to the consumer's utility when she buys fewer than the minimum requirement of K units and receives no discount. Notice that when $q \ge \theta_i$, the consumer obtains the maximum valuation of $s_i\theta_i$. Analogously, the third and the fourth expressions correspond to the consumer's utility when she buys at least the minimum required quantity and receives the discount. Note also that, depending on the relative size of K to θ_i, the second or the third interval of q may be empty.

3.4.2 Fixed-Amount Discount

In a fixed-amount discount $D^F = (m, K)$, the utility that a type-ij consumer draws when purchasing q units of the product is given by:

$$U_{ij}(q, D^F) = \begin{cases} s_i q - pq & \text{if } 0 \le q < \min\{\theta_i, K\} \\ s_i \theta_i - pq & \text{if } \theta_i \le q < K \\ s_i q - pq + m + T_j(q, D^F) & \text{if } K \le q < \theta_i \\ s_i \theta_i - pq + m + T_j(q, D^F) & \text{if } q \ge \max\{\theta_i, K\} \end{cases}$$

Notice that the only difference between the utility under a fixed-amount discount and that under an all-unit discount is the discount amount. That is, the fixed-amount discount m is independent of q as long as $q \geq K$, but the discount amount received under the all-unit scheme, prq, increases with q.

The optimal purchase decision of a consumer is characterized in Proposition 3.1.

Proposition 3.1. *For a given conditional discount $D^k, k \in \{A, F\}$, with a minimum purchase requirement of K, the optimal purchase decision of a type-ij consumer, $i \in \{l, h\}, j \in \{v, d\}$, is characterized by two switching curves: $\sigma_j(\theta_i, D^k)$ and $\bar{\theta}_j(D^k)$, as follows:*

i) If $s_i < \sigma_j(\theta_i, D^k)$, no purchase is optimal.

ii) If $s_i \geq \sigma_j(\theta_i, D^k)$ and $\theta_i < \bar{\theta}_j(D^k)$, it is optimal to buy quantity $\theta_i < K$ at the full price.

iii) If $s_i \geq \sigma_j(\theta_i, D^k)$ and $\theta_i \geq \bar{\theta}_j(D^k)$, it is optimal to buy either K or θ_i at the discount.

The switching curves $\sigma_j(\theta_i, D^k)$ and $\bar{\theta}_j(D^k)$ are increasing in p and K, and decreasing in the depth of the discount. Furthermore, $\bar{\theta}_d(D^k) \leq \bar{\theta}_v(D^k) \leq K$ and $\sigma_d(\theta_i, D^k) \leq \sigma_v(\theta_i, D^k) \leq p$.

Proposition 3.1 states that a type-ij consumer's optimal purchase quantity under a conditional discount depends on her marginal valuation of the product (s_i) and her maximum consumption (θ_i). If the consumer has low marginal valuation compared to the price, then she will not buy the product (part i) of the proposition). If her marginal valuation is sufficiently high, even when her consumption level is low, she will still buy the product at the full price (part ii) of the proposition). When both her valuation and maximum consumption level are high, she will buy at least the required quantity K and receive the discount. If the terms of a discount become more attractive (smaller K or larger r or m), the consumer's purchase quantity increases as the switching curves decrease.

Figure 3.2 and 3.3 together show how deal-prone consumers behave differently from value-conscious consumers for a given conditional discount D^k.[5] Since deal-prone consumers draw additional utility when they buy at a discount, they are more likely than value-conscious consumers to increase their purchase quantity to meet the requirement for the discount. Consequently, the deal-prone switching curves lie below the value-conscious switching curves, as stated in Proposition 3.1 and illustrated in Figure 3.3. This implies that deal-prone consumers always buy no less than value-conscious consumers who have the same valuation and consumption level.

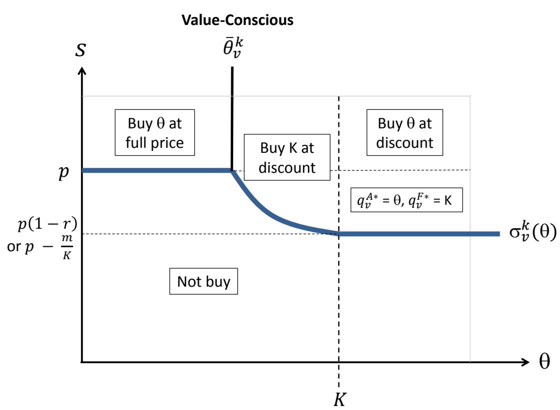

Figure 3.2: Value-conscious purchase quantity under a conditional discount

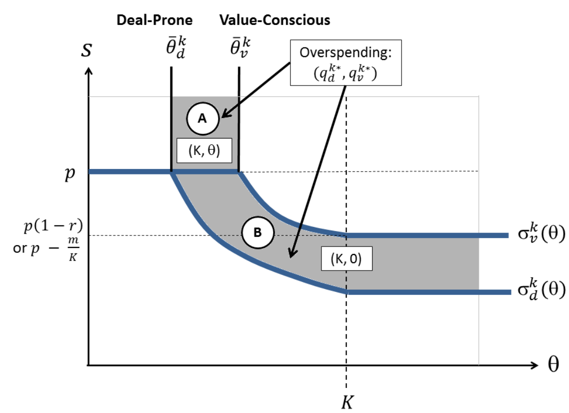

Figure 3.3: Cognitive overspending under a conditional discount

Figure 3.3 highlights the regions (A and B, shaded) where a deal-prone consumer purchases a strictly greater quantity than a value-conscious consumer does. Notice that these regions are bounded by the value-conscious and the deal-prone switching curves. If the consumer valuation falls in Region A (high marginal valuation, low consumption), the value-conscious consumer buys θ at no discount since her maximum consumption level is too far from the minimum requirement K (θ is below the value-

[5]For notational simplicity, we drop the valuation type i and discount D^k from the expressions displayed in the figures.

conscious switching curve). On the other hand, the deal-prone consumer is willing to increase her purchase quantity to K in order to qualify for the discount (θ is above the deal-prone switching curve). In this case, the deal-prone consumer ends up buying much more than what she could consume just to receive the satisfaction (transaction utility) from completing the deal. If the consumer valuation falls in Region B (low marginal valuation, high consumption), the value-conscious consumer does not buy the product since her valuation is too low (s is below the value-conscious switching curve). However, the deal-prone consumer still buys K units to receive the discount (s is above the deal-prone switching curve). In this case, although the deal-prone consumer does not highly value the consumption of the product, she ends up making a purchase anyway due to the transaction utility she receives from the discount.

We can see that when the consumer valuation is confined by the value-conscious and deal-prone switching curves, only deal-prone consumers, not value-conscious consumers, are enticed to purchase more in order to qualify for the conditional discount. This is because in that situation, the acquisition utility from buying at discount is marginally lower than that from buying at no discount (or not buying). Deal-prone consumers are actually better off purchasing a larger quantity to receive the discount since the transaction utility they obtain with the discount is sufficient to increase their overall utility. We define such a situation where the purchase quantity of a deal-prone consumer is strictly greater than the purchase quantity of a value-conscious consumer with the same valuation as "cognitive overspending."

Note that in the other regions outside of Region A and B, both value-conscious and deal-prone consumer behave the same. In the region above the value-conscious switching curves (to the right of Region A and B), both value-conscious and deal-prone consumers purchase at the discount because their valuation falls above their respective switching curves. Likewise, in the region below the deal-prone switching curves (to the left of Region A and B), both value-conscious and deal-prone consumers do not buy at the discount because their valuation falls below their switching curves.

These comparisons between the purchase behavior of value-conscious and deal-prone consumers are summarized in Proposition 3.2.

Proposition 3.2. *For a consumer valuation* (θ, s) *and a conditional discount* $D^k, k \in \{A, F\}$, *with a minimum purchase requirement* K*:*

i) If (θ, s) *falls between the value-conscious and deal-prone switching curves (Region A and B in Figure 3.3), a deal-prone consumer buys* K *units, which is strictly more than what a value-conscious consumer buys (i.e.,* cognitive overspending*).*

ii) In all other cases, both deal-prone and value-conscious consumer behave identically.

Next, we investigate differences between consumer behavior under the all-unit and fixed-amount discount. To rule out the framing effects (e.g. percent-off vs. dollars-off) and make the comparison fair, we compare the all-unit and fixed-amount discount that require the same minimum purchase quantity and offer the same amount of savings when a consumer buys exactly K units, i.e., $D^A = (r, K)$ and $D^F = (m = prK, K)$. When facing these discounts, a consumer who buys K units pays the same amount of $p(1 - r)K$ after receiving the same discount of prK, which triggers the same transaction utility.[6]

Proposition 3.3 discusses the conditions under which the consumer's optimal purchase quantity is the same or different under the all-unit and fixed-amount discount.

Proposition 3.3. *For any consumer's valuation* (θ, s)*, an all-unit discount* $D^A = (r, K)$*, and a fixed-amount discount* $D^F = (m = prK, K)$*:*

i) The all-unit and fixed-amount switching curves are identical. That is, $\sigma_j(\theta, D^A) = \sigma_j(\theta, D^F)$ *and* $\bar{\theta}_j(D^A) = \bar{\theta}_j(D^F)$*, for* $j \in \{v, d\}$*.*

[6]For this, we establish the relationship between the all-unit and fixed-amount deal-prone thresholds, R and M, that $M = pKR$. Hence, $m \geq M$ is equivalent to $r \geq R$.

ii) If $\theta > K$ and $p(1-r) \leq s < p$, then the consumer purchases θ units under the all-unit discount, but K units under the fixed-amount discount. In all other cases, the consumer purchases the same quantity under both discount schemes.

Part i) of the proposition shows that the forms of discounts do not change the region in which a consumer buys at the discount or not. This is because the two discounts set the same condition for the minimum purchase quantity, and the offered discounts trigger the same transaction utility. However, part ii) of the proposition points out that a consumer may in fact purchase different quantities under the two types of discounts in certain situations. More precisely, when a consumer has a high consumption level ($\theta > K$) but moderate willingness to pay ($p(1-r) \leq s < p$), she will buy θ under the all-unit discount, but will buy a smaller quantity of K under the fixed-amount discount. This is because the consumer is willing to pay for the product only at the discounted price, but not the regular price. Under the all-unit discount, she pays the discounted price for every unit. Hence, she is willing to buy as much as her maximum consumption level. On the other hand, under the fixed-amount discount, the consumer essentially has to pay the full price for any units beyond K. Since the full price is too high, the consumer has no incentive to purchase more than what she needs to qualify for the discount.

Proposition 3.1, Proposition 3.2, and Proposition 3.3 altogether fully characterize and compare the consumer purchase behavior under all-unit and fixed-amount discounts. It is worth noting that the effects of deal-proneness on consumer purchase behavior may be either stronger or weaker than the effects of valuation. That is, a deal-prone consumer with a lower valuation may buy a larger or a smaller quantity, compared to what a value-conscious consumer with a higher valuation buys.

3.5 Seller's Problem

We now examine the seller's expected profit when offering a conditional discount D^k, denoted by $\Pi(D^k), k \in \{A, F\}$. Let $\Pi_{ij}(D^k)$ be the seller's expected profit from selling to a type-ij consumer. For instance, $\Pi_{hd}(D^k)$ represents the seller's profit from a high-type deal-prone consumer; likewise, $\Pi_{lv}(D^k)$ represents the seller's profit from a low-type value-conscious consumer. Then, the seller's total expected profit is

$$\begin{aligned} \Pi(D^k) &= \sum_{i\in\{l,h\},j\in\{v,d\}} Pr(i,j)\Pi_{ij}(D^k) \qquad (3.3) \\ &= \beta\gamma\Pi_{hd}(D^k) + \beta(1-\gamma)\Pi_{ld}(D^k) + (1-\beta)\gamma\Pi_{hv}(D^k) + (1-\beta)(1-\gamma)\Pi_{lv}(D^k). \end{aligned}$$

To simplify our analysis, we assume in the base model that the seller's unit cost is normalized to 0 (e.g., the procurement/production cost is sunk, and the seller's objective is to maximize revenues from sales.). But later in Section 3.7.1, we will discuss how to modify our model to reflect when the unit cost is $c > 0$, and show that the presence of a positive unit cost does not change the main insights obtained in this chapter.

For each type of conditional discounts, the seller's problem is to choose the terms of a discount – minimum purchase quantity K, and the discount rate r (all-unit) or m (fixed-amount) – that maximize his expected profit. The seller may also choose to offer no discount by setting $r = 0$ or $m = 0$. While offering a discount can boost the sales volume, the benefit does not come for free. The seller needs to forgo the margin in exchange for the increased sales. This intuition suggests that a conditional discount may not provide additional profits to the seller under all circumstances. Proposition 3.4 identifies exactly when the seller should employ conditional discounts.

Proposition 3.4. [When is it optimal to offer a discount or not?] [7]

i) If there exist deal-prone consumers ($\beta > 0$), no discount is never optimal.

[7] The results in part i) and iia) continue to hold when there are $N > 2$ types of consumer valuation.

ii) If all consumers are value-conscious ($\beta = 0$), then no discount is optimal if and only if a) all consumers are willing to buy the product at the regular price (i.e. $s_l \geq p$) or b) only high-type consumers are willing to buy the product at the regular price and $\gamma \geq \frac{s_l}{p}$.

Part i) of the proposition states that no discount can never be optimal as long as there exist deal-prone consumers in the market. This is because the seller can always set the terms of a discount to offer just enough discount to trigger transaction utility of deal-prone consumers while requiring them to purchase a sufficiently large quantity for the discount to be profitable. Even when all consumers are value-conscious, conditional discounts increase the seller's profit in all cases except when there exist enough consumers in the market who are willing to buy the product at the regular price (part ii)). This is because in that situation, the demand for the product is already high without a discount offer. Hence, it is not worth increasing sales volume by discounting the price.

Proposition 3.4 implies that there are many circumstances where the seller can utilize a conditional discount to extract more profits. The next result discusses which type of discounts the seller should use.

Proposition 3.5. [All-unit discount, fixed-amount discount, and price markdown] [8]

i) When no consumers are willing to buy at the regular price (i.e., $s_h < p$), all-unit discount weakly dominates fixed-amount discount. Furthermore, price markdown is an optimal all-unit discount when $\beta \leq \bar{\beta}$, for some $\bar{\beta} \in [0, 1]$.

ii) In all other cases, fixed-amount discount weakly dominates all-unit discount.

[8] When there are $N > 2$ types of consumer valuation, it continues to hold that all-unit discount weakly dominates fixed-amount discount when no consumers are willing to buy at the regular price, and fixed-amount discount weakly dominates all-unit discount when all consumers are willing to buy at the regular price.

All-unit discount outperforms fixed-amount discount when consumer willingness to pay for the product is sufficiently low that no consumers are willing to buy at the regular price. In this case, only the all-unit discount can induce consumers to buy a quantity strictly greater than the minimum requirement because the price reduction applies to all units purchased. The fixed-amount discount can at most attract consumers to buy exactly the minimum quantity required for the discount because consumers are not willing to pay the full price for any units beyond that.

On the other hand, if there exist some consumers in the market who are already willing to pay for the product at the regular price, then fixed-amount discount is more profitable than all-unit discount. If only the high-type consumers are willing to pay the regular price, the seller's main objective of offering a discount is to attract low-valuation consumers, who originally are not willing to pay the regular price, to buy the product. However, since a discount is offered to all consumers, the high-valuation consumers, who otherwise would buy up to their maximum consumption level at the full price, can also take advantage of the lowered price. Under the all-unit discount, the high-valuation consumers can "free ride" on the discount for every unit they purchase. But under the fixed-amount discount, the maximum discount amount the high-valuation consumers can be awarded is capped. Thus, the seller's margin and total profit are greater with a fixed-amount discount. If all consumers are willing to pay the regular price, notice from Proposition 3.3 part ii) that each type of consumers always purchase the same quantity, either K or θ, under all-unit and fixed-amount discount. If all consumers buy K, then they all receive the same discount of $m = prK$ under both all-unit and fixed-amount discount, and the two discounts result in the same profit to the seller. However, if some consumers buy θ, which is strictly greater than K (This happens when the seller intends to offer the discount to increase the purchase quantity of the low-type only, so $\theta_l < K < \theta_h$.), then the seller has to award them a larger amount of discount under the all-unit scheme. Hence, the fixed-amount scheme is more profitable.

Proposition 3.5 also shows that price markdown is optimal when most consumers in the market are value-conscious. This is because the benefit from inducing deal-prone consumers to overspend is not significant enough to boost the seller's profit when there are not enough deal-prone consumers in the market. In this case, it suffices to use a stand price markdown to increase profits from both value-conscious and deal-prone consumers. The seller's optimal discount schemes discussed in Proposition 3.4 and Proposition 3.5 are summarized in Figure 3.4a to c.

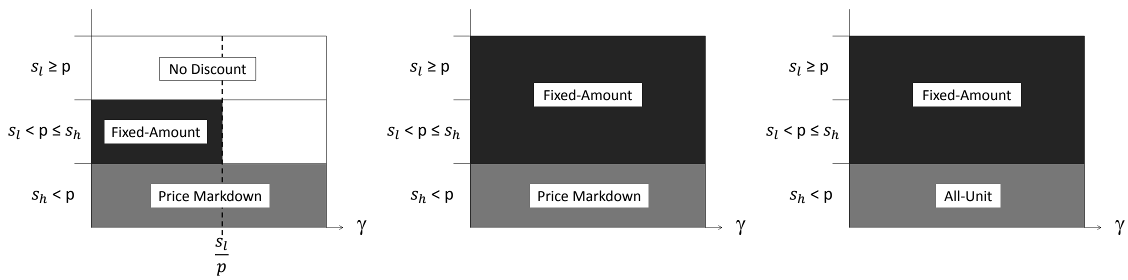

a. $\beta = 0$ b. $0 < \beta \leq \bar{\beta}$ c. $\beta > \bar{\beta}$

Figure 3.4: Seller's optimal discount schemes

In addition to identifying the seller's optimal discount scheme, we discuss when a conditional discount is increasingly more profitable than no discount and price markdown in Proposition 3.6.

Proposition 3.6. *An optimal conditional discount is increasingly more profitable than no discount and price markdown when either more consumers are deal-prone (β increases), or deal-prone consumers have a larger cognitive surplus (t increases).*

A conditional discount is especially beneficial when the market is more deal-prone, signified by either a larger proportion of deal-prone consumers, or a larger degree of responsiveness to deals of deal-prone consumers. This is because conditional discounts can effectively induce deal-prone consumers to overspend.

So far, we have characterized the seller's optimal discount policies when selling to consumers who are heterogenous in both valuation and deal-proneness. To understand how each dimension of consumer heterogeneity affects the seller's discount policies, we analyze two special cases where consumer valuation and deal-proneness are considered in isolation.

3.5.1 Deal-Prone Market with Heterogeneous Valuation

In order to isolate the effects of the heterogeneity in consumer valuation, we consider a special case where all consumers are deal-prone ($\beta = 1$) but they are heterogeneous in valuation. That is, the market consists of two types of consumers: high-valuation deal-prone, and low-valuation deal-prone.[9] Hence, the seller's profit in (3.3) reduces to

$$\Pi(D^k) = \gamma\Pi_{hd}(D^k) + (1-\gamma)\Pi_{ld}(D^k).$$

When choosing the optimal discount terms, the seller needs to decide either to offer a conservative discount, to increase purchase quantity of only the high-type consumers, or to offer a more aggressive discount to increase purchase quantity of the low-type consumers. Lemma 3.1 describes that it is optimal to offer a more generous discount (larger r and/or smaller K) to generate more sales from the low-type consumers only when there is a sufficiently large proportion of them.

Lemma 3.1. *For a conditional discount of type* $k \in \{A, F\}$, *there exists a threshold* $\Gamma^k \in [0, 1]$ *such that offering a discount targeted to only high-type consumers is optimal for* $\gamma > \Gamma^k$. *Otherwise, offering a deeper discount targeted to low-type consumers is optimal.*

We now examine how the threshold Γ^k changes with respect to the magnitude of transaction utility, t. Let $\Gamma^k(t)$ denote the switching curve that characterizes,

[9] The situation where all consumers are value-conscious is a special case of a deal-prone market with $t = 0$.

for a given t, at which γ the seller should target the discount at which segment of consumers. An example of $\Gamma^k(t)$ is shown in Figure 3.5, where the shaded region denotes when the proportion of the high-type consumers is small ($\gamma \leq \Gamma(t)$) and hence it is optimal to target the discount at the low-type segment. Notice that since all consumers are deal-prone, the seller should always offer some form of discounts to increase the purchase quantity of at least one type of consumers, as previously discussed in Proposition 3.4 part i). However, this does not mean that at least one type of consumers should always be induced to overspend (i.e., enticed by transaction utility to purchase more than what a value-conscious consumer would do). To induce cognitive overspending, the seller needs to offer a deep enough discount and set a sufficiently large minimum purchase quantity, which may or may not improve the overall profit. Our next result identifies when cognitive overspending is indeed optimal for the seller.

Proposition 3.7. *When consumers are deal-prone but different in valuation, there exists a continuous switching curve $\Gamma^A(t)$ such that for a given t:*

i) If $\gamma > \Gamma^A(t)$, then $r^* \geq R$ and $K* \geq \theta_h$ (only high-type consumers overspend).

ii) If $\gamma \leq \Gamma^A(t)$, then there exists a threshold $\hat{t}$ such that $r^* \geq R$ and $K^* \geq \theta_l$ (at least low-type consumers overspend) for $t > \hat{t}$.

The same results hold for the optimal fixed-amount discount $D^{F*} = (m^*, K^*)$ when replacing $\Gamma^A(t)$ with $\Gamma^F(t)$, and $r^* \geq R$ with $m^* \geq M$.

Proposition 3.7 implies that it is optimal for the seller to trigger transaction utility and induce cognitive overspending (evident by $r^* \geq R$ or $m^* \geq M$) when either a) there is a sufficiently large proportion of consumers with high valuation (large γ), or b) the magnitude of transaction utility is sufficiently large (large t). Under these conditions, there exist enough consumers who can be induced by transaction utility to significantly increase their purchase quantity due to either a large consumption level (condition a)) or large transaction utility (condition b)). Hence, the seller can improve profit by offering a deep discount while requiring a large minimum purchase quantity to induce cognitive overspending.

We note, however, that when t and γ are both small, it is not always optimal to induce cognitive overspending (i.e., it is possible to have $r^* < R$ or $m^* < M$). Such a situation where cognitive overspending is not optimal is illustrated by the darker region in Figure 3.5. When the cognitive surplus is too small, the consumers are not willing to overspend by much when they receive transaction utility. Hence, it may not be profitable for the seller to offer a deep discount that triggers transaction utility in exchange for a small increase in sales. In particular, when the proportion of neither type of consumers is sufficiently large (moderate γ), it is not profitable for the seller to induce either type of consumers to overspend since they do not have enough mass to generate much larger sales. Notice however that as the cognitive surplus increases, both types of consumers are willing to overspend more when their transaction utility is triggered, making cognitive overspending more profitable for the seller. Hence, the no-overspending region shrinks, and finally disappears.

The exact behavior of the switching curve with respect to t is rather complicated as revealed in Figure 3.5. (The closed-form expressions of the switching curve under each type of conditional discounts are provided in Appendix B.) As a result, for a given γ, the terms of the optimal discount may not be monotone in t. For example, consider $\gamma = 0.55$ in Figure 3.5. Notice that in this example, both types of consumers do not buy at no discount since $s_l < s_h < p$. When t is small (point A), it is not optimal to induce the consumers to overspend. Instead, it is most profitable for the seller to offer a modest discount ($r^* = 0.4 < R$) and require a small purchase quantity ($K^* \leq 2$), just enough to get both the low-type and high-type consumers to purchase their corresponding maximum consumption level. As t increases to a medium value (point B), it becomes profitable to induce the high-type consumers to overspend since they are willing to increase their purchase quantity far beyond their maximum consumption level ($K^* = 7, \theta_h = 4$) when receiving transaction utility ($r^* = 0.5 = R$). It is however still not profitable to induce the low-type consumers to overspend since their consumption level is much lower ($\theta_l = 2$). When t is sufficiently large (point C), even the low-type consumers are willing to increase their purchase quantity by a lot due to large transaction utility. Given a significant presence of the low-type consumers in the market, it is optimal for the seller to set a smaller minimum purchase quantity ($K^* = 4.65$) for the discount, so that the low-type consumers will also overspend.

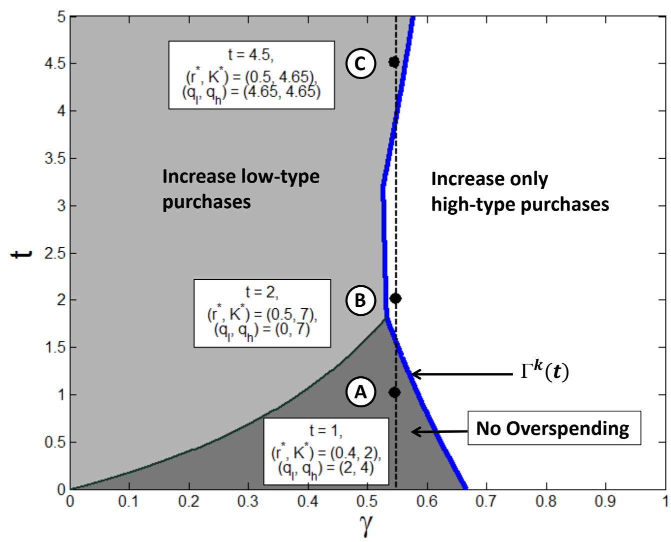

Figure 3.5: Switching curve $\Gamma^A(t)$: $s_l = 2.4, s_h = 3, \theta_l = 2, \theta_h = 4, p = 4, R = 0.5$

Our results for this special case show that it is not always optimal for the seller to induce deal-prone consumers to overspend even when all consumers are deal-prone. Next, we investigate the effects of consumer heterogeneity in deal-proneness on the optimal discount policies.

3.5.2 Homogeneous Valuation with Heterogeneous Deal-Proneness

To isolate the effects of the heterogeneity in consumers' attitude towards a deal, we consider a special case where all consumers have the same valuation, characterized by the same s and θ, but they are either deal-prone or value-conscious. For a given proportion of deal-prone consumers β, the seller's profit in (3.3) can be expressed by

$$\Pi(D^k) = \beta\Pi_d(D^k) + (1-\beta)\Pi_v(D^k).$$

The seller's optimal strategy in targeting the discount at either the deal-prone or value-conscious segment of the market is characterized in Lemma 3.2.

Lemma 3.2. *For a conditional discount of type $k \in \{A, F\}$:*

i) There exists a threshold $\bar{\beta}^k \in [0,1)$ such that offering a discount to increase the purchase quantity of only deal-prone consumers (cognitive overspending) is optimal for $\beta > \bar{\beta}^k$. Otherwise, offering a discount to increase the purchase quantity of all consumers is optimal.

ii) If consumers are willing to buy at the regular price ($s \geq p$), then $\bar{\beta}^k = 0$. Otherwise, $\bar{\beta}^k = \frac{s\theta}{t+s\theta}$.

Lemma 3.2 part i) implies that it is optimal to induce cognitive overspending only when there are enough deal-prone consumers in the market. Notice that in order to induce cognitive overspending, the seller needs to offer a sufficiently deep discount to trigger transaction utility. In exchange, the seller will set a large minimum purchase quantity so that the deal-prone consumers end up purchasing a larger quantity than the value-conscious consumers do. Such a discount does not attract the value-conscious consumers to buy more. Hence, cognitive overspending is not profitable when the proportion of deal-prone consumers is small.

Part ii) of the lemma reveals that if the consumers are willing to buy at the regular price, then cognitive overspending is always optimal (evident by $\bar{\beta}^k = 0$). In this case, since the value-conscious consumers are already willing to buy up to their maximum consumption at the full price, it is in fact optimal to not offer them any discount. If the consumers are not willing to pay the regular price, the threshold $\bar{\beta}^k$ is given by a function that is decreasing in t. This implies that as deal-prone consumers receive larger transaction utility from a deep discount, the proportion of deal-prone consumers that is required for cognitive overspending to be profitable for the seller gets smaller because they are willing to overspend by a larger amount.

Next, we investigate whether the all-unit discount or fixed-amount discount is more profitable. Interestingly, as Proposition 3.8 reveals, the optimal all-unit and fixed-amount discount always induce the same consumer purchase behavior and yield the same profit to the seller.

Proposition 3.8. *When consumers have the same valuation but are different in their deal-proneness, the optimal all-unit discount and fixed-amount discount always result in the same consumer purchase quantities and the same seller's profit.*[10]

To understand this result, consider the following two cases of a conditional discount: a) cognitive overspending is optimal, and b) cognitive overspending is not optimal. In case a), we learn from Proposition 3.2 that the deal-prone consumers buy K while the value-conscious consumers buy less than K. Since no consumers buy more than K, the optimal all-unit and fixed-amount discounts $(D^A = (r, K), D^F = (m = prK, K))$ always induce the same purchase quantity, give the same discounts to the consumers, and yield the same seller's profit. In case b), we know from Lemma 3.2 part ii) that it is only possible when $s < p$. Hence, from Proposition 3.5, the all-unit discount weakly dominates the fixed-amount discount. Notice that the only situation where the all-unit discount $D^A = (r, K)$ can be strictly more profitable than the fixed-amount discount is when both types of consumers buy $\theta > K$. However, in this case, there is always a fixed-amount discount $D^F = (m = pr\theta, \theta)$ which induces both types of consumers to buy the same quantity θ, and results in the same profit of $p(1 - r)\theta$.

[10]This result continues to hold when t is a random variable, uniformly distributed over a finite interval, e.g., $t \sim U[0, \bar{t}]$.

The results from the two special cases (Section 3.5.1 and 3.5.2) show that it is not always optimal for the seller to induce deal-prone consumers to overspend. In a market where all consumers are deal-prone, cognitive overspending can be optimal when the magnitude of transaction utility (t) is large or the proportion of the high-type (γ) is large. In a market where consumers have the same valuation for the product, cognitive overspending can be optimal when the consumers have high valuation (s) or the proportion of the deal-prone consumers (β) is large. While it is intuitive to find that cognitive overspending is likely to be profitable when the market is highly deal-prone (large t or large β), it is quite surprising to find that the seller is also likely to benefit from cognitive overspending when the market has high valuation (large s or large γ). Naturally, one would think that offering a discount to increase consumer purchase quantities is only beneficial when the consumers have low valuation. We show, however, that when some consumers are deal-prone, the seller can extract even more profit when the consumers have high valuation by using a conditional discount to induce the deal-prone consumers to overspend.

Regarding the types of discounts, we find that the all-unit and fixed-amount discount may perform differently only in presence of heterogeneity in consumer valuation. If consumers have the same valuation, even when they are heterogenous in deal-proneness, both all-unit and fixed-amount discount are equally profitable. This implies that the fundamental differences between the two mechanisms of conditional discounts are the different effects they have on consumers with different valuations of the product.

Our next interest is to get a sense of how much profit improvement can be generated by implementing conditional discounts under different retailing scenarios. We employ numerical study to address this in the next section.

3.6 Numerical Study

We conduct two sets of numerical experiments to address a few important managerial questions regarding the use of conditional discounts in practice. More specifically, we are interested in answering the following questions: 1) By how much can a seller improve profits with a conditional discount? 2) What is the profit difference between using all-unit and fixed-amount discounts? and 3) What factors affect the magnitude of profit improvement from offering a conditional discount?

3.6.1 Profit Improvement

In the first numerical study, we compare the performance of different types of discounts under a large number of different retailing scenarios. We generate 1,296 different problem instances by varying each parameter, as summarized in Table 3.1.[11]

θ_l	θ_h	s_l	s_h	p	t	R	γ	β
1	{2, 3, 4}	1	2	{0.5, 1.5, 2.5}	{0.1, 0.3, 0.5, 0.75}	{0, 0.25, 0.5, 0.75}	{0.2, 0.5, 0.8}	{0.2, 0.5, 0.8}

Table 3.1: Problem parameters for the numerical study of profit improvement

For each problem instance, we solve for the optimal price markdown, all-unit, and fixed-amount discount, and compare the seller's profit under different discount schemes. Table 3.2 summarizes how much profit (in percentage) the seller can gain by offering an optimal price markdown, all-unit, and fixed-amount discount, over no discount. We note however that the instances where the no-discount profit is zero ($p = 2.5$) are excluded from the statistics presented in the table as the profit improvement in those cases is infinite. Since only the instances where at least some consumers are willing to pay the regular price are considered, we observe that the profit improvement from offering the fixed-amount discount is greater than that from offering the all-unit discount, supporting the result in Proposition 3.5. Overall, Table 3.2 provides some

[11] We normalize θ_l and s_l to 1, and s_h to 2 since similar effects of changing these parameters can be observed by changing θ_h and p.

evidence to support that in many cases, offering a discount can significantly increase the seller's profit. In particular, when using a conditional discount, the seller can expect to see a profit improvement of as much as 20% on average, which is about 10% higher than the profit improvement obtained from a standard price markdown.

Next, we compare the performance of the all-unit and fixed-amount discount. Out of 1,296 instances, we find that the two discounts perform equally well in 688 (53.08%) instances; all-unit discount performs better in 297 (22.92%) instances; and fixed-amount discount performs better in 311 (24%) instances. This reveals that about half of the time, the seller can employ either form of conditional discounts to increase profits. However, the other half of the time, one discount scheme can perform better than the other, calling the seller's attention to choosing the appropriate type of conditional discounts for the market he is facing. In fact, the profit differences between offering the two discount schemes can be significant, as reported in Table 3.3. (Profit improvement of all-unit over fixed-amount discount refers to $\frac{\text{profit from all-unit - profit from fixed-amount}}{\text{profit from fixed-amount}}$.) In approximately 40% of the instances where the all-unit discount performs better than the fixed-amount discount, we find that the profit improvement is at least 10%. Likewise, an analogous analysis on the profit improvement of the fixed-amount over all-unit discount shows a similar result that when the fixed-amount discount performs better than the all-unit discount, the profit improvement is at least 10% in about 40% of the instances.

Profit Improvement when Using	Statistics				
	Mean	Median	Standard Deviation	Minimum	Maximum
Price Markdown	10.22%	0.00%	25.45%	0.00%	98.00%
All-Unit Discount	19.37%	7.48%	27.99%	0.00%	175.76%
Fixed-Amount Discount	23.94%	10.66%	32.49%	0.00%	179.34%

Table 3.2: Statistics for the seller's profit improvement when using conditional discounts over no discount

Profit improvement of All-unit over Fixed-amount Discounts							
$< -10\%$	$-10 - 0\%$	$0 - 10\%$	$10 - 20\%$	$20 - 30\%$	$30 - 40\%$	$40 - 50\%$	$> 50\%$
106	893	177	53	49	9	7	2

Table 3.3: Profit difference between all-unit and fixed-amount discounts

3.6.2 Effects of Problem Parameters

In the second numerical study, we investigate how each parameter affects the performance of conditional discounts, compared to other discount schemes. We systematically increase the value of each parameter at steps of 0.01, one at a time, from the three base cases ($p > s_h$, $s_l < p \leq s_h$, and $p \leq s_l$). This results in a total of 1,400 problem instances, as summarized in Table 3.4. In each problem, we solve for the optimal discount policies and compare the seller's profits. Table 3.5 presents the results.

Scenario	θ_l	θ_h	s_l	s_h	p	t	R	γ	β
Base cases	1	2	1	2	{0.5, 1.5, 2.5}	0.2	0.2	0.2	0.2
Increase θ_h	1	[2.01, 3]	1	2	{0.5, 1.5, 2.5}	0.2	0.2	0.2	0.2
Increase t	1	2	1	2	{0.5, 1.5, 2.5}	[0.21, 1]	0.2	0.2	0.2
Increase R	1	2	1	2	{0.5, 1.5, 2.5}	0.2	[0.21, 0.8]	0.2	0.2
Increase γ	1	2	1	2	{0.5, 1.5, 2.5}	0.2	0.2	[0.21, 1]	0.2
Increase β	1	2	1	2	{0.5, 1.5, 2.5}	0.2	0.2	2	[0.21, 1]
Increase p	1	2	1	2	[0.51, 2.5]	0.2	0.2	2	0.2

Table 3.4: Problem parameters for the numerical study of effects of parameters on the profit improvement

Scenario	Optimal Policy Vs. No Discount [a]			Optimal Policy Vs. Price Markdown			All-unit Vs. Fixed-Amount		
	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.
Base cases	58.74%	58.74%	81.19%	3.50%	1.33%	4.95%	9.26%	7.79%	10.08%
Increase θ_h	47.38%	1.33%	47.14%	4.53%	1.23%	5.66%	11.43%	10.50%	10.08%
Increase t	65.96%	73.08%	50.76%	8.31%	9.17%	8.33%	9.26%	7.79%	8.25%
Increase R	58.57%	57.77%	57.48%	3.44%	1.33%	3.98%	9.27%	7.79%	7.79%
Increase γ	4.17%	3.01%	11.66%	3.52%	1.72%	4.43%	2.97%	0.07%	5.28%
Increase β	65.46%	68.45%	55.81%	8.33%	9.17%	6.57%	9.36%	7.79%	8.42%
Increase p	81.52%	91.70%	63.87%	4.74%	1.27%	5.38%	8.67%	7.85%	7.50%

Table 3.5: Statistics for profit improvement with respect to changes in problem parameters

[a] The statistics for the comparisons with no discount does not include the case of $p > s_h$ since no discount profit is zero.

As in Table 3.2, the profit improvement of the optimal policy vs. no discount and the profit improvement of the optimal policy vs. price markdown in Table 3.5

measure the gains in profits from using the better one of the all-unit and fixed-amount discount. The profit improvement of the all-unit vs. fixed-amount discount measures the profit difference between using the two types of conditional discounts. From the table, we observe that an increase in θ_h, γ, and R lead to a smaller profit improvement; while, an increase in p, t, and β lead to a larger profit improvement of the optimal conditional discount over no discount and price markdown. This is because as θ_h and γ increase, consumers have larger consumption and higher willingness to pay for the product. These characteristics of consumers contribute to a larger sales volume in the market, making it less necessary for the seller to employ a conditional discount to boost sales. Analogously, when the regular price of the product increases, the sales volume decreases. In this case, offering a discount is needed in order to generate more sales, and requiring a minimum purchase quantity for the discount helps improve sales even further. The profit improvement obtained from the use of conditional discounts is also strengthened when the market is more deal-prone, caused by a decrease in R, an increase in t, or an increase in β, following the results we have discussed in Proposition 3.6.

The profit difference between all-unit and fixed-amount discount is shown to be most affected by a change in θ_h and γ. More precisely, the profit difference increases with θ_h but decreases with γ. To understand this result, recall from the discussion of Proposition 3.5 that the all-unit discount is more profitable than the fixed-amount discount when consumers buy strictly more under the all-unit discount (θ) than under the fixed-amount discount (K, where $K < \theta$). Hence, as θ_h increases, the profit difference becomes larger. On the other hand, the fixed-amount discount is more profitable than the all-unit discount when the high-type consumer buys their maximum consumption level (θ_h, where $\theta_h > K$) but receives a larger discount amount under the all-unit discount ($pr\theta_h$ under all-unit; $m = prK$ under fixed-amount). In this case, the profit difference becomes larger when the discount rate and θ_h are larger.

Notice that the seller is likely to offer a larger discount rate when there are more low-type consumers in the market (γ small). This explains why the profit difference decreases with γ.

3.7 Extension

3.7.1 Positive Unit Cost

In our base model, we assume that the unit cost is normalized to 0. Here, we modify the model to reflect when the seller incurs a unit cost of c for each unit sold, where $0 < c < p$. Notice that a positive unit cost has no effects on the consumer's problem since consumers do not observe the cost. Hence, all results regarding the consumer purchase behavior (Section 3.4) continue to hold.

For the seller's problem, the presence of a positive unit cost affects the seller's profit as follows. The seller's profit from selling q units under an all-unit discount $D^A = (r, K)$ is given by $(p(1-r)-c)q$ if $q \geq K$, and $(p-c)q$ if $q < K$. Likewise, the seller's profit from selling q units under a fixed-amount discount $D^F = (m, K)$ is given by $(p-c)q - m$ if $q \geq K$, and $(p-c)q$ if $q < K$. Hence, with a positive unit cost, we find that the region where no discount is optimal becomes larger, and the region where cognitive overspending is optimal becomes smaller. Consequently, the following results need to be modified to reflect a positive unit cost. Proposition 3.4 part i), which states that no discount is not optimal as long as there exist deal-prone consumers, will hold only when the unit cost is not too large. (We can prove that there exists a threshold $\bar{c}$ such that the result holds as long as $c < \bar{c}$.) Condition b) of Proposition 3.4 part ii) will be changed from $s_l < p \leq s_h$ and $\gamma \geq \frac{s_l}{p}$ to $s_l < p \leq s_h$ and $\gamma \geq \frac{s_l - c}{p - c}$. Notice that the threshold on γ is smaller with a positive unit cost than with no unit cost ($\frac{s_l - c}{p - c} < \frac{s_l}{p}$), implying that the region where no discount is optimal is larger. Likewise, the results for the two special cases (Section 3.5.1 and 3.5.2) will be affected in the same way that no discount is more likely to be optimal and cognitive overspending is less likely to be optimal. Other than these, all other results continue to hold in presence of a positive unit cost.

3.7.2 Concave Consumer Valuation

In our base model, we assume that a consumer type $i \in \{l, h\}$ has a constant marginal valuation of s_i for the first θ_i units of consumption. Here, we consider an alternate form of the consumer valuation in which the marginal valuation is linearly decreasing in the units consumed, as adopted in *Chuang and Sirbu* (1999). More precisely, let v_0 be the marginal valuation of the first unit of goods. Then, for a consumer type i, the marginal valuation of the q^{th} unit of goods is given by

$$v_i(q) = \begin{cases} v_0(1 - \frac{q}{\theta_i}) & \text{if } 0 \leq q \leq \theta_i \\ 0 & \text{if } q > \theta_i \end{cases}$$

This implies that a consumer valuation is a concave increasing function of q:

$$V_i(q) = \int_0^q v_i(x)dx = \begin{cases} v_0 q(1 - \frac{q}{2\theta_i}) & \text{if } 0 \leq q \leq \theta_i \\ \frac{v_0\theta_i}{2} & \text{if } q > \theta_i \end{cases}$$

Solving the consumer's problem with this acquisition utility, we obtain similar results as in Proposition 3.1. That is, the consumer's optimal purchase quantity under a conditional discount $D^A = (r, K^A)$ and $D^F = (m, K^F)$ is characterized by two thresholds: $\underline{\theta}_j(v_0, D^k)$ and $\bar{\theta}_j(v_0, D^k), k \in \{A, F\}$, such that a type-$ij$ consumer buys a quantity smaller than K^k at no discount if $\theta_i < \underline{\theta}_j(v_0, D^k)$, buys exactly K^k at discount if $\underline{\theta}_j(v_0, D^k) \leq \theta_i < \bar{\theta}_j(v_0, D^k)$, and buys more than K^k at discount if $\theta_i \geq \bar{\theta}_j(v_0, D^k)$, $i \in \{l, h\}, j \in \{v, d\}$. Furthermore, if $K^A = K^F$ and $m = prK$, then $\underline{\theta}_j(v_0, D^A) = \underline{\theta}_j(v_0, D^F)$ and $\bar{\theta}_j(v_0, D^A) \leq \bar{\theta}_j(v_0, D^F)$. Figure 3.6a and 3.6b graphically illustrate the consumer's optimal purchase quantity, $q_j(\theta_i)$, under the all-unit and fixed-amount discount.

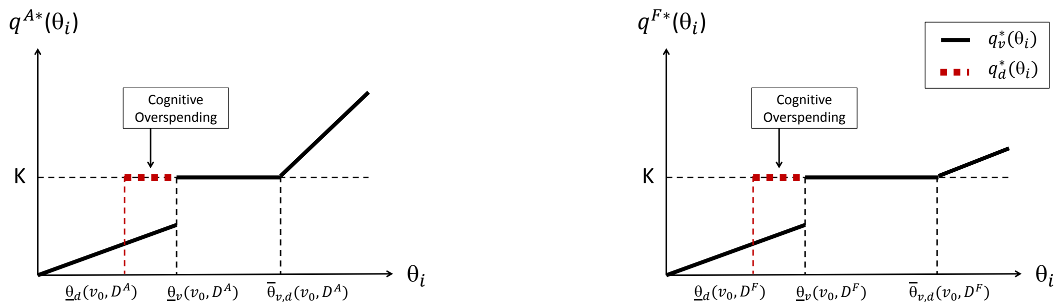

a. b.

Figure 3.6: Optimal purchase quantity under a conditional discount with a concave valuation: a. all-unit, b. fixed-amount

We also find that there exists a region where a deal-prone consumer overspends (buys strictly more than what a value-conscious consumer does) under the all-unit and fixed-amount discount, analogous to Proposition 3.2. Likewise, there exists a region ($\theta > \bar{\theta}_j(v_0, D^A)$) where a consumer purchases strictly more under the all-unit discount than under the fixed-amount discount, similar to Proposition 3.3.

Given that the consumer's problem with concave valuation shares similar characteristics as that in the base model with linear valuation, we can infer similar results for the seller's problem. First, for $D^A = (r, K)$ and $D^F = (prK, K)$, we have that the all-unit discount weakly dominates the fixed-amount discount when consumers valuation (v_0) is sufficiently low, and the fixed-amount discount weakly dominates the all-unit discount otherwise. Furthermore, when all consumers have the same valuation but are different in their deal-proneness, we continue to have that the all-unit and fixed-amount discount yield the same seller's profits.

3.7.3 Endogenous Price

Previously, we have assumed that the retail price p is exogenously given. Here, we extend our model to consider the situation where the seller can optimize the retail price in addition to the discount term. Hence, the seller's problem under each discount scheme involves three decision variables: the regular price p, the discount rate r for all-unit or discount amount m for fixed-amount discount, and the minimum purchase quantity K. Given large degrees of freedom in the discount terms, there can be multiple optimal prices for a given problem. However, we can show for each discount scheme that an optimal price is either $p = s_l$ (all consumers are willing to pay the regular price), $p = s_h$ (only high-type consumers are willing to pay the regular price), or $p = \frac{s_h\theta_h+t}{\theta_h(1-R)}$ (no consumers are willing to pay the regular price).[12]

Let p^A and p^F denote the optimal price under the all-unit and fixed-amount discount, respectively. If $p^A = p^F$, then all results in the base model with exogenous price follow immediately. Otherwise, Proposition 3.5 implies that the all-unit discount can be better than the fixed-amount discount only when $p^A = \frac{s_h\theta_h+t}{\theta_h(1-R)}$. Likewise, the fixed-amount discount can be better than the all-unit discount only when p^F is s_l or s_h. Furthermore, we can show that the result in Proposition 3.5 part i), which states that price markdown is an optimal all-unit discount for $\beta < \bar{\beta}$ still holds, and we continue to have that the optimal all-unit and fixed-amount discount result in identical outcome when consumers have homogeneous valuation but different levels of deal-proneness, as in Proposition 3.8.

3.8 Conclusion

Motivated by consumer behavior in retailing, this chapter discusses how price promotions influence purchase decisions of different types of consumers, and which type of promotions is most profitable to the seller under which market situations. We consider a market where consumers can be heterogeneous in two dimensions: willingness to pay for the product and the deal-proneness to a discount offer, and focus on two popular types of conditional promotions: all-unit discount and fixed-amount discount. Our study shows that deal-prone consumers may be induced to overspend

[12]See Appendix B for the proof.

when offered a conditional discount. As a result, retailers can employ conditional discounts to effectively improve their profitability, especially when facing a market of consumers with low willingness to pay and high deal-proneness. We find, however, that it is not always optimal for the seller to induce consumer overspending. Only when there is a sufficiently large proportion of highly deal-prone or high-valuation consumers in the market that inducing overspending can be profitable.

Another key finding is that, depending on the market, one conditional discount scheme can perform better than the other. We show that when consumers are not willing to pay the regular price for the product, the all-unit discount outperforms the fixed-amount discount because only the all-unit discount can induce consumers to buy strictly more than the minimum quantity required for the discount. On the other hand, when some consumers are already willing to pay the regular price, the fixed-amount discount is more profitable than the all-unit discount. In this case, the fixed-amount discount awards only a limited discount amount, less than that awarded under the all-unit scheme, to the consumers who otherwise would voluntarily pay the regular price. Hence, it allows the seller to maintain a greater profit margin. Based on these findings, an important implication is that an all-unit discount should be adopted to stimulate sales of a high price tag product or a newly-launched item; while, a fixed-amount discount is more effective as a frequent promotion of a low price tag or established brand-named product. Our key results are shown robust to changes in modeling assumptions, suggesting that the insights from this chapter apply to a broad range of realistic retailing situations.

There are a few aspects of price promotions and consumer psychology relevant to our problem that are not included in the current study, but can serve as interesting future research directions. We will discuss two major areas here. The first area involves post-promotion effects. As reported in several studies, price promotions may lead consumers to experience sticker shock due to the downward price expectation they develop after seeing the product sold at a cheaper price during the promotion period (*DelVecchio et al.*, 2007). Such effect may be more or less intense under different discount schemes, and may pose a limitation on the discount rate that should be offered. These factors could affect the profitability of conditional discounts, which may lead to different results from this essay. The second area is framing and promotional vehicles. When consumers make a purchase decision in an actual shopping scenarios, they are sometimes influenced by how a promotion is framed (e.g., dollars-off, percent-off, Buy-X-Get-Y-Free, Buy-X-for-Y) and how it is awarded (e.g., coupons, straight off the shelf) in addition to the discount term. For example, *Chen et al.* (1998) reported that for high-price products, a price reduction framed in dollar terms was perceived more significant than the same price reduction framed in percentage terms. *Dhar and Hoch* (1996) found that coupons lead to higher sales compared to straight-off-the-shelf price discounts. Taking consumer psychological response to different forms of promotions into account, the seller's joint decision on the discount term and its vehicle can be complicated, yet meaningful to study.

CHAPTER 4

Dynamic Pricing or Dynamic Logistics?

4.1 Introduction

An important operational problem faced by retailers is how to most effectively match inventory to customer demand over a selling horizon. This problem becomes especially challenging for retailers who sell their products through multiple channels or stores. While multi-channel retailing is found to be associated with enhanced customer loyalty and sales growth, it requires extra efforts in coordinating channel strategies to efficiently serve customers across channels (*Neslin et al.*, 2006). If not managed well, a multi-channel environment could result in a supply-demand mismatch, and hurt the retailer. For example, Best Buy reported that 2 to 4 percent of its online traffic did not result in a purchase because the inventory was out of stock at distribution centers, and hence, was shown on the website as unavailable. The company estimated, however, that in 80 percent of those cases, the products were in fact available at one of its stores. The Best Buy CEO Hubert Joly admitted that this lost sales due to the stockouts at their online channel represented a very large number (*Schinkel*, Jun. 25, 2013; *Blair*, Jul. 1, 2013).

To efficiently balance inventory and demand in a multi-channel environment, retailers may deploy various tools from marketing and supply chain management. The most common marketing tools involve some form of pricing and promotions, aiming to shape demand to match with available inventory at each channel. These are often implemented through channel-specific discounts or promotions such as online-only sale events at Forever21, and in-store-only coupons at Staples. In fact, a number of major retailers like JC Penney, Kohl's, and Walmart, explicitly state under their pricing policies that they do not match the prices of an item sold in-store and online (*DeNicola*, May 13, 2013). JC Penney explains that it may offer a clearance discount on an online item if it does not sell as quickly as it does it store (*Brownell*, May 31, 2013); whereas, Walmart attributes this pricing policy to differences in distribution, regional competition and other factors (Walmart.com).

Alternatively, retailers can better manage and allocate their inventory across multiple channels in pace of sales. For example, retailers like Macy's, Nordstrom, and Toys 'R' Us leverage their store inventory to fulfill online orders through the practice known as "ship-from-store distribution" (*Lynch*, Jul. 18, 2013). A Canadian clothing company, Roots, utilizes inventory for online orders to fulfill an in-store purchase by offering to ship items, which are out of stock in stores, to an in-store customer's house (*Financial-Post*, Sep. 24, 2013).

Although both pricing and inventory strategies are intended to serve the same purpose of reducing supply-demand mismatch and increasing the firm's profit, many retailers choose to adopt both of these strategies, sometimes in an uncoordinated way. This raises an important research question as to whether managing the supply-demand balance in a multi-channel environment using a pricing tool or an inventory tool is more effective under which situations. Furthermore, if a retailer uses both tools, how do they interact with each other?

In this chapter, we consider a dual-channel retailer selling over a finite horizon. The retailers in this category include many of the major department stores, electronics stores, and fashion retailers, who offer short life-cycle products both online and in physical stores ("brick-and-click" retailers). Facing uncertain and price-sensitive demand at each channel, the retailer may use a pricing strategy (charging different prices at different channels), or a transshipment strategy (transferring inventory between store and distribution center) [1] to maximize total profits from sales. We are interested in answering the following research questions:

1. What is the optimal dynamic pricing and transshipment strategy?

2. What factors affect the benefit of price differentiation and transshipment?

3. Is transshipping inventory more or less effective than differentiating prices under which situations?

4. If the retailer adopts both price-differentiating and transshipment strategies, how is the optimal pricing decision influenced by the transshipping decision, and vice versa?

We model a joint dynamic pricing and transshipping problem, where the retailer incurs a unit transaction cost when selling the product at each channel, and a unit transshipping cost when making a transshipment of inventory between the channels. An arriving customer decides whether to purchase the product from one of the available channels, based on her valuations and the observed prices at the channels. In this setting, although both channels sell the same product, the customer may derive different utilities when purchasing from different channels due to the nature of transactions (e.g. ability to try on) and associated costs (e.g. shipping fee) involved. Hence, the product sold at one channel can be perceived as a different product from the product sold at the other channel. This justifies the assumption that the customer choice model follows the multinomial logit (MNL) model, which is widely used in the multi-product literature.

[1] This is an instance of lateral transshipments (stock movements between locations of the same echelon (*Paterson et al.*, 2011))

To answer the research questions, we first consider the situation where the retailer can adopt either a price differentiation policy or an inventory transshipment policy in the current period, but not both. We characterize the retailer's optimal dynamic pricing and transshipping policies. Our findings show that the optimal price differentiation policy always results in a larger probability of making a sale in the current period, compared to what would happen under the optimal uniform pricing policy. On the other hand, a transshipment in the current period may or may not be profitable, and an optimal transshipment decision may result in either a larger or a smaller probability of making a sale in the current period. We also investigate the factors that affect the benefit from adopting a price differentiation policy or a transshipment policy. Next, we consider the situation where the retailer can utilize both price differentiation and inventory transshipment. We show that transshipment can increase the value of the remaining inventory at the channel from which the transshipment is made, allowing the retailer to charge a higher price for the product at the channel. Transshipment can also be used to replenish stock at the channel that stocks out. This makes it possible for the retailer to continue selling the product at both channels and benefit from price differentiation.

To further investigate the benefit of price differentiation and transshipment, we conduct a numerical study. Our results show that the benefit of price differentiation is generally larger than the benefit of transshipment. However, when the retailer's inventory position is significantly out of balance (e.g. very low inventory at the channel with high customer valuation), transshipment can be more effective than price differentiation. We also find that the benefit of price differentiation and the benefit of transshipment may either substitute or complement each other. When the retailer's inventory position is unfavorable (low inventory at the high-margin channel), the two mechanisms substitute each other since the retailer can use either mechanism to try to adjust his inventory position in the intended direction. On the other hand, when the retailer's inventory position is already favorable (high inventory at the high-margin channel), the same balance of the inventory levels at the channels should be maintained. In this case, the two mechanisms are complementarily employed to influence the inventory position in the opposite direction, so that the balance is kept.

4.2 Literature Review

Our research problem lies in the area of both dynamic pricing and transshipment decisions for multi-channel retailers. These two topics have been studied in the Operations Management as well as the Marketing literature. We will review three main streams of relevant work: dynamic pricing, transshipment for multi-location or multi-channel retailers, and joint dynamic pricing and inventory policies.

Dynamic pricing problems have received much attention from the academia due to its popularity and practicality in many industries. One of the most influential works in this field is *Gallego and van Ryzin* (1994), who consider a dynamic pricing problem of a single product over a finite horizon, and show that the optimal price is strictly decreasing in the stock level but increasing in the length of the selling horizon. A similar problem is considered in *Bitran and Mondschein* (1997), with extensions to periodic-review pricing and pricing policies with announced discounts. Their computational experiments reveal that loss of profits when using periodic review instead of continuous review is small. With announced discounts, the resulting optimal prices allow the store to sell most of the merchandise during the first periods and avoid offering large discounts toward the end of the horizon. A natural extension of these works is to consider a dynamic pricing problem for multiple products. When a firm sells multiple products, the demand for each product may be influenced by the availability and price of the other products that the customers consider as substitutes. This gives rise to research problems on dynamic pricing of substitutable products. A similarity between the setting where a retailer sells substitutable products and our setting where the retailer sells an identical product through multiple channels is that a customer's decision on which product to buy or where to buy it from depends on her valuations and prices of all the possible choices. Hence, we assume our customer choice model is characterized by the MNL model, which is widely used among the substitutable products literature. *Dong et al.* (2009) study a dynamic pricing problem of multiple substitutable products, where the customer's choice is explained by the MNL model. They show that dynamic pricing converges to static pricing as inventory levels of all products approach the number of remaining selling periods. Furthermore, their numerical findings, considering only two substitutable products, suggest that the performance of unified dynamic pricing (charging the same price for all products) is closest to that of the full-scale dynamic pricing especially when the quality difference among the products decreases. *Suh and Aydin* (2011) considers a dynamic pricing problem of two substitutable products; they also adopts the MNL model to describe the customer's choice. They provide analytical results showing that the marginal value of an item is increasing in the remaining time and decreasing in the stock level of either product; however, the optimal prices of an item do not always behave in the same direction as the marginal value. Another multi-product pricing paper using logit models is *Li and Huh* (2011). They consider the nested logit model and show the concavity of the seller's profit function with respect to market shares in a single-period setting. This result can be applied to other settings, including the joint inventory and dynamic pricing problem, to find optimal policies. While these papers consider similar customer choice model and seller's dynamic pricing problem to ours, they do not consider inventory transshipments since the substitutable products in their settings are not identical. Additionally, no analytical results regarding the benefits of price differentiation are provided.

Another group of pricing papers which are relevant to our work is in the area of multi-channel pricing. According to some earlier studies and practitioners' beliefs, retailers generally keep consistent prices across distribution channels to maintain a strong brand and to avoid customers' irritation due to the perception of price unfairness (*Neslin et al.*, 2006; *Campbell*, 1999). However, other studies argue that channel-based price differentiation could be justified by differences in channel characteristics and the fact that consumers derive different utility from various distribution channels (*Chu et al.*, 2007; *Kacen et al.*, 2003). Hence, options to differentiate price levels among channels can in fact create opportunities for firms to improve their pricing strategies (*Sotgiu and Ancarani*, 2004). Based on data collected from multi-channel retailers, *Wolk and Ebling* (2010) find that many multi-channel retailers do engage in channel-based price differentiation, with some indication that this tendency increases over time. Analytical works on multi-channel pricing mostly consider a dual-channel seller selling a product through a physical store and an online store. *Yan* (2008) assumes that customers value a purchase from the physical store more than that from the online store. It is shown that the optimal online price is higher than the physical store price if and only if the marginal cost to sell the product online is far larger than the marginal cost to sell through the physical channel. *Shen and Zhang* (2012) consider a market with two groups of customers: fashion customers who value an online purchase more, and traditional customers who value a physical-store purchase more. They show that when there exist enough fashion customers and the unit cost for the online channel is low, it is profitable for the seller to offer the product through the online channel in addition to the traditional store, and charge a higher price. These papers consider a pricing problem of a dual-channel retailer similar to our essay. However, they study a single-period problem with different customer's choice models from ours, and without inventory or transshipment consideration.

The second stream of relevant literature is on inventory transshipments in a multi-location system. *Rudi et al.* (2001) study a problem of two retail firms at distinct locations selling the same product; the firm who runs out of stock may receive a transshipment from the other firm with surplus inventory at a cost. Under central coordination, this problem is similar to our transshipment problem in the sense that the system's objective is to make transshipment decisions which maximize the total profits. However, they consider a single-period model, where demand at the two locations are independent and retail prices are exogenously given. A multi-period setting is considered in *Hu et al.* (2005), where a centralized-ordering system of N stores periodically decides on order sizes, allocation quantities, and if necessary, emergency transshipments among the stores, in order to minimize the total expected cost until the end of the horizon. Since they focus on inventory problems and the systems's objective is to minimize costs rather than maximize profits, dynamic pricing is not addressed. For additional review of inventory transshipment literature, please see *Paterson et al.* (2011).

The closest literature to the current essay is in the area of joint dynamic pricing and inventory policies. A joint dynamic pricing and inventory problem of a distribution system consisting of multiple geographically dispersed retailers is considered in *Federgruen and Heching* (2002). In each period, the system decides on size of a replenishment, the price to be charged, and the allocation of any arriving order to the retailers. They provide an approximate model where a base-stock/list-price policy is optimal. The optimal price is shown to be nonincreasing in the system-wide inventory position. What differentiates our work form their work is that we allow prices at different locations to differ; whereas, in their model, the price in each period is applied to all stores. *Moon et al.* (2010) study a joint dynamic pricing and inventory problem in a dual-channel supply chain system. A customer chooses to buy the product from the channel where she receives a larger surplus. Under the vertical integration setting, the manufacturer decides on the production quantity as well as the price for each channel. Their inventory problem is rather different from ours because they assume that production can occur at any time. Hence, there are no stockouts, and no need for transshipments in their model. Another relevant work in this area is *Ceryan et al.* (2013), who study a joint dynamic pricing and capacity allocation problem for two substitutable products sharing a flexible resource. Their results suggest that the availability of a flexible resource helps maintain stable price differences across products over time even though the price of each product may fluctuate. The allocation of a flexible resource among substitutable products is similar to the option to transship inventory among the channels in our model. However, in their setting, the products can be replenished in each period; while, in our setting, the inventory is not replenishable. This difference can lead to different implications of inventory decisions on the system's performance. The closest work to ours in terms of inventory model is *Bitran et al.* (1998), considering a dynamic pricing problem of a retail chain of multiple stores, who has an option to transfer merchandise among stores at a cost. They propose a heuristic and numerically show that it performs better than the current practice in a fashion retail chain in Chile. While the transshipment problem in their paper is similar to the transshipment problem that we consider, there are some notable differences in other dimensions. Their paper considers coordinated prices and independent demand among stores. On the other hand, we allow different stores to charge different prices and let a customer's purchase decision depend on her valuations and prices of all stores. Furthermore, we analytically characterize the optimal pricing and transshipping policies, and investigate their benefits to the retailer.

To our best knowledge, we are the first to consider a joint dynamic price differentiating and inventory transshipping policy for a dual-channel retailer. This model enables us to answer important questions regarding best practices in pricing and transshipping strategies, which have not been addressed by existing literature, for instance: How do transshipping policies affect pricing policies, and vice versa? Is transshipping more or less effective than pricing, under what situations? We address these questions in subsequent sections.

4.3 Model

We consider a retailer who sells a seasonal product through two channels (e.g. "brick-and-click" retailer who sells products both online and at a physical store, retailer who has two physical stores at different locations) over a finite horizon of length T. Since the selling horizon is short, we assume no replenishment can take place during the horizon. No salvage value is obtained for unsold units at the end of the horizon. Stocks of the product are kept at two separate locations, dedicated for demand arriving at each channel. At the beginning of period $t = 1, ..., T$, the inventory level is denoted by $\mathbf{I^t} = (I_1^t, I_2^t)$, where $I_i^t \geq 0$ is the level of inventory to satisfy demand at channel $i \in \{1, 2\}$. The retailer decides on i) whether to transship any stock from one channel to the other[2], and ii) how much to charge for the product sold at each channel.

The retailer's transshipment decision is characterized by $\mathbf{s^t} = (s_{12}^t, s_{21}^t)$, where s_{ij}^t is the amount of inventory being transshipped from channel i to channel j, $i \neq j$. For simplicity, any transshipment is assumed to occur instantaneously and before a customer arrival in each period.[3] Let $\mathbf{Y^t} = (Y_1^t, Y_2^t)$, where $Y_i^t = I_i^t + s_{ji}^t - s_{ij}^t, i, j \in \{1, 2\}$ denote the inventory level at channel i after a transshipment is made in period t. Notice that a transaction can occur at channel i only when $Y_i^t > 0$. We let $A(\mathbf{Y^t}) = \{i : Y_i^t > 0\}$ denote the set of channels with the product available in stock.

The retailer's pricing decision is characterized by $\mathbf{p^t} = (p_1^t, p_2^t)$, where p_i^t is the price of a unit of product sold at channel i in period t. In each period, a pricing policy $\mathbf{p^t}$ is made based on the updated inventory level $\mathbf{Y^t}$, and is announced before a customer arrival.

[2]A transshipment in our setting refers to the practice where inventory dedicated for a channel is used to satisfy demand at the other channel with or without the actual transfer of inventory between two warehouses.

[3]Transshipment lead times can be incorporated with slight modifications.

4.3.1 Customer Choice Model

We assume that each period is short enough that at most one customer arrives, and each customer buys at most one unit of the product. The probability that a customer arrives and demands the product in a single period t is $\lambda^t \in [0, 1]$. Due to easy access to price and availability information nowadays, we assume that an arriving customer can observe the product price and availability at both channels before making a purchase decision. Hence, p_1^t, p_2^t, and $A(\mathbf{Y^t})$ are known to the customer.[4] The customer only considers purchasing the product from one of the channels with product availability ($i \in A(\mathbf{Y^t})$). We adopt a multinomial logit model (MNL), which has been extensively used in the marketing and operations management literature, to describe the choice of an arriving customer to make a purchase from one of the available channels, or neither. The MNL model nicely captures both the known and random factors that influence the purchase decision of a utility-maximizing customer in a dual-channel setting while still providing desirable properties, which make the analyses tractable.

Let $U_i = v_i - p_i + \zeta_i$ be a customer's net utility from purchasing a unit of the product from channel $i \in A(\mathbf{Y^t})$ at price p_i, where v_i denotes the customer's net valuation from the purchase, and ζ_i is a Gumbel error term with mean 0 and shape parameter 1. The customer's net valuation v_i represents the product valuation adjusted for channel-specific (dis)utilities. These include, for instance, the ability to try on the product, traveling time, customer service, time until the product arrives, etc. Although the customer's valuation of the product itself should generally be the same for both channels, the customer may have shopping preferences for a channel over the other, resulting in different net valuation of the purchase at each channel. The random variable ζ_i represents the utility influenced by unobservable characteristics.

[4]Customers do not need to know the exact inventory level at each channel, but they can observe whether the product is in stock or out of stock at each channel.

We assume that ζ_1 and ζ_2 are independent and identically distributed.

Let μ_i be the probability that an arriving customer purchases from channel $i \in \{1,2\}$, and μ_0 be the probability that the customer does not purchase. Then, we have the following well-known results from the MNL model (*Luce*, 1959; *McFadden*, 1974), adjusted for the condition for the product availability (*Suh and Aydin*, 2011):

$$\begin{aligned} \mu_i(\mathbf{p^t}, A(\mathbf{Y^t})) &= \begin{cases} \frac{\exp(v_i - p_i^t)}{1+\sum_{j\in A(\mathbf{Y^t})}\exp(v_i - p_i^t)} & \text{if } i \in A(\mathbf{Y^t}) \\ 0 & \text{if } i \notin A(\mathbf{Y^t}) \end{cases} \\ \mu_0(\mathbf{p^t}, A(\mathbf{Y^t})) &= \frac{1}{1+\sum_{j\in A(\mathbf{Y^t})}\exp(v_i - p_i^t)} \end{aligned} \tag{4.1}$$

For notational simplicity, we will sometimes write $\mu_i(\mathbf{p^t}, A(\mathbf{Y^t}))$ as μ_i when $\mathbf{p^t}$ and $A(\mathbf{Y^t})$ are clearly stated in the context.

4.3.2 Retailer's Problem

In a period $t = 1, ..., T$, the retailer's problem is to decide whether to make a transshipment from one channel to the other: $\mathbf{s^t} = (s_{12}^t, s_{21}^t)$, and what price to charge for the product sold at each channel: $\mathbf{p^t} = (p_1^t, p_2^t)$. Notice that if the product is out of stock at channel i ($Y_i^t = 0$), then p_i^t becomes irrelevant since a customer will never consider buying the product from channel i. We let $p_i^t \to \infty$ for $i \notin A(\mathbf{Y^t})$.

The unit transshipment cost from channel i to channel j is $m_{ij} \geq 0$. To avoid unnecessary transshipments, we assume that $s_{12}^t \cdot s_{21}^t = 0$. Since a transshipment occurs instantaneously, we can innocuously restrict the retailer's transshipment decision to be such that a positive transshipment may be made only when the current inventory at the destination channel is zero.[5] Furthermore, since the retailer can sell at most one unit in each period, there is no need to transship more than one unit at a time. As a result, the transshipment decisions in our model can be simplified to $s_{ij}^t \in \{0,1\}$

[5] This practice corresponds to "reactive transshipment," as opposed to "proactive transshipment", in *Paterson et al.* (2011).

for all $t = 1, ..., T$.

We assume that the production cost of the product is sunk. However, the retailer incurs a total unit transaction cost of c_i^t when a unit of product is sold at channel $i \in \{1, 2\}$.

Let $V^t(\mathbf{I^t})$ denote the retailer's expected discounted profit-to-go under the optimal policy in period t with the initial inventory level $\mathbf{I^t} = (I_1^t, I_2^t), t = 1, ..., T$, and a discount rate $\beta \in [0, 1]$. Then, the retailer's dynamic program when he can adopt both price differentiating and inventory transshipping policies is given by:

$$
\begin{aligned}
V^t(\mathbf{I^t}) &= \max_{\mathbf{p^t},\mathbf{s^t}} \left\{ \lambda^t \left[\sum_{i\in\{1,2\}} \mu_i \left[p_i^t - c_i^t + \beta V^{t-1}\left(Y_i^t - 1, Y_{-i}^t\right)\right] \right] \right. \\
&\quad \left. + \left(\lambda^t \mu_0 + (1-\lambda^t)\right) \beta V^{t-1}(Y_1^t, Y_2^t) - \sum_{i,j\in\{1,2\}} m_{ij} s_{ij}^t \right\} \\
&= \max_{\mathbf{p^t},\mathbf{s^t}} \left\{ \lambda^t \mu_1 \left[p_1^t - c_1^t + \beta V^{t-1}\left(Y_1^t - 1, Y_2^t\right)\right] \right. \\
&\quad + \lambda^t \mu_2 \left[p_2^t - c_2^t + \beta V^{t-1}\left(Y_1^t, Y_2^t - 1\right)\right] \\
&\quad \left. + \left(\lambda^t \mu_0 + \left(1-\lambda^t\right)\right) \beta V^{t-1}\left(Y_1^t, Y_2^t\right) - m_{12} s_{12}^t - m_{21} s_{21}^t \right\} \\
&= \max_{\mathbf{p^t},\mathbf{s^t}} \left\{ \beta V^{t-1}\left(Y_1^t, Y_2^t\right) + \lambda^t \mu_1 \left[p_1^t - c_1^t + \beta V^{t-1}\left(Y_1^t - 1, Y_2^t\right) - \beta V^{t-1}\left(Y_1^t, Y_2^t\right)\right] \right. \\
&\quad \left. + \lambda^t \mu_2 \left[p_2^t - c_2^t + \beta V^{t-1}\left(Y_1^t, Y_2^t - 1\right) - \beta V^{t-1}\left(Y_1^t, Y_2^t\right)\right] - m_{12} s_{12}^t - m_{21} s_{21}^t \right\}
\end{aligned}
$$

$$
\begin{aligned}
\text{s.t.} \quad & Y_i^t = I_i^t + s_{ji}^t - s_{ij}^t \text{ for } i, j \in \{1, 2\}, t = 1, ..., T \\
& \qquad\qquad\qquad \text{(inventory level after transshipment)}, \\
& 0 \le s_{ij}^t \le I_i^t,\ s_{ij}^t = 0 \text{ if } I_j^t > 0 \text{ for } i, j \in \{1, 2\}, t = 1, ..., T \\
& \qquad\qquad\qquad \text{(transshipment constraints)}, \\
& p_i^t \ge 0 \text{ for } i \in \{1, 2\}, t = 1, ..., T \qquad \text{(nonnegative pricing)}, \\
& V^0(\mathbf{I}) = 0 \text{ for all } \mathbf{I} \qquad \text{(end of horizon)}.
\end{aligned}
$$

Notice that since the customer only considers purchasing from a channel where the product is available, we have $\mu_i > 0$ if and only if $Y_i^t \geq 1$. To simplify the expressions, we define $\Delta_i^t(I) = V^{t-1}(I) - V^{t-1}(I - \mathbf{e_i})$, for $i \in \{1, 2\}$, where $\mathbf{e_1} = (1, 0)$ and $\mathbf{e_2} = (0, 1)$. That is, $\Delta_i^t(I)$ is the marginal value of the product at channel i, given the inventory level I and the remaining time t. Then, we can rewrite the retailer's profit-to-go as:

$$\begin{aligned} V^t(\mathbf{I^t}) &= \max_{\mathbf{p^t},\mathbf{s^t}} \Big\{ \beta V^{t-1}\left(Y_1^t, Y_2^t\right) + \lambda^t \mu_1 \left[p_1^t - c_1^t - \beta \Delta_1^t \left(Y_1^t, Y_2^t\right)\right] \\ &\quad + \lambda^t \mu_2 \left[p_2^t - c_2^t - \beta \Delta_2^t \left(Y_1^t, Y_2^t\right)\right] - m_{12} s_{12}^t - m_{21} s_{21}^t \Big\} \\ &= \max_{\mathbf{p^t},\mathbf{s^t}} \left\{ \beta V^{t-1}\left(Y_1^t, Y_2^t\right) - m_{12} s_{12}^t - m_{21} s_{21}^t + \lambda^t J^t(\mathbf{p^t}, \mathbf{s^t}, \mathbf{Y^t}) \right\} \end{aligned} \tag{4.2}$$

Here, we let J^t denote the terms in the retailer's profit-to-go which involve the retailer's pricing decisions.

Furthermore, to help explain the effects of pricing on the retailer's profit more intuitively, we define "sale ratio" as follows.

Definition 4.1. Sale ratio $[R(\mathbf{p^t}, A(\mathbf{Y^t}))]$ is the probability of making a sale divided by the probability of not making a sale in the current period, for given prices and product availability. Mathematically, $R(\mathbf{p^t}, A(\mathbf{Y^t})) := \frac{1-\mu_0(\mathbf{p^t}, A(\mathbf{Y^t}))}{\mu_0(\mathbf{p^t}, A(\mathbf{Y^t}))}$.

When the product price is set optimally, we refer to the resulting sale ratio as the optimal sale ratio $[R^{t*} = R(\mathbf{p^{t*}}, A(\mathbf{Y^t}))]$.

In the next section, we will characterize the retailer's optimal pricing and transshipping policy.

4.4 Optimal Pricing and Transshipping Policies

To understand how the price differentiation and inventory transshipment mechanism work independently or concurrently to help improve the retailer's profit in a dual-channel environment, we consider situations where the retailer implements each mechanism alone, and when he can implement both mechanisms together. To represent different policy options, we will denote a joint pricing-transshipping policy with AB, where $A \in \{u, d\}$ denotes the pricing option (u for uniform pricing, d for differentiated pricing), and $B \in \{n, t\}$ denotes the transshipment option (n for no transshipment, t for transshipment). For example, if the retailer adopts a policy of type dt, he can set different prices at the two channels, and make an inventory transshipment from the non-empty location to the empty location. On the other hand, under a policy of type un, the retailer always sets the same price at both channels, and cannot transfer any inventory between the channels. Table 4.1 summarizes the four types of policies we consider in this essay.

Price Differentiation	No Transshipment	With Transshipment
Yes	dn	dt
No	un	ut

Table 4.1: Possible pricing and transshipping policies

As a benchmark, we first study the base model where the retailer adopts a uniform pricing policy without transshipment.

4.4.1 Base Model: Uniform Pricing and No Transshipment

When the retailer charges the same price for the product sold at both channels and when transshipment is not an option, the retailer's only decision is the optimal uniform price for each period. In this case, the retailer's dynamic program is given by equation (4.2), with the following constraints:

$$Y_i^t = I_i^t \text{ for } i, j \in \{1, 2\}, t = 1, ..., T,$$

$$s_{ij}^t = 0 \text{ for } i, j \in \{1, 2\}, t = 1, ..., T \text{ (no transshipment)},$$

$$p_1^t = p_2^t = p^t \geq 0 \text{ for } t = 1, ..., T \text{ (uniform pricing)},$$

$$V^0(\mathbf{I}) = 0 \text{ for all } \mathbf{I}.$$

Let V_{un}^t denote the retailer's optimal profit-to-go under uniform pricing and no transshipment. Then,

$$V_{un}^t(\mathbf{I^t}) = \max_{p^t} \left\{ \beta V_{un}^{t-1}\left(I_1^t, I_2^t\right) + \lambda^t J^t(p^t, \mathbf{I^t}) \right\}. \tag{4.3}$$

We derive the retailer's optimal uniform dynamic pricing policy under two scenarios: i) when inventory is in stock at both channels, and ii) when inventory runs out at one channel. (If inventory runs out at both channels, the selling horizon ends immediately.) In the first scenario, the retailer can expect to make a sale at each channel with a positive probability. However, in the second scenario, the retailer can never make a sale at the channel where the stockout occurs, regardless of the price. Without loss of generality, when analyzing the stockout scenario, we assume that the inventory runs out at channel 1. The results for the case when the inventory runs out at channel 2 is analogous. Proposition 4.1 characterizes the optimal uniform dynamic pricing and the optimal profit.

Proposition 4.1. *For any remaining time $t = 1, ..., T$ and inventory level $\mathbf{I^t}$, the optimal uniform pricing policy without transshipment is as follows:*

1. *If $I_1^t > 0$ and $I_2^t > 0$, the optimal uniform price p_{un}^{t*} is unique and it satisfies $p_{un}^{t*} = \frac{\exp(v_1)(c_1^t + \beta\Delta_1^t(\mathbf{I^t})) + \exp(v_2)(c_2^t + \beta\Delta_2^t(\mathbf{I^t}))}{\exp(v_1)+\exp(v_2)} + \frac{1}{\mu_0(p_{un}^{t*})}$. The retailer's expected discounted profit-to-go under the optimal policy is given by $V_{un}^t(\mathbf{I^t}) = [\exp(v_1 - p_{un}^{t*}) + \exp(v_2 - p_{un}^{t*})]\lambda^t + \beta V_{un}^{t-1}(\mathbf{I^t}) = R_{un}^{t*}\lambda^t + \beta V_{un}^{t-1}(\mathbf{I^t})$.*

2. *If $I_1^t = 0$ and $I_2^t > 0$, the optimal uniform price p_{un}^{t*} is unique and it satisfies $p_{un}^{t*} = c_2^t + \beta\Delta_2^t(\mathbf{I^t}) + \frac{1}{\mu_0(p_{un}^{t*})}$. The retailer's expected discounted profit-to-go under the optimal policy is given by $V_{un}^t(\mathbf{I^t}) = \exp(v_2 - p_{un}^{t*})\lambda^t + \beta V_{un}^{t-1}(\mathbf{I^t}) = R_{un}^{t*}\lambda^t + \beta V_{un}^{t-1}(\mathbf{I^t})$.*

When there is no stockout, the optimal uniform price consists of two components. The first component is a weighted average of the sum of the marginal value and the unit transaction cost associated with each channel. This quantity depends on how likely a customer will buy from each channel. Since the price is uniform, what determines the purchase decision is simply the net valuation of the purchase, v_1 and v_2. The optimal price is more influenced by the marginal value and transaction cost of the channel from which a customer is more likely to make a purchase. The other component of the optimal price is a base premium, $\frac{1}{\mu_0(p_{un}^{t*})}$, which goes up with the probability that the customer makes a purchase from one of the channels.

When there is a stockout at a channel, the optimal uniform price has similar characteristics as in the no stockout case. That is, it includes a base premium on top of the marginal value and the unit transaction cost of the product. The only difference is that the optimal price in the stockout case involves only the marginal value and the unit transaction cost of the product at the channel with available inventory. This is because the product can never be sold from the stockout channel. The base premium remains the same as that under the no stockout case.

Next, we investigate the retailer's optimal pricing policy when the price at the two channels can be differentiated.

4.4.2 Price Differentiation and No Transshipment

Suppose the retailer can charge different prices for the product sold at different channels, but cannot transfer inventory between channels. Then, the retailer's pricing problem becomes similar to a standard dynamic pricing problem of two substitutable products (e.g., Suh and Aydin 2011). The retailer's dynamic program in this case is similar to (4.3) except that p_1^t and p_2^t in each period can be different, and hence, Δ_i^t under this *pn* policy can be different from that under the *un* policy. Let V_{dn}^t denote the retailer's optimal profit-to-go under price differentiation and no transshipment.

Then,

$$V_{dn}^t(\mathbf{I^t}) = \max_{\mathbf{p^t}} \left\{ \beta V_{dn}^{t-1}\left(I_1^t, I_2^t\right) + \lambda^t J^t(\mathbf{p^t}, \mathbf{I^t}) \right\}, \ \mathbf{p^t} = (p_1^t, p_2^t). \tag{4.4}$$

Under this pricing policy, only the prices at the channels where the product is available are relevant. Hence, when there is a stockout at channel 1, the retailer essentially sets only one price at channel 2. This optimal price corresponds to the optimal uniform price in the stockout case, as characterized by Proposition 4.1 part 2. The retailer can take advantage of price differentiation only when the product is available at both channels. We characterize the optimal dynamic pricing policy in Proposition 4.2.

Proposition 4.2. *For any remaining time $t = 1, ..., T$ and inventory level $\mathbf{I^t}$, the optimal dynamic pricing policy without transshipment is as follows:*

1. *If $I_1^t > 0$ and $I_2^t > 0$, the optimal price pair (p_1^{t*}, p_2^{t*}) is unique and it satisfies $p_i^{t*} = c_i^t + \frac{1}{\mu_0(p_1^{t*}, p_2^{t*})} + \beta\Delta_i^t(\mathbf{I^t})$. The retailer's expected discounted profit-to-go under the optimal price differentiation policy is given by $V_{dn}^t(\mathbf{I^t}) = [\exp(v_1 - p_1^{t*}) + \exp(v_2 - p_2^{t*})]\lambda^t + \beta V_{dn}^{t-1}(\mathbf{I^t}) = R_{dn}^{t*}\lambda^t + \beta V_{dn}^{t-1}(\mathbf{I^t})$.*

2. *If $I_1^t = 0$ and $I_2^t > 0$, the optimal price at channel 2, p_2^{t*}, is the same as the optimal uniform price for the same inventory level.*

Under the price differentiation policy, when there is no stockout, we see that the optimal price for each channel has a similar structure as the optimal uniform price discussed in Proposition 4.1 part 1. That is, the optimal channel price consists of the marginal value and unit transaction cost, specific to the channel, as well as the base premium, which is the same for both channel. Alternatively, we can view this result as the optimal margin $p_i^{t*} - c_i^t$ for each channel is the sum of the marginal value of the product in the channel and a base premium, which is channel-independent.

Now that we have characterized the retailer's optimal pricing policies, we can compare the retailer's profit under the price differentiation and uniform pricing policy to see the benefit of price differentiation for a dual-channel retailer. Since many problem variables are time-dependent, we focus our comparisons on the retailer's profit under different pricing policies in a single period in order to provide clean results that highlight the benefit of the price differentiation strategy. For this, we consider the benefit from price differentiation as how much profit the retailer can gain if he is allowed to price differentiate in period t only, as supposed to using a pure uniform pricing policy.[6]

Corollary 4.1 shows that the optimal sale ratio under price differentiation is always greater than or equal to the optimal sale ratio under uniform pricing. This implies that the retailer's ability to price discriminate between the two channels allows him to strategically convert more of the customer traffic into sales. In fact, as shown in Proposition 4.3, the larger the sale ratio he can induce using price differentiation, the more profit he can generate.

Corollary 4.1. *For any remaining time $t = 1, ..., T$ and inventory level $\mathbf{I^t}$, the optimal sale ratio under price differentiation is greater than or equal to the optimal sale ratio under uniform pricing: $R_{dn}^{t*} \geq R_{un}^{t*}$.*

Proposition 4.3. *For any remaining time $t = 1, ..., T$ and inventory level I^t where $I_1^t > 0$ and $I_2^t > 0$, the benefit from using price differentiation in the current period is monotonically increasing with the sale ratio difference, $R_{dn}^{t*} - R_{un}^{t*}$.*

Due to differences in customer preferences and nature of the transactions taken place at different channels, the customer willingness to pay and the costs incurred by the retailer when selling the product at each channel generally differ across the channels. Price differentiation allows the retailer to reflect such differences on the product prices. That is, the retailer can charge a high price at the channel where the customer has high willingness to pay and/or the transaction cost is high in order

[6] Alternatively, we can consider how much the retailer will lose if constrained to offer a uniform price in period t only. The same results hold.

to extract more surplus and protect the margin. On the other hand, for the channel with a low customer willingness to pay and a low transaction cost, the retailer can generate more profit by increasing the sales volume. Hence, a lower price should be offered. By charging the prices that customers are willing to pay, the retailer can attract more customers to purchase the product while maintaining a reasonable level of sales margin. This cannot be done as effectively under uniform pricing since the same price is charged for the product sold at both channels. While the price may seem attractive for the channel with a higher customer willingness to pay, the price will be deemed too expensive for the other channel with a lower customer willingness to pay. To balance the sales volume and margin at both channels using a single price, the retailer cannot attract as much customers to purchase, resulting in a smaller sale ratio under uniform pricing than under price differentiation. Since the retailer's profit is directly correlated to the sale ratio, the profit difference between the two pricing strategies is amplified as the difference in sale ratio gets larger.

Next, we investigate the optimal transshipment policy and the benefit of transshipment.

4.4.3 Uniform Pricing with Transshipment

When the retailer can transfer inventory between the channels but cannot charge different prices, his profit is as given by (4.2) except that $p_1^t = p_2^t = p^t$ for all t. Let V_{ut}^t denote the retailer's optimal profit-to-go under uniform pricing and transshipment. Then,

$$\begin{aligned} V_{ut}^t(\mathbf{I^t}) \quad &= \quad \max_{p^t, \mathbf{s^t}} \left\{ \beta V_{ut}^{t-1}\left(Y_1^t, Y_2^t\right) - m_{12}s_{12}^t - m_{21}s_{21}^t + \lambda^t J^t(p^t, \mathbf{s^t}, \mathbf{Y^t}) \right\} \quad (4.5) \\ \text{s.t.} \quad & Y_i^t = I_i^t + s_{ji}^t - s_{ij}^t \text{ for } i,j \in \{1,2\}, t = 1,...,T, \\ & s_{ij}^t \in \{0,1\},\ s_{ij}^t = 0 \text{ if } I_j^t > 0 \text{ for } i,j \in \{1,2\}, t = 1,...,T, \\ & V^0(\mathbf{I}) = 0 \text{ for all } \mathbf{I}. \end{aligned}$$

Under the assumption that a transshipment occurs instantaneously and the transshipment costs and discount rate are time-independent, it is easy to see that the retailer's optimal strategy in the current period is to not make any transshipment as long as there is some inventory at both channels. Thus, $s_{ij}^{t*} = 0, i, j \in \{1, 2\}$ whenever $I_1^t > 0$ and $I_2^t > 0$. We will consider a transshipment decision only when there is a stockout. As before, we assume that the stockout occurs at channel 1: $I_1^t = 0$ and $I_2^t > 0$. When there is a stockout at channel 1, a transshipment in the current period may or may not be profitable. Since the retailer can sell at most one unit in each period ($\lambda^t \leq 1$), he only needs to transfer at most one unit of inventory from channel 2 to channel 1. Hence, the retailer's transshipment decision in period t is essentially to decide whether to transfer a unit of inventory from channel 2 to channel 1 or not. That is, the optimal transshipment decision in period t is $\mathbf{s^{t*}} = (s_{12}^{t*} = 0, s_{21}^{t*} \in \{0, 1\})$.

Since a transshipment results in a change in the inventory level at both channels, there is a tradeoff between the marginal value of a unit of inventory before and after the transshipment. We define the quantity "transshipment marginal value" to capture this tradeoff, which plays an important role in determining optimal transshipment strategies.

Definition 4.2. Transshipment marginal value ($\Delta_{12}^t(\mathbf{I^t})$) is the difference between the marginal value of the inventory at channel 1 after the transshipment is made, and the marginal value of the inventory at channel 2 before the transshipment is made. Mathematically, $\Delta_{12}^t(\mathbf{I^t}) := \Delta_1^t(1, I_2^t - 1) - \Delta_2^t(0, I_2^t)$.

The retailer's optimal transshipment strategy and the benefit of transshipment under uniform pricing are characterized in Proposition 4.4. Again, to highlight the effect of transshipment on the retailer's profit, we consider the benefit from transshipment as the profit gain when the retailer is allowed to transship a unit of inventory to the stockout channel in a single period. To exclude the effect from pricing policies, we assume the retailer uses the optimal uniform pricing policy in all periods.

Proposition 4.4. *For any remaining time $t = 1, ..., T$ and inventory level $\mathbf{I^t}$ such that $I_1^t = 0$ and $I_2^t > 0$, let $R_{un}^{t*}(R_{ut}^{t*})$ be the optimal sale ratio before (after) the transshipment; $p_{un}^{t*}(p_{ut}^{t*})$ be the optimal uniform price before (after) the transshipment.*

1. *It is optimal for the retailer to transship a unit from channel 2 to channel 1 ($s_{21}^{t*} = 1$) in period t if and only if $m_{21} \leq (R_{ut}^{t*} - R_{un}^{t*})\lambda^t + \beta\Delta_{12}^t(\mathbf{I^t})$.*

2. *Suppose a transshipment is made from channel 2 to channel 1 in period t. Then, $p_{ut}^{t*} \geq p_{un}^{t*}$ if $\beta\Delta_{12}^t(\mathbf{I^t}) \geq c_2^t - c_1^t$ and a) $I_2^t > 1$ and $\Delta_2^t(1, I_2^t - 1) \geq \Delta_2^t(0, I_2^t)$, or b) $I_2^t = 1$ and $v_1 \geq v_2$.*

3. *The benefit from making a transshipment in period t is monotonically decreasing in m_{21} and c_1^t, but is monotonically increasing in the sale ratio difference $R_{ut}^{t*} - R_{un}^{t*}$.*

When a transshipment is made, the retailer's inventory position is changed. This in turn influences the retailer's pricing decisions both in the current period and the future periods. Hence, the overall effects of a transshipment on the retailer's profit can be viewed as coming from two parts: the effects in the current period, and the effects in the future periods. After a transshipment takes place in the current period, the retailer charges an optimal price corresponding to the updated inventory position. This price may be different from what he would have charged if the transshipment was not made. The difference in prices leads to different customer purchase decisions, which subsequently result in the difference in the retailer's current period profit. This profit difference is therefore captured by the difference in the sale ratio before and after the transshipment, multiplied by the customer arrival rate: $(R_{ut}^{t*} - R_{un}^{t*})\lambda^t$. For future periods, the effects of transshipment essentially come from the tradeoff between the marginal value of a unit of inventory before and after the transshipment. Hence, the expected profit difference before and after the transshipment coming from the future periods is given by the discounted transshipment marginal value, $\beta\Delta_{12}^t(\mathbf{I^t})$. To decide whether a transshipment in the current period is profitable or not, the retailer needs to compare the overall potential gain in profit with a transshipment to the transshipment cost he has to incur. As characterized in Proposition 4.4 part 1, it is optimal for the retailer to make a transshipment if the expected gain in profit is large enough to cover the transshipment cost.

Part 2 of the proposition discusses an implication of the transshipment marginal value on the optimal uniform price before and after the transshipment. Recall from Proposition 4.1 that the optimal uniform price is larger the larger the unit transaction cost and the marginal value of the product. Without a transshipment, the optimal price depends on the transaction cost and the marginal value at channel 2 only since channel 1 stocks out. If a transshipment is made, then the optimal price will also depend on the transaction cost and the marginal value at channel 1. Hence, if the transshipment results in a significant increase in the marginal value of inventory, or if the transaction cost at channel 1 is large relative to the transaction cost at channel 2, then the retailer should charge a higher price after the transshipment. The same insight holds for a special case where there is only one unit of inventory left ($I_1^t = 0$ and $I_2^t = 1$). However, in this case, the optimal price after the transshipment is essentially the price at channel 1 alone since channel 2 stocks out. Hence, for the retailer to charge a higher price after the transshipment, the customer valuation at channel 1 should also be greater than the customer valuation at channel 2.

The benefit of transshipment depends on several factors as discussed in part 3 of the proposition. Intuitively, we see that the transshipment cost negatively affects the benefit because it is an additional cost that the retailer has to incur when making a transshipment. The benefit also goes down with the transaction cost at channel 1. This is because the main purpose of transshipping a unit of inventory from channel 2 to channel 1 is to trade the potential loss in profit from not selling the unit at channel 2 with the potential gain from selling the unit at channel 1. A higher transaction cost at channel 1 translates to a lower potential gain of selling a unit at channel 1, and subsequently, a smaller benefit from making a transshipment to channel 1.

A quantity that positively affects the benefit of transshipment is the sale ratio difference. As pointed out earlier, the sale ratio difference is directly linked to the current period gain in profit from the transshipment. Hence, the larger sale ratio generated by the transshipment, the larger benefit to the retailer. This result is similar to what we have seen from Proposition 4.3 that the benefit from price differentiation monotonically increases with the sale ratio difference. There is however a notable difference between the benefit from price differentiation and the benefit from transshipment. Recall from Corollary 4.1 that the optimal sale ratio under price differentiation is always greater than that under uniform pricing, implying the sale ratio difference between the two pricing policies is always positive. This is not the case when we consider the sale ratio difference between uniform pricing with transshipment and uniform pricing without transshipment. When the stockout channel does not have a significantly higher customer willingness to pay or a much lower transaction cost, the optimal price after the transshipment may result in a smaller sale ratio, signifying that the retailer does not benefit from the transshipment in the current period. Figure 4.1 illustrates a scenario when the optimal sale ratio after a transshipment is smaller than the optimal sale ratio before a transshipment ($R_{ut}^{t*} - R_{un}^{t*} < 0$). Since the retailer's current period profit is smaller with the transshipment, the benefit from transshipment monotonically decreases with the customer arrival rate in the current period. Notice however that in this example, the benefit of transshipment from future periods is still positive ($\Delta_{12}^{t}(\mathbf{I^t}) > 0$). Hence, when the customer arrival rate in the current period is sufficiently small, the overall benefit of transshipment is still positive.

So far, we have considered the retailer's optimal decisions and benefit from implementing price differentiation and inventory transshipment policies independently.

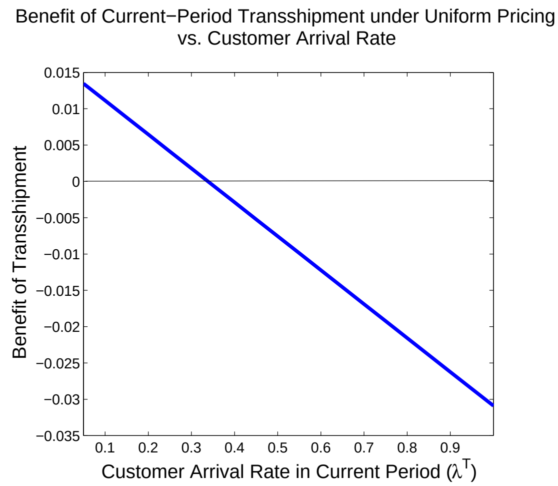

Figure 4.1: Benefit from transshipment in the current period vs. the customer arrival rate: $T = 5, I_1^5 = 0, I_2^5 = 2, \beta = 0.95, v_1 = 8.5, v_2 = 10, c_1^t = 1, c_2^t = 2, m_{12} = 0.2, m_{21} = 0.05, \lambda^t = 0.3$ for $t = 4, 3, 2, 1$

In the next section, we investigate the scenario where the retailer adopts both price differentiating and inventory transshipping policies together, and discuss how the two mechanisms interact with each other.

4.4.4 Price Differentiation with Transshipment

When evaluating the benefit of price differentiation and transshipment separately in a single period, we know that price differentiation is beneficial only when both channels hold some inventory. On the other hand, for a transshipment decision to be relevant in the current period, there must be a stockout at a channel. These different requirements prevent us from conducting a fair comparison of the two mechanisms when implemented in a single period. However, we are able to provide results in Theorem 4.1 which describe how a price differentiating decision as well as its benefit are influenced by a transshipment decision when the retailer adopts a joint price differentiating-inventory transshipping policy.

Theorem 4.1. *For any remaining time $t = 1, ..., T$ and inventory level I^t such that $I_1^t = 0$ and $I_2^t > 0$:*

1. *If the transshipment in period t increases the marginal value of the product at channel 2, then the optimal price of the product at channel 2 after the transshipment is higher than the optimal price without the transshipment. (If $\Delta_2^t(1, I_2^t - 1) \geq \Delta_2^t(0, I_2^t)$, then $p_{2,11}^{t*} \geq p_{2,10}^{t*}$).* [7]

2. *Under the optimal transshipment decision in the current period, the benefit from price differentiation is monotonically decreasing in the transshipment cost in the current period.*

Part 1 of Theorem 4.1 discusses an interesting dynamic between the price differentiation and transshipment policy. Since a transshipment allows the retailer to transfer inventory from the channel with an abundance to the channel with a shortage, the mechanism can result in an increased marginal value of inventory at the originating channel. This subsequently justifies the retailer charging a higher price for the product sold at the channel since he incurs less risk of overstocking. In a way, this result explains how the transshipment mechanism works to help the retailer avoid marking down prices at the channel with excessive inventory.

Transshipment also enables the retailer to leverage price differentiation to improve profits when facing a stockout situation. Without transshipment, the retailer is constrained to operate in only one channel, from which the benefit from price differentiation cannot be realized. A transshipment makes it possible for the retailer to replenish inventory at the stock-out channel, allowing him to continue selling at both channels and extract more profits using channel-based price discrimination. While transshipment can be beneficial, we learn from Proposition 4.4 part 1 that when the transshipment cost is too high, it is not optimal for the retailer to transship. In this case, the retailer makes more profit from selling at a single channel, without exercising price differentiation. This explains the result in part 2 of Theorem 4.1 that the benefit

[7] This result is relevant only when $I_2^t > 1$. If $I_2^t = 1$, then after the transshipment is made from channel 2 to channel 1, the inventory level at channel 2 becomes zero and $p_{2,10}^{t*}$ is irrelevant.

of price differentiation decreases as the transshipment cost increases. Notice however that this result applies to the situation where the retailer has only one opportunity (in the current period) to transship and/or price differentiate the product. If the retailer is allowed to implement both mechanisms over a longer period of time, the result may be different, as shown in Figure 4.2. In this example, we observe that the benefit from price differentiation ($V_{dn}^T - V_{un}^T$) increases with the transshipment cost when the cost is not too high. This is because for this particular setting, the retailer can employ either the price differentiation or the transshipment mechanism to balance inventory and demand at each channel over the selling horizon. In other words, the two mechanisms can substitute each other. (We will discuss more in the next section when the benefit from the two mechanisms may substitute or compliment each other.) Hence, when the transshipment cost is larger, the retailer increasingly prefers to use price differentiation rather than transshipment, resulting in larger benefits from price differentiation. When the transshipment cost is too large, however, the retailer finds it not profitable to transship even in a stockout situation, which could occur in future periods. Hence, the benefit from price differentiation is less likely to be realized in future periods, following the same insight from Theorem 4.1 part 2.

When considering the benefit of transshipment, many retailers may only think about the ability to use inventory at one channel to cover the possible loss in demand arriving at the other channel. The less visible, but potentially more substantial benefit of transshipment in enabling the retailer to realize the benefit from pricing to a larger extent is often understated or overlooked. Hence, our results in Theorem 4.1 as well as Proposition 4.4 part 2 provide important managerial insights to help retailers understand the transshipment and pricing mechanisms better.

Although transshipment can affect the optimal price levels, we note that the overall behavior of optimal prices with respect to inventory level and remaining time is not significantly changed by the transshipment option. As illustrated in Figure 4.3 a. and

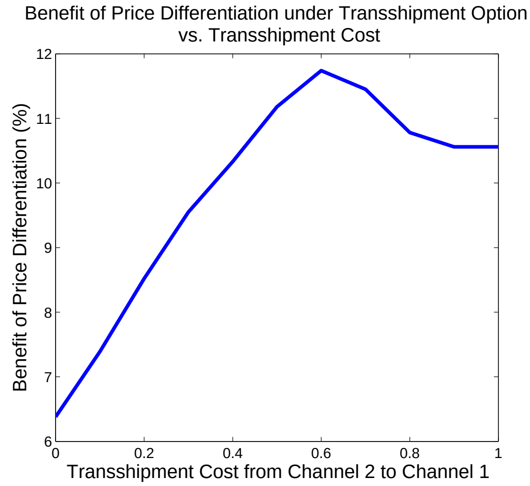

Figure 4.2: Benefit from price differentiation vs. transshipment cost: $T = 5, I_1^5 = 1, I_2^5 = 2, \beta = 0.95, v_1 = 4, v_2 = 2, c_1^t = 2, c_2^t = 0.5, m_{12} = 0.1, \lambda^t = 0.8$ for $t = 5, 4, ..., 1$

b., the optimal prices with and without transshipment options share similar trends with respect to the inventory level at channel 1. Likewise, Figure 4.3 c. and d. show that the optimal prices with and without transshipment options behave similarly with respect to the remaining time. Notice also that, unlike the case of a single channel (e.g. *Gallego and van Ryzin* (1994)), the optimal prices in our dual-channel setting are not necessarily monotone in either the inventory level or the remaining time. For example, as the inventory level at channel 1 increases, the optimal price at channel 2 may either decrease or increase. This is because an increase in stock at channel 1 requires the retailer to lower the price at channel 1 to reduce the risk of having unsold inventory at the end of the selling horizon. The lower price at channel 1 draws some customers away from channel 2. Hence, the retailer may need to lower the price at channel 2 as well in order to avoid overstocking at channel 2 due to the substitution effect. On the other hand, if the inventory level at channel 1 becomes a lot higher, the potential loss from overstocking at channel 1 is significantly larger than that from overstocking at channel 2. In this case, it may be optimal for the retailer to increase the price at channel 2, in addition to lowering the price at channel

1, to send even more customers from channel 2 to buy from channel 1. We note that similar non-monotonicity results have been observed in *Suh and Aydin* (2011) for the optimal prices of two substitutable products without transshipment options, where the customer choice also follows a MNL model.

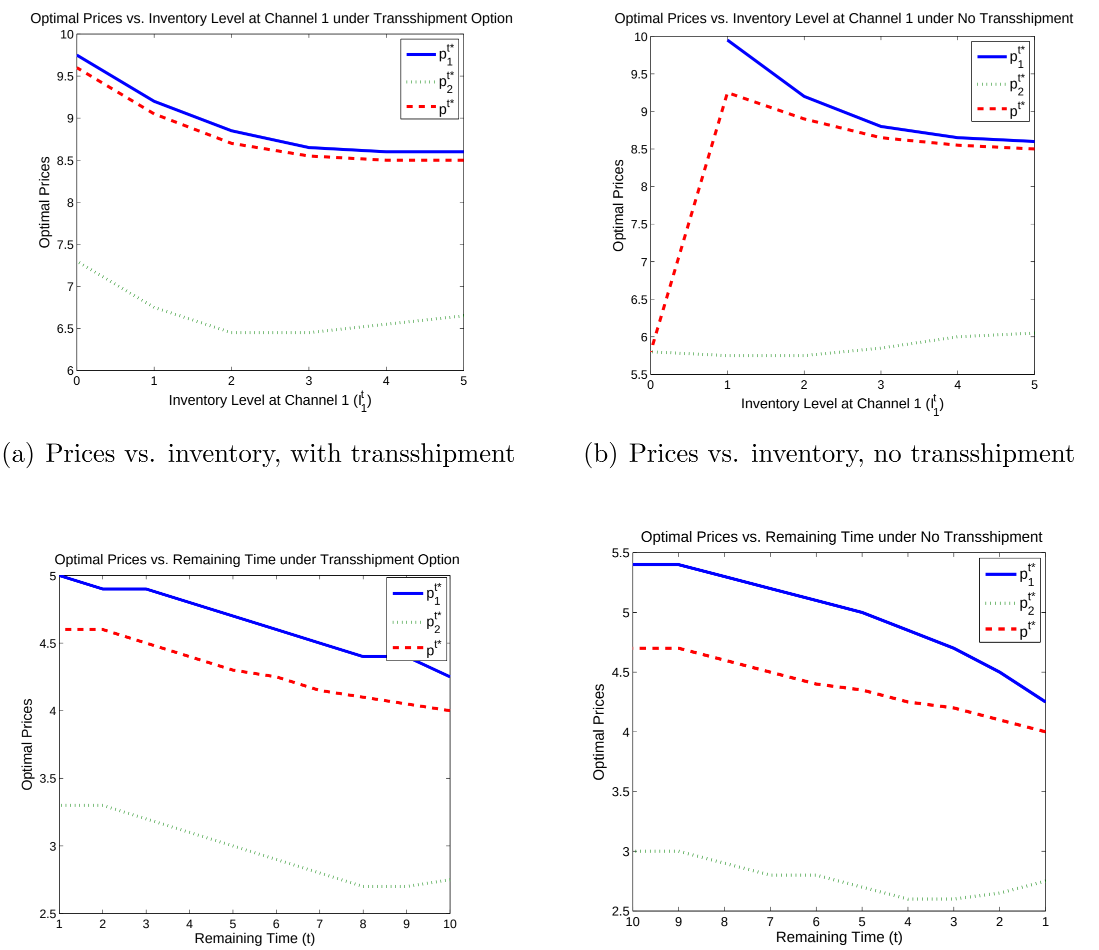

(a) Prices vs. inventory, with transshipment

(b) Prices vs. inventory, no transshipment

(c) Prices vs. remaining time, with transshipment

(d) Prices vs. remaining time, no transshipment

Figure 4.3: Optimal prices vs. inventory level and remaining time: $I_2^T = 2, \lambda^t = 0.8, \beta = 0.95$. In Figure a-b, $T = 5, v_1 = 10, v_2 = 6, c_1^t = 3, c_2^t = 1, m_{12} = 0.5, m_{21} = 0.5$. In Figure c-d, $I_1^T = 1, v_1 = 4, v_2 = 2, c_1^t = 2, c_2^t = 0.5, m_{12} = 0.1, m_{21} = 0.2$.

We have discussed how the transshipment and pricing mechanisms work together to most effectively balance inventory and demand in a dual-channel setting. Our next interest is to compare the effectiveness of the transshipment and price differentiation policies in improving the retailer's profit. To get a sense of how much benefit the retailer can gain from each policy, we employ numerical study to determine the retailer's optimal profits under different scenarios. The results are presented in the next section.

4.5 Numerical Study

We conduct three sets of numerical examples, 200 instances each. Each set is given a different initial inventory level, but the problem parameters are either the same, or randomly chosen based on the same distributions for all three sets. The problem parameters are summarized below.

Set 1: $I_1^T = 2, I_2^T = 3$

Set 2: $I_1^T = 1, I_2^T = 4$

Set 3: $I_1^T = 4, I_2^T = 1$

For all three sets, $T = 5, \beta = 0.95, v_1 \sim U[1,4), v_2 \sim U[1,2), m_{12} = 0.1, m_{21} = 0.2, c_1^t \sim U[0.5, 2.5), c_2^t \sim U[0.5, 1.5), \lambda^t \sim U[0,1)$ for $t = 5, 4, ..., 1$.

Notice that the problem parameters are chosen to simulate a situation where channel 1 is a premium channel (higher customer valuation and higher transaction cost, on average). Hence, for a brick-and-click retailer, channel 1 in our model represents the physical store whereas channel 2 represents the online store. Note also that we keep the total initial inventory the same across the three sets of experiments, and only vary the inventory distribution between the two channels.

Our program randomly generates 200 problem instances for each set, and solves for optimal solutions under the four types of pricing and transshipping policies, listed in Table 4.1. The statistics of the optimal prices in the current period and the profit improvement from using price differentiation and/or transshipment policies are summarized in Table 4.2 and Table 4.5, respectively.

Table 4.2: Statistics for the optimal prices in the current period

		No Transshipment		With Transshipment	
		Mean	Standard Deviation	Mean	Standard Deviation
Price Differentiation	p_1^{T*}	3.47	0.54	3.42	0.53
	p_2^{T*}	2.77	0.32	2.80	0.35
Uniform Pricing	p^{T*}	3.24	0.45	3.21	0.42

Table 4.2a: Optimal prices under Set 1

		No Transshipment		With Transshipment	
		Mean	Standard Deviation	Mean	Standard Deviation
Price Differentiation	p_1^{T*}	3.69	0.54	3.42	0.47
	p_2^{T*}	2.69	0.30	2.80	0.39
Uniform Pricing	p^{T*}	3.34	0.48	3.19	0.38

Table 4.2b: Optimal prices under Set 2

		No Transshipment		With Transshipment	
		Mean	Standard Deviation	Mean	Standard Deviation
Price Differentiation	p_1^{T*}	3.29	0.50	3.33	0.51
	p_2^{T*}	3.09	0.35	2.92	0.39
Uniform Pricing	p^{T*}	3.23	0.37	3.18	0.39

Table 4.2c: Optimal prices under Set 3

First, we discuss the optimal prices in Table 4.2. We observe that in Set 1 and Set 2, the average optimal price at channel 1 under transshipment options is lower than that under no transshipment; while, the average optimal price at channel 2 under transshipment options is higher than that under no transshipment. This is due to the fact that the initial inventory level in Set 1 and Set 2 are such that there is less inventory at channel 1, which is the premium channel. Hence, if the retailer has an option to transship, it is likely that most of the transshipments will be made from channel 2 to channel 1. In response to a higher expected stock level at channel 1 and a lower expected stock level at channel 2 as a result of future transshipments, it is optimal for the retailer to charge a lower price at channel 1 and a higher price at channel 2 in the current period. Furthermore, notice that the gaps between the average optimal prices with and without transshipment options are larger in Set 2 than in Set 1. This is because the initial inventory level at channel 1 in Set 2 is less

	Set 1		Set 2		Set 3	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Benefit of Price Diff. under No Transshipment $\left(\frac{V_{10}^T - V_{00}^T}{V_{00}^T}\right)$	5.18%	5.44%	5.27%	5.11%	2.97%	3.86%
Benefit of Price Diff. under Transshipment $\left(\frac{V_{11}^T - V_{01}^T}{V_{01}^T}\right)$	5.04%	5.55%	4.71%	5.56%	3.93%	5.40%
Benefit of Transshipment under Uniform Pricing $\left(\frac{V_{01}^T - V_{00}^T}{V_{00}^T}\right)$	1.35%	2.64%	8.39%	12.12%	2.27%	3.53%
Benefit of Transshipment under Price Diff. $\left(\frac{V_{11}^T - V_{10}^T}{V_{10}^T}\right)$	1.20%	2.33%	7.69%	10.97%	3.21%	4.88%
Joint Benefit of Price Diff. and Transshipment $\left(\frac{V_{11}^T - V_{00}^T}{V_{00}^T}\right)$	6.41%	5.38%	13.20%	11.01%	6.34%	7.48%

Table 4.3: Benefit of price differentiation and transshipment

than that in Set 1. Hence, transshipments are likely to take place more often in Set 2, resulting in larger effects of future transshipments on the current prices.

Now, consider the optimal prices in Set 3. We find that under transshipment options, the average optimal price at channel 1 is higher while the average optimal price at channel 2 is lower, compared to the optimal prices without transshipment. These results are in contrast with what we observe from Set 1 and Set 2. This is because, unlike Set 1 and Set 2, Set 3 is given an initial inventory level with significantly more inventory at channel 1. In this case, it is much more likely that channel 2 will stock out in the future periods. Hence, if the retailer has an option to transship, most of the transshipments are likely to be made from channel 1 to channel 2. Since the expected stock level at channel 1 is lower whereas the expected stock level at channel 2 is higher under transshipment options, it is optimal for the retailer to charge a higher price at channel 1 and a lower price at channel 2 in the current period.

It is interesting to see that in all three sets of experiments, the average optimal uniform price in the current period is lower with transshipment options than without transshipment. This is mainly driven by the fact that under transshipment options, the retailer can always make a transshipment to the stockout channel when needed in a future period, and still be able to charge a reasonable price for the product, as we have learned from Proposition 4.4 part 2. Since stocking out at a channel in the future periods is not a huge concern, the opportunity cost of selling a unit in the current period (marginal value of a unit of inventory) is lower, compared to under no transshipment. Hence, it is optimal for the retailer to charge a lower price in the current period, when a transshipment is not needed (since the initial inventory level in all three sets are such that both channels carry some inventory), in order to generate more sales earlier on and reduce the risk of overstocking at the end of the selling horizon.

Next, we discuss the profit improvements from using price differentiation and inventory transshipment. From Table 4.5, we see that the benefits of price differentiation in Set 1 and Set 2 are higher than that in Set 3 on average. This is because the initial inventory level in Set 3 is already quite in balance with the arriving demand. Since the customer valuation is generally higher at channel 1, even when a uniform price is offered, more customers are likely to buy from channel 1, where there is more inventory to serve. Hence, there is only a small room for profit improvement using price differentiation. On the other hand, in Set 1 and Set 2, the retailer starts with less inventory at channel 1. If a uniform price is offered, the retailer will quickly run out of inventory at channel 1. Therefore, charging different prices at the two channels (specifically, higher prices at channel 1) has a high potential in helping the retailer more efficiently balance inventory with demand.

The benefit from transshipment is also largely influenced by the initial inventory level. Notice from Table 4.5 that transshipment improves the retailer's profit the most in Set 2, and the least in Set 1. This is because in Set 2, the initial inventory level is notably out of balance as there is a lot less inventory at channel 1, which is the premium channel. In this case, the retailer is in critical need of more inventory at channel 1. Hence, the ability to transfer inventory between the channels greatly improves profits. In Set 3, the initial inventory level is preferable to the retailer as there is a lot more inventory at channel 1. However, since channel 2 has very little inventory, the chance of stocking out is high. Hence, transshipment is still very much necessary in this setting. On the other hand, Set 1 is given an initial inventory level where both channels have almost equal amount of inventory in stock. Since the chance of stocking out is not particularly high at either channel, inventory transfers are not required very often, resulting in small benefits of transshipment.

While both price differentiation and inventory transshipment can significantly improve the retailer's profit, the performance of the two mechanisms can vary, depending on the problem parameters. We note that in the majority of the problem instances (70% in Set 1, 54.5% in Set 2, and 62% in Set 3, out of 200 instances conducted in each set), the benefit of price differentiation is larger than the benefit of transshipment. This is likely to be due to the fact we pointed out earlier that price differentiation is always profitable whereas transshipment may or may not be profitable. Furthermore, we notice that transshipment is especially beneficial when the initial inventory level is such that there is a lot more inventory at channel 2, and channel 1 has high customer valuation but low transaction cost, relative to channel 2. When these conditions do not hold, the benefit of transshipment is generally smaller than the benefit of price differentiation. This is because, when the inventory level is substantially uneven (large difference between the inventory levels at the two channels) in a way that there is less inventory at the high-margin channel, the retailer can resolve the situation more effectively by correcting his inventory position. If he only price differentiates between the channels, he could prevent stocking out at the high-margin channel by charging a high price. However, doing so results in less sales at the high-margin channel and less profits overall. Based on this reason, it is not surprising to observe that the benefit of transshipment is larger than the benefit of price differentiation in more problem instances in Set 2.

Finally, it is worth noting that the benefit from price differentiation and the benefit from transshipment may either substitute or complement each other. To see this, notice from Table 4.5 that in Set 1 and Set 2, the benefit of price differentiation is larger when the retailer does not have an option to transship. Likewise, in Set 1 and Set 2, the benefit of transshipment is larger under uniform pricing than under price differentiation. This implies that the benefit of price differentiation and the benefit of transshipment generally substitute each other under the settings of Set 1 and Set 2. The main driver of this result is the fact that Set 1 and Set 2 are given an initial inventory level that is unfavorable to the retailer (less inventory at the high-margin channel). In such a situation, the retailer would want to rebalance his inventory position by trying to reduce the inventory at channel 2 and maintain or increase the inventory at channel 1. Both price differentiation and transshipment mechanisms can be employed for this same purpose. Hence, the benefit of the two mechanisms substitute each other. However, the opposite result is observed in Set 3, where the benefit of price differentiation is larger when the retailer can transship, and the benefit of transshipment is larger when the retailer can price differentiate. This is because in Set 3, the initial inventory level is favorable for the retailer as there is more inventory at the high-margin channel. In this case, it is in the retailer's best interest to try to maintain the same balance in his inventory position. To most effectively achieve this, the retailer needs to employ both price differentiation and transshipment mechanisms to complement each other's effect. More precisely, the retailer would maintain a high price at channel 1 to extract most surplus. However, doing so increases the chance of having a stock out at channel 2 and a leftover at channel 1 in the future periods. Hence, transshipments will be needed to replenish the stock at channel 2, allowing the retailer to continue to benefit from price differentiation.

4.6 Extension

4.6.1 Ex-post Transshipment

So far, we have assumed that if the retailer decides to make a transshipment in the current period, the transshipment then takes place before a customer arrival, and an arriving customer observes the actual product availability at the channels. In practice, some retailers may have more flexibility in the timing of a transshipment, and hence, are able to delay the transshipment until after the customer makes a decision to purchase. In this case, even when the product is out of stock at a channel, the retailer can announce the product availability at both channels, but only transship the product to the stock-out channel when a customer decides to buy from that channel. If an arriving customer does not buy or buys from the in-stock channel, then the retailer does not need to make a transshipment in that period. For example, if the product is out of stock at the distribution center, used to fulfill demand at the online store, the retailer may continue to display the product as available online, and ship the product from the physical store only after a customer places an online order. We refer to this transshipment policy as "ex-post transshipment," and use x to denote the policy.

To see how the results regarding transshipment (Proposition 4.4 and Theorem 4.1) are affected by the ex-post transshipment policy, we consider an inventory level $\mathbf{I^t} = (0, I_2^t)$, where $I_2^t > 0$, and assume as before that the retailer may transship the product from channel 2 to channel 1 only in the current period. Under the ex-post transshipment policy, if the retailer decides to continue selling at both channels, he makes a unit transshipment from channel 2 to channel 1 only when the customer chooses to buy from channel 1. Hence, the retailer's pricing problem under price differentiation in period t is given by:

$$
\begin{aligned}
\hat{V}_{dx}^{t}(\mathbf{I^t}) &= \max_{\mathbf{p^t}} \Big\{ \lambda^t \mu_1 \left[p_1^t - c_1^t + \beta V_{un}^{t-1}\left(0, I_2^t - 1\right) - m_{21}\right] \\
&\quad + \lambda^t \mu_2 \left[p_2^t - c_2^t + \beta V_{un}^{t-1}\left(0, I_2^t - 1\right)\right] \\
&\quad + \left(\lambda^t \mu_0 + \left(1 - \lambda^t\right)\right) \beta V_{un}^{t-1}\left(0, I_2^t\right) \Big\} \\
&= \max_{\mathbf{p^t}} \Big\{ \beta V_{un}^{t-1}\left(0, I_2^t\right) + \lambda^t \mu_1 \left[p_1^t - c_1^t - \beta \Delta_2^t\left(0, I_2^t\right) - m_{21}\right] \\
&\quad + \lambda^t \mu_2 \left[p_2^t - c_2^t - \beta \Delta_2^t\left(0, I_2^t\right)\right] \Big\} \\
&= \max_{\mathbf{p^t}} \Big\{ \beta V_{un}^{t-1}\left(0, I_2^t\right) + \lambda^t J_{dx}^t(\mathbf{p^t}, \mathbf{I^t}) \Big\} \qquad (4.6) \\
s.t. \quad & p_i^t \geq 0 \text{ for } i \in \{1, 2\}, t = 1, ..., T, \\
& V_{dx}^0(\mathbf{I}) = 0 \text{ for all } \mathbf{I}.
\end{aligned}
$$

The retailer's uniform pricing problem is given by (4.6) with $p_1^t = p_2^t = p^t$.

Comparing $\hat{V}_{dx}^t$ to the retailer's pricing problem with ex-ante transshipment in period t, given by (4.2) with $s_{21}^t = 1, s_{12}^t = 0, \mathbf{Y^t} = (1, I_2^t - 1)$, and $V^{t-1} = V_{un}^{t-1}$, we see that the terms that depend on prices (J^t) under the ex-post transshipment are the same as those under the ex-ante transshipment with $\beta\Delta_1^t(1, I_2^t - 1)$ replaced by $\beta\Delta_2^t(0, I_2^t) + m_{21}$, and $\Delta_2^t(1, I_2^t - 1)$ replaced by $\Delta_2^t(0, I_2^t)$. With these two modifications, the optimal prices under ex-post transshipment can be characterized in the same way as under ex-ante transshipment. Notice also that by replacing $\beta\Delta_1^t(1, I_2^t - 1)$ with $\beta\Delta_2^t(0, I_2^t) + m_{21}$, the transshipment marginal value $\Delta_{12}^t(\mathbf{I^t})$ is modified to $\frac{m_{21}}{\beta}$ under the ex-post transshipment.

The results in Proposition 4.4 and Theorem 4.1 continue to hold under ex-post transshipment with the two term modifications discussed above. For example, the condition when it is optimal to make a transshipment (in this case, to offer the product at both channels) in the current period, given in Proposition 4.4 part 1., is modified from $m_{21} \leq (R_{ut}^{t*} - R_{un}^{t*})\lambda^t + \beta\Delta_{12}^t(\mathbf{I^t})$ to $(R_{ux}^{t*} - R_{un}^{t*})\lambda^t \geq 0$, which reduces to $R_{ux}^{t*} - R_{un}^{t*} \geq 0$ since $\lambda^t \geq 0$. In Theorem 4.1 part 1., the sufficient condition for the optimal price at channel 2 after the transshipment to be higher than the optimal price without transshipment is modified from $\Delta_2^t(1, I_2^t - 1) \geq \Delta_2^t(0, I_2^t)$ to $\Delta_2^t(0, I_2^t) \geq \Delta_2^t(0, I_2^t)$, which is always true. Hence, with ex-post transshipment, the optimal price at channel 2 with the transshipment is always higher than the optimal price without the transshipment.

4.7 Conclusion

Channel-based price differentiation and inventory transshipment are two of the most common mechanisms used by multi-channel retailers to balance inventory and demand arriving at each channel. Price differentiation can be employed to shift the demand from the channel with low inventory to the channel with more inventory in order to prevent overstocking or understocking. On the other hand, inventory transshipment does not affect the demand, but it allows the retailer to directly correct his inventory position by transferring inventory from the channel with more inventory to the channel with less inventory. Since both mechanisms can be implemented to serve the same purpose of balancing the inventory and demand across the channels, we are interested in investigating what factors affect the benefit from each mechanism, when one mechanism is likely to be more effective than the other, and how the two mechanisms influence each other if they are adopted concurrently.

We model a dual-channel retailer's joint dynamic pricing and transshipping problem over a finite horizon. In each period, the retailer decides how much to charge for the product sold at each channel, and whether to make any inventory transshipment between the channels. An arriving customer decides whether to purchase the product from one of the available channels, based on her valuations and the observed prices at the channels. Her purchase decision follows the multinomial logit model.

We characterize the retailer's optimal dynamic pricing and transshipping policies. We find that price differentiation and transshipment policies have different implications on the product sales. The optimal price differentiation policy always results in a larger probability of making a sale in the current period, compared to the optimal uniform pricing policy. On the other hand, the optimal transshipment decision may result in a larger or smaller probability of making a sale in the current period. When price differentiation and inventory transshipment policies are implemented together, we show that transshipment can increase the value of the remaining inventory, and subsequently the price, at the channel from which the transshipment is made. Additionally, since a transshipment is required in a stock-out situation for the benefit of price differentiation to be realized, we find that the benefit of price differentiation may decrease in the transshipment cost when the cost is large.

Our numerical study helps compare the benefit of price differentiation and inventory transshipment under various scenarios. The results show that the benefit of price differentiation is generally larger than the benefit of transshipment. However, transshipment can be more effective than price differentiation when the retailer holds significantly less inventory at the high-margin channel. When adopted together, the price differentiation and inventory transshipment mechanisms may either substitute or complement each other. When the retailer's inventory position requires some correction, either price differentiation or inventory transshipment can be used. Hence, the two mechanisms substitute each other in this case. On the other hand, when the current inventory balance at the channels should be maintained in the same proportion, the two mechanisms work together complementarily to improve the retailer's profit.

This study considers a joint dynamic pricing and inventory problem in a setting and context that are different from existing literature. In particular, we focus on the price differentiation and inventory transhipment mechanisms in a dual-channel retailing, which are widely observed in practice, but not well studied in the literature. Our results help explain how each mechanism works in a dual-channel environment, how the two mechanisms are similar or different, and when one should be chosen over the other, or both are required. Although our current model assumes a specific customer choice model and a simple retailer's cost structure, we believe that the managerial insights obtained from this study can be applied to many realistic dual-channel situations to help make appropriate pricing and inventory decisions.

There are several directions for future research on this topic that would further provide valuable contribution to the literature. A natural extension of the current model is to consider a multi-channel setting (more than two channels/stores). Having multiple channels is likely to result in a significantly higher complexity level of the problem, especially on the transshipment decisions. However, certain modeling assumptions may help reduce the decision space and make it possible to achieve some interesting results. One could also extend the model to consider a different customer choice model. For example, the heterogeneity in customer valuation for the product sold at each channel may be distributed over a general distribution; different customers may react differently when they find that their preferred channel stocks out; observing different prices at the channels may have certain effects on how customers evaluate their purchases. Lastly, the model can be extended to incorporate additional costs and operational constraints that could affect the retailer's pricing and transshipping decisions. These include, for instance, costs associated with managing different prices at the channels, fixed transshipment costs, inventory holding costs, inventory holding capacity, transshipment lead times.

CHAPTER 5

Conclusion

While there are many pricing mechanisms used in practice, different mechanisms may not deliver the same benefits to sellers under different business situations. Hence, choosing an appropriate mechanism can be a challenging problem for a seller. This dissertation aims to provide insights into how different pricing mechanisms affect the buyer's behavior and the seller's profitability to ultimately help business managers make effective operational decisions.

To address various types of pricing problems commonly occurring in a supply chain, this dissertation considers three different seller's problems in a series of three essays.

The first essay (Chapter 2) examines whether and when the supplier benefits from more accurate demand forecasts obtained by the buyer under different contract structures. An important finding of the essay is that there indeed exist contracts under which the supplier can always benefit from the buyer's more accurate demand forecasts. However, depending on how certain the supplier is about the buyer's forecasting capability, the contract structure may be more sophisticated than what is common in practice. This finding complements the existing supply chain contract literature which has considered simpler forms of contracts and reported that the supplier's profit can be hurt by the buyer's improved demand forecasts.

The second essay (Chapter 3) discusses how price promotions influence purchase decisions of different types of consumers, and which type of promotions is most profitable to the seller under which market situations. Two forms of conditional discounts, all-unit and fixed-amount discount, are considered. This study shows that both all-unit and fixed-amount discount are equally effective in inducing deal-prone consumers to overspend. However, when consumers are heterogenous in their willingness to pay for the product, one form of discounts may outperform the other. Hence, the seller's choice of discount types can significantly determine how much profit he obtains. The findings from this essay provide guidance on when it is optimal to offer a conditional discount, and which form of discounts to offer to maximize profits.

The third essay (Chapter 4) considers a dual-channel retailer who can employ either price differentiation or inventory transshipment to balance inventory and demand at each channel. This study investigates what factors affect the benefit from each mechanism, when one mechanism is likely to be more effective than the other, and how the two mechanisms influence each other if they are implemented concurrently. It is shown that while price differentiation always increases sales volume, inventory transshipment may sometimes result in a smaller sales volume in the current period. In terms of benefits to the retailer, whether the mechanism of price differentiation or inventory transshipment results in a larger profit improvement, and whether the benefit of the two mechanisms, when adopted together, substitute or complement each other primarily depend on the retailer's initial inventory position. The results obtained from this study can help retailers make good judgement when implementing joint pricing and inventory policies in a dual-channel environment.

On the academic side, this dissertation studies new research problems and provides original results that contribute to the existing Operations Management literature as well as other related fields such as Marketing and Economics. On the practical side, the insights discussed in this work can help managers craft their strategies to achieve successful operations in many realistic business situations.

APPENDICES

APPENDIX A

Additional Results and Proofs for Chapter 2

A.1 Characterization of Optimal Contracts

A.1.1 Early Static Contract

The supplier's optimization problem with early static contract is given by

$$\begin{aligned}
\max_{\mathbf{q},\mathbf{t}} \; & \sum_{i\in\{L,H\}} p_i^1(-c_1 q_i + t_i) \\
s.t. \; & \sum_{j\in\{L,H\}} p_{ij}^1[\Gamma(D_j, q_i) - t_i] \geq 0,\; i = \{L, H\} \\
& \sum_{j\in\{L,H\}} p_{ij}^1[\Gamma(D_j, q_i) - t_i] \geq \sum_{j\in\{L,H\}} p_{ij}^1[\Gamma(D_j, q_{-i}) - t_{-i}],\; i = \{L, H\} \\
& q_i, t_i \geq 0 \; i = \{L, H\}
\end{aligned}$$

One can show that the supplier's problem is equivalent to the following reduced problem. (We provide the detailed discussion of how to obtain this reduced form when we study the supplier's problem with dynamic contract, the derivation for early and late static contracts are similar.)

$$\begin{aligned}
\max_{\mathbf{q},\mathbf{t}} \quad & \sum_{i\in\{L,H\}} p_i^1(-c_1 q_i + t_i) \qquad\qquad (A.1)\\
s.t. \quad & \sum_{j\in\{L,H\}} p_{Lj}^1[\Gamma(D_j, q_L) - t_L] \geq 0,\\
& \sum_{j\in\{L,H\}} p_{Hj}^1[\Gamma(D_j, q_H) - t_H] \geq \sum_{j\in\{L,H\}} p_{Hj}^1[\Gamma(D_j, q_L) - t_L],\ i = \{L, H\}\\
& q_H \geq q_L
\end{aligned}$$

From this, we can characterize the optimal early static contract, which is presented in Proposition A.1.

Proposition A.1. *The optimal early static contract quantities can be obtained by solving the following equations.*

$$(1 - (1 + p_H^1)\theta_1)\Gamma'(D_H, q_L) + ((1 + p_H^1)\theta_1 - p_H^1)\Gamma'(D_L, q_L) - c_1 p_L^1 = 0$$
$$\theta_1\Gamma'(D_H, q_H) + (1 - \theta_1)\Gamma'(D_L, q_H) - c_1 = 0$$

The supplier charges the retailer $t_L = (1 - \theta_1)\Gamma(D_H, q_L) + \theta_1\Gamma(D_L, q_L)$, $t_H = \theta_1\Gamma(D_H, q_H) - (2\theta_1 - 1)\Gamma(D_H, q_L) + (1 - \theta_1)\Gamma(D_L, q_H) + (2\theta_1 - 1)\Gamma(D_L, q_L)$. *The supplier always produces exactly the quantity selected by the buyer.*

Proof of Proposition A.1: Consider the reduced early static optimization problem given by (A.1). The optimal contract quantities can be obtained from solving the first-order conditions because $\Gamma(D, q)$ satisfies Property 4 in Assumption 1, which implies the unimodality of the supplier's profit in contract quantities. If $\Gamma(D, q)$ also satisfies $\Gamma'(D_i, q) \geq \Gamma'(D_j, q)$ for any $D_i \succcurlyeq D_j$, as is the case for the buyer's revenue in both exogenous price and endogenous price models, the solution of q_L and q_H that satisfy the first-order conditions always satisfy the monotonicity constraint $q_L \leq q_H$. To see this, let Π_{ES} be the objective function of this optimization problem. From the first-order conditions given in the proposition, we have,

$$\begin{aligned}\frac{\partial \Pi_{ES}}{\partial q_L}|_{q_H} &= (1-\theta_1)\Gamma'(D_H, q_H) + \theta_1 \Gamma'(D_L, q_H) - c_1 \\ &\le \theta_1 \Gamma'(D_H, q_H) + (1-\theta_1)\Gamma'(D_L, q_H) - c_1 = 0\end{aligned}$$

where the inequality is due to $\theta_1 \ge \max\{p_L, p_H\} \ge \frac{1}{2}$ and $\Gamma'(D_H, q_H) \ge \Gamma'(D_L, q_H)$, and the last equality follows from the first-order condition with respect to q_H. Since the objective function is unimodal in q_L, $\frac{\partial \Pi_{ES}}{\partial q_L}|_{q_H} \le 0$ implies $q_H \ge q_L$.

The supplier produces the exact quantity selected by the buyer in period 1 because the production is made after the screening.

A.1.2 Late Static Contract

The supplier's optimization problem with late static contract is given by

$$\begin{aligned}\max_{\mathbf{q},\mathbf{t},\rho} \ -c_1\rho \quad &+ \sum_{i\in\{L,H\}} p_i^1 \sum_{j\in\{L,H\}} p_{ij}(t_j - c_2(q_j-\rho)^+) \\ s.t. \sum_{k\in\{L,H\}} & p_{jk}^2 \Gamma(D_k, q_j) \ge t_j \\ \sum_{k\in\{L,H\}} & p_{jk}^2 \Gamma(D_k, q_j) - t_j \ge \sum_{k\in\{L,H\}} p_{jk}^2 \Gamma(D_k, q_{-j}) - t_{-j} \\ & \rho, q_j, t_j \ge 0, \ j = \{L, H\}\end{aligned}$$

The reduced optimization problem under late static contract is given by

$$\begin{aligned}\max_{\mathbf{q},\ \mathbf{t},\rho} \ -c_1\rho \quad &+ \sum_{i\in\{L,H\}} p_i^2 (t_i - c_2(q_i-\rho)^+) \\ s.t. \sum_{j\in\{L,H\}} & p_{Lj}^2 [\Gamma(D_j, q_L) - t_L] \ge 0, \\ \sum_{j\in\{L,H\}} & p_{Hj}^2 [\Gamma(D_j, q_H) - t_H] \ge \sum_{j\in\{L,H\}} p_{Hj}^2 [\Gamma(D_j, q_L) - t_L], \\ & q_H \ge q_L\end{aligned} \tag{A.2}$$

The optimal late static contract is characterized in Proposition A.2.

Proposition A.2. *The supplier's period 1 optimal production under late static contract is*

$$\rho^*(\boldsymbol{q}) = \begin{cases} q_H & \text{if } \frac{c_1}{c_2} \le p_H^2 \\ q_L & \text{if } \frac{c_1}{c_2} > p_H^2 \end{cases}$$

The optimal late static contract quantities can be obtained by solving the following equations:

$$(1-(1+p_H^2)\theta_2)\Gamma'(D_H, q_L) + ((1+p_H^2)\theta_2 - p_H^2)\Gamma'(D_L, q_L) = (c_1 - p_H^2 c_2)^+$$
$$p_H^2[\theta_2\Gamma'(D_H, q_H) + (1-\theta_2)\Gamma'(D_L, q_H)] = \min\{c_1, p_H^2 c_2\}$$

If $q_L > q_H$, it is optimal to bunch the low- and high-type quantity and offer $q_L = q_H = \bar{q}$ which satisfies $(1-\theta_2)\Gamma'(D_H, \bar{q}) + \theta_2\Gamma'(D_L, \bar{q}) - c_1 = 0$.

The supplier charges the buyer $t_L = (1-\theta_2)\Gamma(D_H, q_L) + \theta_2\Gamma(D_L, q_L)$, $t_H = \theta_2\Gamma(D_H, q_H) - (2\theta_2 - 1)\Gamma(D_H, q_L) + (1-\theta_2)\Gamma(D_L, q_H) + (2\theta_2 - 1)\Gamma(D_L, q_L)$.

Proof of Proposition A.2: We know from (A.2) that $q_H \ge q_L$ holds at the optimal solution. If the expected saving is greater than the loss, $p_H^2 c_2 \ge c_1$, the supplier improves his expected profit by producing the larger quantity of q_H in period 1.

The transfer payments are derived from the binding constraints stated above. The unimodality of $-\Gamma''(D, q)$ implies the concavity of the supplier's profit function in the contract quantities. Plugging in the expressions of the optimal transfer payments into the objective function, we can use first-order conditions to derive the optimal contract quantities. First, we solve the first-order conditions ignoring the constraints $q_H \ge q_L$. However, if the inequality $q_H \ge q_L$ is violated, then we know that bunching is optimal for the two quantities. Hence, we replace q_L and q_H by $\bar{q}$ in the supplier's profit, and solve the first-order condition with respect to $\bar{q}$ to obtain the optimal solution of $q_L = q_H = \bar{q}$.

A.1.3 Dynamic Contract

We will first show that the supplier's optimization problem for dynamic contract can be reduced to the following problem.

$$\max_{\mathbf{q},\mathbf{t},\rho} \left\{ \sum_{i=\{L,H\}} p_i^1(-c_1\rho_i + t_i) + \sum_{i\in\{L,H\}} p_i^1 \sum_{j\in\{L,H\}} p_{ij}[t_{ij} - c_2(q_i + q_{ij} - \rho_i)^+] \right\} \tag{A.3}$$

$$s.t. \quad \sum_{j\in\{L,H\}} p_{Lj} \sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_L + q_{Lj}) - t_{Lj}] \geq t_L \text{ (Period 1 Participation Constraints)}$$

$$\sum_{j\in\{L,H\}} p_{Hj} \sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_H + q_{Hj}) - t_{Hj}] - t_H \geq \text{ (Period 1 Incentive Constraint)}$$

$$p_{HH} \sum_{k\in\{L,H\}} p_{Hk}^2[\Gamma(D_k, q_L + q_{LH}) - t_{LH}] + p_{HL} \sum_{k\in\{L,H\}} p_{Lk}^2[\Gamma(D_k, q_L + q_{LL}) - t_{LL}] - t_L$$

$$\sum_{k\in\{L,H\}} p_{Lk}^2[\Gamma(D_k, q_i + q_{iL}) - t_{iL}] \geq \sum_{k\in\{L,H\}} p_{Lk}^2\Gamma(D_k, q_i),\ i = \{L,H\}$$

(Period 2 Participation Constraints)

$$\sum_{k\in\{L,H\}} p_{Hk}^2[\Gamma(D_k, q_i + q_{iH}) - t_{iH}] \geq \sum_{k\in\{L,H\}} p_{Hk}^2[\Gamma(D_k, q_i + q_{iL}) - t_{iL}],\ i = \{L,H\}$$

(Period 2 Incentive Constraints)

$$\rho_i, q_i, q_{ij}, t_i, t_{ij} \geq 0 \ i = \{L,H\},\ j = \{L,H\} \text{ (Nonnegativity Constraints)}$$

$$q_{HH} \geq q_{HL}, q_{LH} \geq q_{LL}$$

Let PC_i and PC_{ij} $i,j \in \{L,H\}$ be the participation constraints of buyer type i and type ij in period 1 and 2, respectively. Let $IC_{ij,\ i(-j)/(-i)k,(-i)l}$ $i,j,k,l \ \in \{L,H\}$ be the period 1 incentive constraint of the type i buyer not to deviate to contract $(-i)k$ in period 2 when she observes ij, and not to deviate to $(-i)l$ when she observes $i(-j)$. Let IC_{ij} be the period 2 incentive constraint for the buyer observing ij not to deviate to $i(-j)$.

Period 1 PC_H is implied by period 1 PC_L, period 1 $IC_{HH,HL/LH,LL}$, and period 2 IC_{LH} as follows: $\sum_{j\in\{L,H\}} p_{Hj} \sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_H + q_{Hj}) - t_{Hj}] - t_H \geq p_{HH} \sum_{k\in\{L,H\}} p_{Hk}^2[\Gamma(D_k, q_L + q_{LH}) - t_{LH}] + p_{HL} \sum_{k\in\{L,H\}} p_{Lk}^2[\Gamma(D_k, q_L + q_{LL}) - t_{LL}] - t_L \geq p_{LH} \sum_{k\in\{L,H\}} p_{Hk}^2[\Gamma(D_k, q_L + q_{LH}) - t_{LH}] + p_{LL} \sum_{k\in\{L,H\}} p_{Lk}^2[\Gamma(D_k, q_L + q_{LL}) - t_{LL}] - t_L \geq 0$. The first inequality is period 1 $IC_{HH,HL/LH,LL}$. The second inequality follows from period 2 IC_{LH} and the fact that $p_{HH} \geq p_{LH}$. The third inequality is period 1 PC_L. Next, period 1 $IC_{HH,HL/LL,LL}$ is implied by period 1 $IC_{HH,HL/LH,LL}$ and period 2 IC_{LH} since $p_{HH}\sum_{k\in\{L,H\}} p^2_{Hk}[\Gamma(D_k, q_L+q_{LH})-t_{LH}]+p_{HL}\sum_{k\in\{L,H\}} p^2_{Lk}[\Gamma(D_k, q_L+q_{LL})-t_{LL}]-t_L \geq p_{HH}\sum_{k\in\{L,H\}} p^2_{Hk}[\Gamma(D_k, q_L+q_{LL})-t_{LL}]+p_{HL}\sum_{k\in\{L,H\}} p^2_{Lk}[\Gamma(D_k, q_L+q_{LL})-t_{LL}]-t_L$.

Similarly, period 1 $IC_{HH,HL/LH,LH}$ is implied by period 1 $IC_{HH,HL/LH,LL}$ and period 2 IC_{LL}; period 1 $IC_{HH,HL/LL,LH}$ is implied by period 1 $IC_{HH,HL/LH,LL}$ and period 2 IC_{LH} and IC_{LL}.

For the low-type period 1 incentive constraints, similar to high type period 1 incentive constraints, we can show that $IC_{LH,LL/HH,HH}$ is implied by $IC_{LH,LL/HH,HL}$ and period 2 IC_{HL}; $IC_{LH,LL/HL,HL}$ is implied by $IC_{LH,LL/HH,HL}$ and period 2 IC_{HH}; and $IC_{LH,LL/HL,HH}$ is implied by period 2 IC_{HL} and IC_{HH}.

Period 2 constraints can be reduced as follows. PC_{HH} is implied by PC_{HL}, IC_{HH}, and $p^2_{HH} \geq p^2_{LH}$. PC_{LH} is implied by PC_{LL}, IC_{LH}, and $p^2_{HH} \geq p^2_{LH}$. Adding IC_{HH} and IC_{HL}, we see that IC_{HL} is implied by IC_{HH} and $q_{HH} \geq q_{HL}$. Similarly, adding IC_{LH} and IC_{LL}, we see that IC_{LL} is implied by IC_{LH} and $q_{LH} \geq q_{LL}$.

Finally, note that at the optimal solution, period 1 PC_L, period 1 $IC_{HH,HL/LH,LL}$, period 2 IC_{HH}, and period 2 IC_{LH} must be binding, or else the supplier can profitably increase t_L, t_H, t_{HH}, and t_{LH}, correspondingly. Notice that the binding period 1 PC_L together with $p_{LH} \leq p_{HH}$ and $q_{HH} \geq q_{HL}$ implies $IC_{LH,LL/HH,HL}$. Hence, all four period 1 low-type incentive constraints can be removed.

Proof of Proposition 2.1: Suppose an optimal dynamic contract is given by $DC := (q_i, t_i)\{(q_{iH}, t_{iH}), (q_{iL}, t_{iL})\}$, $i = L, H$. Now, consider another contract, $DC' := (0, t_i)\{(q_i+q_{iH}, t_{iH}), (q_i+q_{iL}, t_{iL})\}$, $i = L, H$. Since the total quantities and payments at both DC and DC' are the same for all types, the two contracts result in the same profits to the supplier and the buyer if the buyer chooses the contract meant for her type. We will show that DC' satisfies all the constraints, so that it is also an optimal contract. To see this, observe from the supplier's reduced optimization problem given by (A.3) that all constraints of DC and DC' are the same except for period 2 participation constraints. Notice however that the period 2 participation constraints of DC' are implied by the constraints of DC since $\sum_{k\in\{L,H\}} p^2_{Lk}[\Gamma(D_k, q_i + q_{iL}) - t_{iL}] \geq \sum_{k\in\{L,H\}} p^2_{Lk}\Gamma(D_k, q_i) \geq \sum_{k\in\{L,H\}} p^2_{Lk}\Gamma(D_k, 0),\ i = \{L, H\}$ by Property 2 in Assumption 1. Hence, DC' satisfies all the constraints and it is equivalent to DC.

Proof of Proposition 2.2: By Proposition 2.1, there exists an optimal dynamic contract in the following form $DC := (0, t_i)\{(q_{iH}, t_{iH}), (q_{iL}, t_{iL})\},\ i = L, H$. Now, consider the following dynamic contract, $DC' := (0, t_i + t_{iL})\{(q_{iH}, t_{iH} - t_{iL}), (q_{iL}, 0)\},\ i = L, H$.

Since the total quantities and payments in both DC and DC' are the same for all types, the two contracts result in the same profits to the supplier and the buyer if the buyer chooses the contract meant for her type. We will show that period 1 and period 2 constraints hold under DC'. Period 1 participation and incentive constraints of DC' are the same as those of DC because the total transfers $t_i + t_{ij}$ from the buyer to the supplier are the same under both contracts. Since period 2 transfer payments of DC' are smaller than those of DC, period 2 participation constraints of DC' are implied by period 2 participation constraints of DC. Period 2 incentive constraints of DC' are equivalent to those of DC. Therefore, DC' is a feasible contract and it is equivalent to the optimal dynamic contract DC.

Proof of Proposition 2.3: Suppose that in period 1 the buyer chooses type i contract. Then, the supplier produces at least q_{iL} units in period 1. We know that $q_{iH} \geq q_{iL},\ i = L, H$ in a dynamic contract from (A.3). Therefore, if the cost of producing an additional unit at c_1 is less than the expected saving, $p_{iH}c_2$, then the supplier produces exactly q_{iH} in period 1. The optimal transfer payments are obtained from the binding constraints. Since $\Gamma(D, q)$ satisfies Property 4 in Assumption 1, the supplier's profit is unimodal in the contract quantities. Hence, the optimal contract quantities can be derived from solving the first-order conditions, provided in the proposition. If any of the two monotonicity constraints of the contract quantities $q_{iH} \geq q_{iL}$ is violated, we replace the two quantities with the same variable ($\bar{q}_i = q_{iH} = q_{iL}$) and resolve the first-order condition. This case corresponds to bunching of contract quantities.

A.2 Contract Preferences

Proof of Theorem 2.1: We consider the reduced problem for the dynamic contract given in (A.3). Let $S^*(\theta_2) := \{q_i^*(\theta_2), q_{ij}^*(\theta_2), \rho_i^*(\theta_2), t_i^*(\theta_2), t_{ij}^*(\theta_2)\}, i, j \in \{L, H\}$ denote an optimal solution to the supplier's optimization problem under dynamic contract when the buyer's second period information accuracy is θ_2, with $q_i^* = 0$ and $t_{iL}^* = 0, i \in \{L, H\}$. Let $\Pi_{DC}^*(\theta_2) := \Pi_{DC}(S^*(\theta_2), \theta_2)$ denote the corresponding supplier's optimal profit. We will show that there exists a solution $S(\theta_2'')$ such that $\Pi_{DC}^*(\theta_2'') \geq \Pi_{DC}(S(\theta_2''), \theta_2'') \geq \Pi_{DC}^*(\theta_2')$, for $\theta_2'' > \theta_2'$.

Suppose an optimal solution at $\theta_2 = \theta_2'$ is given by $S^*(\theta_2') = \{q_{ij}^*(\theta_2'), \rho_i^*(\theta_2'), t_i^*(\theta_2'), t_{iH}^*(\theta_2')\}, i, j \in \{L, H\}$. We construct a solution $S(\theta_2'') = \{q_{ij}^*(\theta_2'), \rho_i^*(\theta_2'), t_i^*(\theta_2'), t_{iH}(\theta_2'')\}, i, j \in \{L, H\}$ with

$$
\begin{aligned}
t_{LH}(\theta_2'') &= \theta_2''[\Gamma(D_H, q_{LH}^*(\theta_2')) - \Gamma(D_H, q_{LL}^*(\theta_2'))] \\
&\quad +(1 - \theta_2'')[\Gamma(D_L, q_{LH}^*(\theta_2')) - \Gamma(D_L, q_{LL}^*(\theta_2'))] \qquad \text{(A.4)} \\
t_{HH}(\theta_2'') &= \theta_2''[\Gamma(D_H, q_{HH}^*(\theta_2')) - \Gamma(D_H, q_{HL}^*(\theta_2'))] \\
&\quad +(1 - \theta_2'')[\Gamma(D_L, q_{HH}^*(\theta_2')) - \Gamma(D_L, q_{HL}^*(\theta_2'))]
\end{aligned}
$$

Notice that we adopt the same contract quantities, period 1 production, and transfer payments from the optimal contract for $\theta_2 = \theta_2'$ except for the period 2 transfer payments $t_{LH}(\theta_2'')$ and $t_{HH}(\theta_2'')$. We will first show that this solution is feasible for the problem with $\theta_2 = \theta_2''$.

Since $t_L^*(\theta_2')$ and $t_H^*(\theta_2')$ are optimal for θ_2', by Proposition 2.3 they must satisfy

$$\begin{aligned} t_L^*(\theta_2') &= (1-\theta_1)\Gamma(D_H, q_{LL}^*(\theta_2')) + \theta_1\Gamma(D_L, q_{LL}^*(\theta_2')) \\ t_H^*(\theta_2') &= \theta_1\Gamma(D_H, q_{HL}^*(\theta_2')) + (1-\theta_1)\Gamma(D_L, q_{HL}^*(\theta_2')) \\ &\quad -(2\theta_1 - 1)[\Gamma(D_H, q_{LL}^*(\theta_2')) - \Gamma(D_L, q_{LL}^*(\theta_2'))] \end{aligned}$$

Both the period 1 low-type participation constraint and the period 1 high-type incentive constraint $IC_{HH,HL/LH,LL}$ are reduced to $0 \geq 0$. Hence, it follows immediately that these constraints are satisfied when $\theta_2 = \theta_2''$. The period 2 participation constraints are satisfied since $q_i^* = 0$. Additionally, the period 2 IC_{HH} and IC_{LH} constraints are satisfied by our choice of $t_{HH}(\theta_2'')$ and $t_{LH}(\theta_2'')$ as provided above. Therefore, $S(\theta_2'')$ is a feasible solution.

Next, we will show $\Pi_{DC}(S(\theta_2''), \theta_2'') \geq \Pi_{DC}^*(\theta_2')$. Since the contract quantities, period 1 production, and period 1 transfer payments in $S(\theta_2'')$ are chosen to be identical to those in $S^*(\theta_2')$, and hence are independent of θ_2, we will drop the script θ_2 from these quantities for notational simplicity. We will write only t_{HH} and t_{LH}, which are given by (A.4), as functions of θ_2. Then, the supplier's profit function with $S(\theta_2) = \{q_{ij}^*(\theta_2'), \rho_i^*(\theta_2'), t_i^*(\theta_2'), t_{iH}(\theta_2)\}, i, j \in \{L, H\}$ with t_{iH} given by (A.4) can be written as $\hat{\Pi}_{DC}(\theta_2) = p_H^1[p_{HH}(t_{HH}(\theta_2) - c_2(q_{HH} - \rho_H)^+) + t_H - c1\rho_H] + p_L^1[p_{LH}(t_{LH}(\theta_2) - c_2(q_{LH} - \rho_L)^+) + t_L - c_1\rho_L]$. Notice that $\hat{\Pi}_{DC}(\theta_2'') = \Pi_{DC}(S(\theta_2''), \theta_2'')$ and $\hat{\Pi}_{DC}(\theta_2') = \Pi_{DC}^*(\theta_2')$. Hence, we can show $\Pi_{DC}(S(\theta_2''), \theta_2'') \geq \Pi_{DC}^*(\theta_2')$ by showing $\frac{d\hat{\Pi}_{DC}(\theta_2)}{d\theta_2} \geq 0$ for any $\theta_2 \in (\theta_2', 1]$, where $\frac{d\hat{\Pi}_{DC}(\theta_2)}{d\theta_2} = p_H^1[\frac{-(2\theta_1-1)}{(2\theta_2-1)^2}(t_{HH}(\theta_2) - c_2(q_{HH} - \rho_H)^+) + \frac{(\theta_1+\theta_2-1)}{(2\theta_2-1)}(\Gamma(D_H, q_{HH}) - \Gamma(D_H, q_{HL}) - \Gamma(D_L, q_{HH}) + \Gamma(D_L, q_{HL}))] + p_L^1[\frac{(2\theta_1-1)}{(2\theta_2-1)^2}(t_{LH}(\theta_2) - c_2(q_{LH} - \rho_L)^+) + \frac{(\theta_2-\theta_1)}{(2\theta_2-1)}(\Gamma(D_H, q_{LH}) - \Gamma(D_H, q_{LL}) - \Gamma(D_L, q_{LH}) + \Gamma(D_L, q_{LL}))]$.

Notice that by Property 3 in Assumption 1, the term $(\Gamma(D_H, q_{HH}) - \Gamma(D_H, q_{HL}) - \Gamma(D_L, q_{HH}) + \Gamma(D_L, q_{HL}))$ is nonnegative because $q_{HH} \geq q_{HL}$ and $D_H \succcurlyeq D_L$. Similarly, $(\Gamma(D_H, q_{LH}) - \Gamma(D_H, q_{LL}) - \Gamma(D_L, q_{LH}) + \Gamma(D_L, q_{LL})) \geq 0$ because $q_{LH} \geq q_{LL}$

and $D_H \succcurlyeq D_L$. The term $t_{HH}(\theta_2) - c_2(q_{HH} - \rho_H)^+$ and $t_{LH}(\theta_2) - c_2(q_{LH} - \rho_L)^+$ are also nonnegative because $t_{HH}(\theta_2)$ and $t_{LH}(\theta_2)$ are increasing in θ_2, implying $t_{iH}(\theta_2) > t^*_{iH}(\theta'_2) \geq c_2(q_{iH} - \rho_L)^+, i \in \{L, H\}$, where the second inequality follows from the fact that $t^*_{iH}(\theta'_2)$ is optimal for q_{iH} and ρ_i, and hence, must be profitable for the supplier to sell the iH contract in period 2. Given that $\theta_2 \geq \theta_1 \geq \max\{p_L, 1-p_L\}$, $\theta_2 \geq \theta_1 \geq \frac{1}{2}$. Hence, all of the following terms $2\theta_1 - 1, 2\theta_2 - 1, \theta_1 + \theta_2 - 1$, and $\theta_2 - \theta_1$ are nonnegative. Then, the only possibly negative term in $\frac{d\hat{\Pi}_{DC}(\theta_2)}{d\theta_2}$ is the first term.

Observe from Proposition 2.3 that if $\frac{c_1}{c_2} \leq p_{HH}$, then $\rho_H = q_{HH}$ and the optimal q_{HL} is the solution of $[\frac{(1-\theta_2)(\theta_2-\theta_1)}{(2\theta_2-1)}]\Gamma'(D_H, q_{HL}) + [\frac{\theta_2(\theta_2-\theta_1)}{(2\theta_2-1)}]\Gamma'(D_L, q_{HL}) = 0$. Since the left hand side of the equation is always nonnegative, in this case it is optimal to offer the maximum possible quantity for q_{HL}, which is $q_{HL} = q_{HH} = \bar{q}_H$. In other words, whenever $\frac{c_1}{c_2} \leq p_{HH}$, bunching is optimal for the high-type contract, leading to $t_{HH}(\theta_2) - c_2(q_{HH} - \rho_H)^+ = 0$. Then, it follows that $\frac{d\hat{\Pi}_{DC}(\theta_2)}{d\theta_2} \geq 0$.

When $\frac{c_1}{c_2} > p_{HH}$, $\rho_H = q_{HL}$. This also implies $\rho_L = q_{LL}$ since $p_{LH} \leq p_{HH} < \frac{c_1}{c_2}$; the first inequality follows from $\theta_1 \geq \frac{1}{2}$. We will prove by contradiction that $\frac{d\hat{\Pi}_{DC}(\theta_2)}{d\theta_2} \geq 0$. $\frac{d^2\hat{\Pi}_{DC}(\theta_2)}{d\theta_2^2}$ is given by $\frac{d^2\hat{\Pi}_{DC}(\theta_2)}{d\theta_2^2} = (\frac{-4}{2\theta_2-1})\frac{d\hat{\Pi}_{DC}(\theta_2)}{d\theta_2} + \frac{4\theta_2}{(2\theta_2-1)^2}[\Gamma(D_H, q_{HH}) - \Gamma(D_H, q_{HL}) - \Gamma(D_L, q_{HH}) + \Gamma(D_L, q_{HL}) + \Gamma(D_H, q_{LH}) - \Gamma(D_H, q_{LL}) - \Gamma(D_L, q_{LH}) + \Gamma(D_L, q_{LL})]$. Suppose $\frac{d\hat{\Pi}_{DC}(\theta_2)}{d\theta_2} < 0$ for some $\hat{\theta}_2$. Then, from $\frac{d^2\hat{\Pi}_{DC}(\theta_2)}{d\theta_2^2}$, we must have that $\frac{d^2\hat{\Pi}_{DC}(\theta_2)}{d\theta_2^2} > 0$ at $\hat{\theta}_2$ since the first term is positive and the second term is nonnegative. Given this, whenever $\frac{d\hat{\Pi}_{DC}(\theta_2)}{d\theta_2} < 0$, we must have $\frac{d^2\hat{\Pi}_{DC}(\theta_2)}{d\theta_2^2} > 0$, which implies that $\frac{d\hat{\Pi}_{DC}(\theta_2)}{d\theta_2}$ is minimized at the minimum $\theta_2 \leq \hat{\theta}_2$. Since we are considering the case of $\frac{c_1}{c_2} > p_{HH} = \frac{\theta_1+\theta_2-1}{2\theta_2-1}$, and since p_{HH} is monotonically decreasing in θ_2, the minimum θ_2 corresponds to the maximum p_{HH}, which is $p_{HH} = \frac{c_1}{c_2}$. But we have seen earlier that at this point it is optimal to bunch the high-type contracts, resulting in $\frac{d\hat{\Pi}_{DC}(\theta_2)}{d\theta_2} \geq 0$, which is a contradiction. Thus, there exists no such $\hat{\theta}_2$ where $\frac{d\hat{\Pi}_{DC}(\theta_2)}{d\theta_2}\big|_{\hat{\theta}_2} < 0$. The supplier's profit monotonically increases in the buyer's second period information accuracy.

Proof of Proposition 2.4: We will first show that under the conditions stated in the proposition, the supplier's profit under the optimal early static contract is decreasing in θ_1 at $\theta_1 = p_L$ (i.e., $\frac{\partial \Pi_{ES}}{\partial \theta_1}|_{\theta_1 = p_L} < 0$). From equation (A.1) and the results in Proposition A.1, we can derive:

$$\begin{aligned} \frac{\partial \Pi_{ES}}{\partial \theta_1} &= \Gamma(D_H, q_L)\left[-\theta_1 \frac{dp_H^1}{d\theta_1} - 1 - p_H^1\right] + \Gamma(D_L, q_L)\left[-(1-\theta_1)\frac{dp_H^1}{d\theta_1} + 1 + p_H^1\right] \\ &\quad + \Gamma(D_H, q_H)\left[\theta_1 \frac{dp_H^1}{d\theta_1} + p_H^1\right] + \Gamma(D_L, q_H)\left[(1-\theta_1)\frac{dp_H^1}{d\theta_1} - p_H^1\right] \\ &\quad - \frac{dp_H^1}{d\theta_1} c_1 (q_H - q_L) \end{aligned} \tag{A.5}$$

Notice that when $\theta_1 = p_L$, we have $p_H^1 = 0$ and $\frac{dp_H^1}{d\theta_1} = \frac{1}{2p_L - 1}$. Plugging these into (A.5), we obtain

$$\begin{aligned} \frac{\partial \Pi_{ES}}{\partial \theta_1}|_{\theta_1 = p_L} &= \frac{1}{(2p_L - 1)}[(1-3p_L)\Gamma(D_H, q_L) + (3p_L - 2)\Gamma(D_L, q_L) + \\ &\quad p_L \Gamma(D_H, q_H) + (1-p_L)\Gamma(D_L, q_H) - c_1(q_H - q_L)] \end{aligned}$$

Then, it follows immediately that whenever $c_1(q_H - q_L) + (2p_L - 1)[\Gamma(D_H, q_L) - \Gamma(D_L, q_L)] > p_L[\Gamma(D_H, q_H) - \Gamma(D_H, q_L)] + p_H[\Gamma(D_L, q_H) - \Gamma(D_L, q_L)]$, $\frac{\partial \Pi_{ES}}{\partial \theta_1}|_{\theta_1 = p_L} < 0$. This implies that there exists $\bar{\theta} > p_L$ such that the supplier's profit under the optimal early static contract is decreasing in θ_1 for $\theta_1 \leq \bar{\theta}$. Hence, for a $\theta \leq \bar{\theta}$, we have that $\Pi_{ES}|_{\theta_1 = \theta} < \Pi_{ES}|_{\theta_1 = p_L}$. Now, consider the supplier's profit under an optimal dynamic contract when the buyer's first-period accuracy is $\theta_1 = p_L$ and the buyer's second-period accuracy is a $\theta_2 \in (p_L, \theta)$. Notice that Π_{DC} in this case cannot be smaller than $\Pi_{ES}|_{\theta_1 = p_L}$ since the early static contract is a special case of the dynamic contract. Hence, we have $\Pi_{DC}|_{\theta_1 = p_L, \theta_2 < \theta} \geq \Pi_{ES}|_{\theta_1 = p_L} > \Pi_{ES}|_{\theta_1 = \theta}$.

Proof of Proposition 2.5: The low-type buyer is the lowest type under early static and dynamic contract, and hence, makes zero expected profit under those two contract types in period 1. However, if offered a late static contract and if she observes a high demand signal in period 2, the low-type buyer can make a non-zero profit because (q_H, t_H) is not the lowest-type contract. More precisely, the expected profit of the buyer observing LH demand signal is given by $\sum_{j \in \{L,H\}} p_{Hj}^2 (\Gamma(D_j, q_H) - t_H)$, which is positive by the incentive compatibility constraint of the high-type. This occurs with positive probability, p_{LH}. Therefore, the expected profit of a low-type buyer under late static contract may be strictly more than that from early static and dynamic contracts.

Proof of Proposition 2.6: 1. Given the expression of optimal transfer payments for early static contract in Proposition A.1, we can derive the high-type buyer's expected profit from early static contract as $\pi_{ES} := (2\theta_1 - 1)[\Gamma(D_H, q_L) - \Gamma(D_L, q_L)]$. Likewise, from the result in Proposition 2.3, we can derive the high-type buyer's expected profit from dynamic contract as $\pi_{DC} := (2\theta_1 - 1)[\Gamma(D_H, q_{LL}) - \Gamma(D_L, q_{LL})]$. Hence, the profit difference is

$$\pi_{ES} - \pi_{DC} = (2\theta_1 - 1)[\Gamma(D_H, q_L) - \Gamma(D_H, q_{LL}) + \Gamma(D_L, q_L) - \Gamma(D_L, q_{LL})].$$

Notice that to show $\pi_{ES} - \pi_{DC} \geq 0$, it suffices to show that $q_L \geq q_{LL}$. Under an optimal dynamic contract characterized in Proposition 2.3, $\bar{q}_L$ is the solution to the sum of the equations for q_{LL} and q_{LH}. Note however that the equation that characterizes $\bar{q}_L$ is identical to the equation that characterizes the optimal q_L in early static contract. Hence, $q_L = \bar{q}$. Let q_{Lj} be the solution to the first-order condition of the dynamic contract: $y_{Lj}(q) = 0,\ j \in \{L, H\}$ (i.e., $y_{LL}(q) = 0$ is equal to the equation for q_{LL} in Proposition 2.3). Then, $\bar{q}_L$ and q_L is the solution to $y_{LL}(q) + y_{LH}(q) = 0$. Notice that $y_{LH}(q_{LL}) \geq 0$ by the unimodality of the supplier's profit in the contract quantity, and the fact that $q_{LL} \leq q_{LH}$. This implies $y_{LL}(q_{LL}) + y_{LH}(q_{LL}) = y_{LH}(q_{LL}) \geq 0$. Hence, it follows that $q_{LL} \leq \bar{q}_L = q_L$, which in turn implies that $\pi_{ES} - \pi_{DC} \geq 0$.

2. From part 1., we have $\pi_{DC} := (2\theta_1 - 1)[\Gamma(D_H, q_{LL}) - \Gamma(D_L, q_{LL})]$. By Property 3 in Assumption 1, π_{DC} increases with q_{LL} since $D_H \succcurlyeq D_L$. We will show the result by showing $\frac{\partial q_{LL}}{\partial \theta_2} \leq 0$. By the implicit function theorem, $\frac{\partial q_{LL}}{\partial \theta_2} = \frac{-\partial y_{LL}/\partial \theta_2}{\partial y_{LL}/\partial q_{LL}}$, where $y_{LL}(\theta_2, q_{LL}) = 0$ denotes the first-order condition with respect to q_{LL} given in

Proposition 2.3. We will show the result for the two cases of i) $\frac{c_1}{c_2} \leq p_{LH}$, and ii) $\frac{c_1}{c_2} > p_{LH}$.

i) $\frac{c_1}{c_2} \leq p_{LH}$: In this case, we have $\frac{\partial y_{LL}}{\partial \theta_2} = \frac{(2\theta_2(1-\theta_2)-\theta_1)}{(2\theta_2-1)^2}[\Gamma'(D_H, q_{LL}) - \Gamma'(D_L, q_{LL})] - \frac{(2\theta_1-1)}{(2\theta_2-1)^2}\Gamma'(D_L, q_{LL})$. First, we will show that $A := (2\theta_2(1-\theta_2)-\theta_1) \leq 0$. To see this, notice that $\frac{dA}{d\theta_2} = -2(2\theta_2-1) \leq 0$ since $\theta_2 \geq \theta_1 \geq \max\{p_L, p_H\} \geq \frac{1}{2}$. Hence, A is maximized at $\theta_2 = \theta_1$. Note that $A|_{\theta_2=\theta_1} = -\theta_1(2\theta_1-1) \leq 0$, implying $A \leq 0$ for any θ_1 and θ_2. Next, notice from $\frac{\partial y_{LL}}{\partial \theta_2}$ that $\Gamma'(D_H, q_{LL}) - \Gamma'(D_L, q_{LL}) \geq 0$ by Property 3 in Assumption 1, and $\Gamma'(D_L, q_{LL}) \geq 0$ by Property 2 in Assumption 1. Thus, we can conclude that $\frac{\partial y_{LL}}{\partial \theta_2} \leq 0$. Now, consider $\frac{\partial y_{LL}}{\partial q_{LL}}$. Let Π_{DC} denote the supplier's profit under the dynamic contract. Then, by definition $y_{LL} = \frac{\partial \Pi_{DC}}{\partial q_{LL}}$, and $\frac{\partial y_{LL}}{\partial q_{LL}} = \frac{\partial^2 \Pi_{DC}}{\partial q_{LL}^2}$. Since the supplier's profit is unimodal in q_{LL} by Property 4 in Assumption 1, it follows that $\frac{\partial^2 \Pi_{DC}}{q_{LL}^2} \leq 0$ at an optimal q_{LL}. Hence, $\frac{\partial y_{LL}}{\partial q_{LL}} \leq 0$. Together with $\frac{\partial y_{LL}}{\partial \theta_2} \leq 0$ that we have shown earlier, we have that $\frac{\partial q_{LL}(\theta_2)}{\partial \theta_2} \leq 0$.

ii) $\frac{c_1}{c_2} > p_{LH}$: In this case, we have $\frac{\partial y_{LL}}{\partial \theta_2} = \frac{(2\theta_2(1-\theta_2)-\theta_1)}{(2\theta_2-1)^2}[\Gamma'(D_H, q_{LL}) - \Gamma'(D_L, q_{LL})] - \frac{(2\theta_1-1)}{(2\theta_2-1)^2}[\Gamma'(D_L, q_{LL}) - c_2]$. Since $\frac{c_1}{c_2} > p_{LH}$, $\rho_L^* = q_{LL}$ by Proposition 2.3. Hence, if the buyer orders the LH contract in the second period, the supplier needs to produce $q_{LH} - q_{LL}$ at the cost of c_2, and will receive a transfer payment of t_{LH}. For the LH contract to be profitable to the supplier, the production cost must be no greater than the transfer payment. That is, $c_2(q_{LH} - q_{LL}) \leq t_{LH}$. From the expression of t_{LH} in Proposition 2.3, it is implied that $c_2 \leq \theta_2[\frac{\Gamma(D_H, q_{LH}) - \Gamma(D_H, q_{LL})}{q_{LH} - q_{LL}}] + (1-\theta_2)[\frac{\Gamma(D_L, q_{LH}) - \Gamma(D_L, q_{LL})}{q_{LH} - q_{LL}}] \leq \theta_2\Gamma'(D_H, q_{LL}) + (1-\theta_2)\Gamma'(D_L, q_{LL})$, where the second inequality follows from Property 4 in Assumption 1. Applying this to the expression of $\frac{\partial y_{LL}}{\partial \theta_2}$ derived above, we have $\frac{\partial y_{LL}}{\partial \theta_2} \leq [\frac{(2\theta_2(1-\theta_2)-\theta_1+\theta_2(2\theta_1-1))}{(2\theta_2-1)^2}][\Gamma'(D_H, q_{LL}) - \Gamma'(D_L, q_{LL})]$. Let $B := (2\theta_2(1-\theta_2) - \theta_1 + \theta_2(2\theta_1 - 1))$. Notice that $\frac{dB}{d\theta_1} = 2\theta_2 - 1 \geq 0$. Hence, B is maximized at $\theta_1 = \theta_2$, where $B|_{\theta_1=\theta_2} = 0$. This implies $B \leq 0$ for any θ_1 and

θ_2. Note also that $\Gamma'(D_H, q_{LL}) - \Gamma'(D_L, q_{LL}) \geq 0$ by Property 3 in Assumption 1. Hence, $\frac{\partial y_{LL}}{\partial \theta_2} \leq 0$. Next, consider $\frac{\partial y_{LL}}{\partial q_{LL}}$, which is identical to that in i) since $c_1 - p_{LH}c_2$ is independent of q_{LL}. We have shown in i) that $\frac{\partial y_{LL}}{\partial q_{LL}} \leq 0$. Thus, $\frac{\partial q_{LL}(\theta_2)}{\partial \theta_2} \leq 0$.

Proof of Theorem 2.2: We will first show that the supplier's optimal strategy when $K \leq \tilde{K}$ is to announce upfront that a side payment will be given to any buyer who updates forecast, and the optimal strategy when $K > \tilde{K}$ is to offer an early static contract. Let Π_{ES} and Π_{DC} be the supplier's expected profits from an optimal early static and dynamic contract, respectively. We denote the supplier's profit when the update is free by $\Pi_x^f,\ x \in \{ES, DC\}$, and when the update is costly by $\Pi_x^c,\ x \in \{ES, DC\}$. Since an early static contract is a special case of dynamic contract, we have $\Pi_{DC}^f \geq \Pi_{ES}^f$. Let $\tilde{K} = \Pi_{DC}^f - \Pi_{ES}^f \geq 0$. If the update cost is K, and the supplier pays the update cost, it is easy to see that the supplier's profit from dynamic contract is reduced by K. If the buyer pays the update cost, the supplier's profit from dynamic contract is reduced by K since the supplier has to reduce the transfers to $t_{ij} - K$, to ensure the buyer's participation. The supplier's profit from early static contract remains the same because there is no update cost involved. Thus, $\Pi_{DC}^c = \Pi_{DC}^f - K$, and $\Pi_{ES}^c = \Pi_{ES}^f$. Consequently, $\Pi_{DC}^c - \Pi_{ES}^c = \tilde{K} - K$. If $K \geq \tilde{K}$, early static contract is more profitable to the supplier than dynamic contract. Hence, the supplier always offers an early static contract without a side payment. If $K < \tilde{K}$, then the supplier offers a dynamic contract with a side payment of $S < \tilde{K} - K$ to any buyer who is willing to obtain a forecast update. Notice that under such contract, the low-type buyer strictly prefers to obtain a forecast update. The high-type buyer does not have an incentive to deviate from updating, which can be supported by an off-the-path equilibrium belief such that if the supplier observes no-update, then he believes the buyer is high-type with probability one, and will offer the first-best high-type early static contract, resulting in zero profit to the deviating high-type. Hence, the pooling equilibrium where both types update is a PBE.

Next, we will show that if $K \leq \tilde{K}$, separating equilibrium does not exist when the supplier announces upfront that a side payment will be given if the buyer updates forecast. Suppose there exists a separating equilibrium where the low-type announces an update while the high-type announces no update. Then, the supplier's optimal contract for the high-type is the first-best early static contract. For the low-type, the optimal contract is the first-best dynamic contract with a side payment. In this case, the high-type has an incentive to mimic the low-type since the low-type's contract gives the high-type positive profits. Therefore, such an equilibrium does not exist. Now consider the opposite separating equilibrium where the low-type buyer announces to not update; the high-type announces to update. In this equilibrium, the low-type buyer is offered the first-best early static contract; the high-type buyer is offered the first-best dynamic contract with a side payment. Let the high-type profit from taking the low-type's first-best early static contract be M. Notice that it is optimal for the supplier to offer a side payment smaller than M to maximize his expected profit. In this case, it is profitable for the high-type buyer to deviate to the no-update strategy. Hence, this separating equilibrium does not exist.

We can show there is no separating equilibrium for the case of $K > \tilde{K}$ in the same way.

Proof of Theorem 2.3: 1. Observe that the constraints for the two-dimensional screening contract in the supplier's problem include all the constraints of both early static and dynamic contract. Hence, the set of constraints that are implied and removed from the problem is the same as what we have shown in Appendix A.1. Based on the resulting reduced problem, we can show that the period 2 incentive compatibility constraints IC_{HH} and IC_{LH} are binding at optimality. Otherwise, the supplier can increase t_{HH} and t_{LH}, respectively, without violating any other constraints. The binding IC_{iH} implies that the type iH buyer's expected profit from choosing the contract type iH and iL in period 2 are the same, $i \in \{L, H\}$.

Now, consider the incapable type i buyer's expected profit from the dynamic contract, given by

$$\upsilon_{iH}\left[\sum_{j\in\{L,H\}}\sum_{k\in\{L,H\}} p_{ij}p_{jk}^2\Gamma(D_k, q_i^D + q_{iH}) - t_{iH}\right] +$$
$$\upsilon_{iL}\left[\sum_{j\in\{L,H\}}\sum_{k\in\{L,H\}} p_{ij}p_{jk}^2\Gamma(D_k, q_i^D + q_{iL}) - t_{iL}\right] - t_i^D.$$

We can see that the difference in the expected profit to the capable buyer and the incapable buyer essentially comes from the mismatch of the chosen contract quantity and the demand type predicted by the improved forecast: $\upsilon_{iH}\sum_{k\in\{L,H\}} p_{iL}p_{Lk}^2\Gamma(D_k, q_i^D + q_{iH})$ and $\upsilon_{iL}\sum_{k\in\{L,H\}} p_{iH}p_{Hk}^2\Gamma(D_k, q_i^D + q_{iL})$. That is, with probability $\upsilon_{iH}p_{iL}$, the incapable buyer chooses the iH contract but would have observed signal L in period 2 if she were capable. With probability $\upsilon_{iL}p_{iH}$, the buyer chooses the iL contract but would have observed signal H in period 2 if she were capable. Given what we showed earlier that the period 2 constraint IC_{iL} can be removed and IC_{iH} is binding, we know that the low-subtype (iL) deviation to the high-subtype (iH) contract is not profitable; while, the high-subtype deviation to the low-subtype contract yields the same expected profit as that from the high-subtype contract. This implies with probability $\upsilon_{iH}p_{iL}$, the incapable buyer may end up losing money. The optimal strategy for the incapable buyer is to always choose the low-subtype contract, iL, in period 2; $\upsilon_{iL} = 1$.

2. From $\upsilon_{iL} = 1$ in part 1., we have that $\sum_{l\in\{L,H\}}\sum_{j\in\{L,H\}} \upsilon_{il}p_{ij}\sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i^D + q_{il}) - t_{il}] - t_i^D = \sum_{j\in\{L,H\}} p_{ij}\sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i^D + q_{iL}) - t_{iL}] - t_i^D$. Since the period 2 constraints IC_{iH} are binding at optimality, it follows that $\sum_{j\in\{L,H\}} p_{ij}\sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i^D + q_{iL}) - t_{iL}] - t_i^D = \sum_{j\in\{L,H\}} p_{ij}\sum_{k\in\{L,H\}} p_{jk}^2[\Gamma(D_k, q_i^D + q_{ij}) - t_{ij}] - t_i^D$. That is, the RHS of the first forecasting capability incentive constraint is

equal to the LHS of the second forecasting capability constraint, $i \in \{L, H\}$. This shows that both forecasting capability incentive constraints are binding at optimality since the LHS of the first constraint is the same as the RHS of the second constraint. Therefore, the equality of the capable and incapable expected profit for each buyer type i is shown.

Now, consider the low-type buyer. Given the remaining constraints in the reduced problem, we can check that the period 1 participation constraints for both incapable and capable low-type buyer must be binding at optimality. Otherwise, the supplier can profitably and feasibly increase t_L and t_L^D accordingly. This implies that the low-type expected profit is zero for both capable and incapable buyer. For the high-type buyer, the period 1 participation constraints are not binding, implying the high-type expected profit is positive for both capable and incapable buyer.

3. We will first show that the supplier's profit monotonically increases with the capability probability ϕ. The supplier's profit for a given capability probability $\phi \in [0, 1]$ is $\Pi(\phi) = \phi\Pi_C(\phi) + (1 - \phi)\Pi_I(\phi)$ where $\Pi_C(\phi)$ is the expected profit from contracting with a capable buyer and $\Pi_I(\phi)$ is the expected profit from contracting with an incapable buyer. Notice that $\Pi_C(\phi) \geq \Pi_I(\phi)$ for any $\phi \in [0, 1]$. Otherwise, the supplier can increase his profit by offering the same early static contract to the capable buyer while satisfying all constraints because the early static contract is a special case of the dynamic contract. Next, we show that $\Pi(\phi') \geq \Pi(\phi)$ for any $\phi' > \phi$. Notice that if $\phi' > \phi$, then we have $\phi'\Pi_C(\phi) + (1 - \phi')\Pi_I(\phi) \geq \phi\Pi_C(\phi) + (1 - \phi)\Pi_I(\phi) = \Pi(\phi)$ because $\Pi_C(\phi) \geq \Pi_I(\phi)$. This implies that at probability ϕ', the supplier can at least offer the same optimal contract for ϕ to receive at least as much profit as $\Pi(\phi)$. Note that all constraints will be satisfied since the forecasting capability probability does not affect the buyer's profit as long as the same contract is offered. Hence, it follows that $\Pi(\phi') \geq \Pi(\phi)$.

Next, we will show that the supplier's profit monotonically increases with the buyer's period 2 forecast accuracy θ_2. It is easy to see that the results in Proposition 2.1 and Proposition 2.2 for the dynamic contract continue to hold in this model. Hence, we will consider the two-dimensional screening contract where the dynamic contract is in this simplified form with $q_i^D, t_{iL} = 0$.

Let $S^*(\theta_2)$ denote an optimal contract when the buyer's second-period information accuracy is θ_2. Let $\Pi^*(\theta_2) := \Pi(S^*(\theta_2), \theta_2)$ denote the corresponding supplier's optimal profit. We will show that when the buyer's second-period accuracy is $\theta_2'' > \theta_2'$, there exists a solution $S(\theta_2'')$ such that $\Pi^*(\theta_2'') \geq \Pi(S(\theta_2''), \theta_2'') \geq \Pi^*(\theta_2')$.

Suppose an optimal solution at θ_2' is given by $S^*(\theta_2') = \{q_i^*(\theta_2'), t_i^*(\theta_2'), q_{ij}^*(\theta_2'), t_i^{D*}(\theta_2'), t_{iH}^*(\theta_2'), \rho_i^*(\theta_2')\}, i, j \in \{L, H\}$. We construct a solution $S(\theta_2'') = \{q_i^*(\theta_2'), t_i^*(\theta_2'), q_{ij}^*(\theta_2'), t_i^{D*}(\theta_2'), t_{iH}(\theta_2''), \rho_i^*(\theta_2')\}, i, j \in \{L, H\}$ where

$$\begin{aligned} t_{LH}(\theta_2'') &= \theta_2''[\Gamma(D_H, q_{LH}^*(\theta_2')) - \Gamma(D_H, q_{LL}^*(\theta_2'))] + (1 - \theta_2'')[\Gamma(D_L, q_{LH}^*(\theta_2')) - \Gamma(D_L, q_{LL}^*(\theta_2'))] \\ t_{HH}(\theta_2'') &= \theta_2''[\Gamma(D_H, q_{HH}^*(\theta_2')) - \Gamma(D_H, q_{HL}^*(\theta_2'))] + (1 - \theta_2'')[\Gamma(D_L, q_{HH}^*(\theta_2')) - \Gamma(D_L, q_{HL}^*(\theta_2'))] \end{aligned}$$

Notice that we adopt the same early static contract and dynamic contract as the optimal contract for θ_2' except that we modify the period 2 transfer payments $t_{LH}(\theta_2'')$ and $t_{HH}(\theta_2'')$ for the dynamic contract. This construction technique is the same as what we have done in the proof of Theorem 2.1, allowing us to employ the results we have shown there.

We will first show that this solution is feasible for the supplier's problem with $\theta_2 = \theta_2''$. From the proof of Theorem 2.1, it follows immediately that this contract satisfies all the constraints in the original early static and dynamic contract. It remains to show that the contract satisfies the forecasting capability incentive constraints. By construction of $t_{iH}(\theta_2'\prime)$ and given that the period 2 IC_{iH} constraints are binding for $i \in \{L, H\}$, we have that the forecasting capability incentive constraints under $S(\theta_2'')$ are the same as those under $S^*(\theta_2')$. Hence, $S(\theta_2'')$ is a feasible contract.

Next, we will show $\Pi(S(\theta_2''), \theta_2'') \geq \Pi^*(\theta_2')$. Since the early static contract in $S(\theta_2'')$

is the same as that in $S^*(\theta_2')$, we have that the supplier's profit from the early static contract under $S(\theta_2'')$ is the same as that under $S^*(\theta_2')$. Hence, it suffices to show that the supplier's profit from the dynamic contract under $S(\theta_2'')$ is no less than that under $S^*(\theta_2')$. Let Π_{DC} denote the supplier's profit from the dynamic contract. We will show $\Pi_{DC}(S(\theta_2''), \theta_2'') \geq \Pi_{DC}^*(\theta_2')$.

As in the proof of Theorem 2.1, since the differences between $S(\theta_2'')$ and $S^*(\theta_2')$ only arise from $t_{iH}, i \in \{L, H\}$, we define $\hat{\Pi}_{DC}(\theta_2)$ as the supplier's profit from the dynamic contract where only t_{iH} are functions of θ_2, and the rest of the quantities and transfer payments are given. We will show $\frac{d\hat{\Pi}_{DC}(\theta_2)}{d\theta_2} \geq 0$ for any $\theta_2'' \in (\theta_2', 1]$. Since we construct $S(\theta_2'')$ in the same way as in the proof of Theorem 2.1, we have the same expression of $\frac{d\hat{\Pi}_{DC}(\theta_2)}{d\theta_2}$. Furthermore, notice that the characterization of the optimal period 1 production is the same as presented in Proposition 2.3 since the tradeoff between producing a unit in period 1 and in period 2 remains the same. Then, for the case of $\frac{c_1}{c_2} > p_{HH}$, the arguments in the proof of Theorem 2.1 also apply to this model since they are independent of the contract structure. It remains to show that the arguments in the case of $\frac{c_1}{c_2} \leq p_{HH}$ also apply. For this, it suffices to show that the equation that characterizes q_{HL} in the optimal two-dimensional screening contract is the same as that in the original dynamic contract.

First, we will argue that in an optimal two-dimensional screening contract, the period 1 incentive constraint for the incapable high-type (early static) must be binding. To see this, observe that at least one of the two period 1 incentive constraints for the high-type must be binding. Otherwise, the supplier can increase profit by increasing both t_H and t_H^D proportionally without violating other constraints. Let IC_H^{ES} denote the early static constraint, and IC_H^{DC} denote the dynamic constraint. Notice that the LHS of both IC_H^{ES} and IC_H^{DC} are equal due to the fact that the forecasting capability incentive constraints of the high-type are binding. The RHS of IC_H^{ES} is given by $(2\theta_1 - 1)[\Gamma(D_H, q_L) - \Gamma(D_L, q_L)]$ and the RHS of IC_H^{DC} is given by $(2\theta_1 - 1)[\Gamma(D_H, q_{LL}) - \Gamma(D_L, q_{LL})]$, following from the binding PC_L, IC_{LH}, and IC_{HH}. Hence, we can show that IC_H^{ES} is binding by showing $q_L \geq q_{LL}$. Suppose the contrary that $q_L < q_{LL}$ so that it is IC_H^{DC} that is binding. Then, t_H^D is derived from the binding IC_H^{DC}, and t_H is derived from the binding high-type forecasting capability incentive constraints. However, this results in the following first-order conditions:

$$\begin{aligned} y_L(q_L) &:= (1-(1+p_H^1)\theta_1)\Gamma'(D_H, q_L) + ((1+p_H^1)\theta_1 - p_H^1)\Gamma'(D_L, q_L) - c_1 p_L^1 \\ &\quad + p_H^1(2\theta_1 - 1)[\Gamma'(D_H, q_L) - \Gamma'(D_L, q_L)] = 0 \\ y_{\bar{L}}(\bar{q}_L) &:= \phi[(1-(1+p_H^1)\theta_1)\Gamma'(D_H, \bar{q}_L) + ((1+p_H^1)\theta_1 - p_H^1)\Gamma'(D_L, \bar{q}_L) - c_1 p_L^1] \\ &\quad -(1-\phi)p_H^1(2\theta_1 - 1)[\Gamma'(D_H, \bar{q}_L) - \Gamma'(D_L, \bar{q}_L)] = 0, \end{aligned}$$

where $\bar{q}_L$ is the optimal low-type quantity for the dynamic contract when bunching is optimal for the low-type. Notice that $y_{\bar{L}}(q_L) = -p_H^1(2\theta_1 - 1)[\Gamma'(D_H, q_L) - \Gamma'(D_L, q_L)] \leq 0$. Given the unimodality of the supplier's profit in $\bar{q}_L$, this implies $q_L \geq \bar{q}_L$. By the same argument as in the proof of Proposition 2.6 part 1, we can show that $\bar{q}_L \geq q_{LL}$. Hence, $q_L \geq q_{LL}$, which is a contradiction. This shows that IC_H^{ES} is binding at optimality. Then, it follows that t_H is derived from the binding IC_H^{ES}, and t_H^D is derived from the binding high-type forecasting capability incentive constraints. More precisely, t_H is as given in Proposition A.1, which is independent of q_{HL}, and $t_H^D = \theta_1\Gamma(D_H, q_{HL}) + (1-\theta_1)\Gamma(D_L, q_{HL}) - (2\theta_1 - 1)[\Gamma(D_H, q_L) - \Gamma(D_L, q_L)]$. Notice that $\frac{dt_H^D}{dq_{HL}}$ is the same as in the original dynamic contract.

Finally, consider the rest of the transfer payments in the two-dimensional screening contract. It is easy to see that t_L, t_L^D, and t_{LH} are independent of q_{HL}. The expression of t_{HH} is given by the binding IC_{HH}, which is the same as in the original dynamic contract. Hence, the first-order condition which characterizes q_{HL} is the same as given in Proposition 2.3. Then, the result that $\frac{d\hat{\Pi}_{DC}(\theta_2)}{d\theta_2} \geq 0$ for the case of $\frac{c_1}{c_2} \leq p_{HH}$ also applies to this model.

4. We will first show that the high-type buyer's profit is monotonically decreasing in the capability probability ϕ. From part 2., we know that both incapable and capable high-type buyer receive the same expected profit. From the proof of part 3., we have shown that under an optimal contract, the period 1 incentive constraint for the incapable high-type, IC_H^{ES}, is binding. Hence, the high-type buyer's profit is given by the RHS of IC_H^{ES}, which is $(2\theta_1 - 1)[\Gamma(D_H, q_L) - \Gamma(D_L, q_L)]$. We will show the result by showing $\frac{\partial q_L}{\partial \phi} \leq 0$. By the implicit function theorem, $\frac{\partial q_L}{\partial \phi} = \frac{-\partial y_L/\partial \phi}{\partial y_L/\partial q_L}$, where $y_L(\phi, q_L) = 0$ denotes the first-order condition with respect to q_L. We have shown in part 3. that the transfer payments t_H, t_L, and t_H^D are functions of q_L. Plugging in the expression of the transfer payments in the supplier's profit, we can derive $y_L(\phi, q_L)$ as

$$\begin{aligned} y_L(\phi, q_L) \quad &:= \quad (1-\phi)[(1-(1+p_H^1)\theta_1)\Gamma'(D_H, q_L) + ((1+p_H^1)\theta_1 - p_H^1) \\ &\qquad \Gamma'(D_L, q_L) - c_1 p_L^1] - \phi p_H^1(2\theta_1 - 1)[\Gamma'(D_H, q_L) - \Gamma'(D_L, q_L)] = 0 \\ &=: \quad (1-\phi)A(q_L) - \phi p_H^1(2\theta_1 - 1)B(q_L). \end{aligned} \tag{A.6}$$

Notice that $B(q_L) \geq 0$ by Property 3 in Assumption 1. Since $y_L(\phi, q_L) = 0$ at optimality, this implies $A(q_L) \geq 0$. Hence, it follows that $\frac{\partial y_L}{\partial \phi} = -A(q_L) - p_H^1(2\theta_1 - 1)B(q_L) \leq 0$. Now, consider $\frac{\partial y_L}{\partial q_L}$. Let Π denote the supplier's profit under an optimal two-dimensional contract. Then, by definition, $y_L = \frac{\partial \Pi}{\partial q_L}$, and $\frac{\partial y_L}{\partial q_L} = \frac{\partial^2 \Pi}{\partial q_L^2}$. Since the supplier's profit is unimodal in q_L by Property 4 in Assumption 1, it follows that $\frac{\partial^2 \Pi}{\partial q_L^2} \leq 0$ at an optimal q_L. Hence, $\frac{\partial y_L}{\partial q_L} \leq 0$. Together with what we have shown earlier that $\frac{\partial y_L}{\partial \phi} \leq 0$, we have that $\frac{\partial q_L}{\partial \phi} \leq 0$.

Next, we will show that the high-type buyer's profit is independent of the second forecast accuracy, θ_2. Given the expression for the high-type profit above, we can show the result by showing $\frac{\partial q_L}{\partial \theta_2} = 0$. It is easy to see from equation (A.6), which characterizes the optimal q_L, that q_L is independent of θ_2. Hence, the result follows immediately.

APPENDIX B

Additional Results and Proofs for Chapter 3

B.1 Proofs for Chapter 3

Proof of Proposition 3.1: We will first show the results for an all-unit discount $D^A = (r, K)$. Given the utility function in (3.3), we solve the consumer's maximization problem. Notice that $U_{ij}(q, D^A)$ is linear in q within each of the four intervals, implying that the optimal purchase quantity is a boundary solution, which is either $0, K$, or θ_i. After applying some algebra, we derive the closed-form expression of $\sigma_j(\theta_i, D^A)$ and $\bar{\theta}_j(D^A)$ as follow.

$$\sigma_j(\theta_i, D^A) := \begin{cases} p & \text{if } \theta_i < \bar{\theta}_j(D^A) \\ \frac{p(1-r)K - T_j(D^A)}{\theta_i} & \text{if } \bar{\theta}_j(D^A) \leq \theta_i < K \\ p(1-r) - \frac{T_j(D^A)}{K} & \text{if } \theta_i \geq K \end{cases} \tag{B.1}$$

$$\text{and } \bar{\theta}_j(D^A) := K(1-r) - \frac{T_j(D^A)}{p}$$

$$\text{where } T_j(D^A) := \begin{cases} 0 & \text{if } j = v \text{ or } r < R \\ t & \text{if } j = d \text{ and } r \geq R \end{cases}$$

It is immediate to see that $\sigma_j(\theta_i, D^A)$ and $\bar{\theta}_j(D^A)$ increase in p and K, and decrease in r. Additionally, it follows directly from (B.1) that $\bar{\theta}_d(D^k) \leq \bar{\theta}_v(D^k) \leq K$ and $\sigma_d(\theta_i, D^k) \leq \sigma_v(\theta_i, D^k) \leq p$.

To show part i) to iii) of the proposition, we will consider the three cases of i) $s < \sigma_j(\theta_i, D^A)$, ii) $s \geq \sigma_j(\theta_i, D^A)$ and $\theta_i < \bar{\theta}_j(D^A)$, and iii) $s \geq \sigma_j(\theta_i, D^A)$ and $\theta_i \geq \bar{\theta}_j(D^A)$, respectively. To determine the optimal purchase quantity in each case, we compare the utility at $q = 0, \theta_i$, and K. Note that the utility at $q = 0$ is always 0. The utility at $q = \theta$ and K are summarized in Table B.1.

θ_i	$\sigma_j(\theta_i, D^A)$	Utility when $q = \theta_i$ $U_{ij}(\theta_i, D^A)$	Utility when $q = K$ $U_{ij}(K, D^A)$
$\theta_i < \bar{\theta}_j(D^A)$	p	$(s_i - p)\theta_i$	$s_i\theta_i - p(1-r)K + T_j(D^A)$ $= s_i\theta_i - p\bar{\theta}_j(D^A)$
$\bar{\theta}_j(D^A) \leq \theta_i < K$	$\frac{p(1-r)K - T_j(D^A)}{\theta_i}$	$(s_i - p)\theta_i$	$s_i\theta_i - p(1-r)K + T_j(D^A)$ $= (s_i - \sigma_j(\theta_i, D^A))\theta_i$
$\theta_i \geq K$	$p(1-r) - \frac{T_j(D^A)}{K}$	$(s_i - p(1-r))\theta_i + T_j(D^A) =$ $(s_i - \sigma_j(\theta_i, D^A))\theta_i - (\frac{\theta_i - K}{K})T_j(D^A)$	$s_iK - p(1-r)K + T_j(D^A)$ $= (s_i - \sigma_j(\theta_i, D^A))K$

Table B.1: Consumer's utility from purchasing $q = \theta_i$ and K

i) $s_i < \sigma_j(\theta_i, D^A)$

It is straightforward to see from Table B.1 that the utility from purchasing θ_i and K are negative in all three intervals of θ_i because $s_i < \sigma_j(\theta_i, D^A) \leq p$. Thus, no purchase is optimal.

ii) $s_i \geq \sigma_j(\theta_i, D^A)$ and $\theta_i < \bar{\theta}_j(D^A)$

From Table B.1, $U_{ij}(\theta_i, D^A) \geq 0$ since $s \geq \sigma_j(\theta_i, D^A) = p$. Notice also that $U_{ij}(\theta_i, D^A) - U_{ij}(K, D^A) = p(\bar{\theta}_j(D^A) - \theta_i) > 0$ since $\theta_i < \bar{\theta}_j(D^A)$. Thus, it is optimal to buy θ_i at the full price.

iii) $s_i \geq \sigma_j(\theta_i, D^A)$ and $\theta_i \geq \bar{\theta}_j(D^A)$

From Table B.1, if $\bar{\theta}_j(D^A) \leq \theta_i < K$, we have $U_{ij}(K, D^A) \geq 0$ and $U_{ij}(K, D^A) - U_{ij}(\theta_i, D^A) = p(\theta_i - \bar{\theta}_j(D^A)) \geq 0$ from $s_i \geq \sigma_j(\theta_i, D^A)$ and $\theta_i \geq \bar{\theta}_j(D^A)$. Thus, it is optimal to buy K to receive the discount. If $\theta_i \geq K$, we have $U_{ij}(\theta_i, D^A) - U_{ij}(K, D^A) = (s_i - p(1-r))(\theta_i - K)$. This implies it is better to buy θ_i if $s_i \geq p(1-r)$, and it is better to buy K if $s_i < p(1-r)$. Note also that $U_{ij}(\theta_i, D^A) \geq 0$ if $s_i \geq p(1-r)$, and $U_{ij}(K, D^A) \geq 0$ because $s_i \geq \sigma_j(\theta_i, D^A)$. Thus, it is optimal to buy either K or $\theta_i > K$ and receive the discount if $\theta_i \geq K$.

The results for a fixed-amount discount $D^F = (m, K)$ can be shown in an analogous manner, where $\sigma_j(\theta_i, D^F)$ and $\bar{\theta}_j(D^F)$ are constructed by replacing r with $\frac{m}{pK}$ in the expressions of $\sigma_j(\theta_i, D^F)$ and $\bar{\theta}_j(D^F)$, respectively.

Proof of Proposition 3.2: The results follow directly from Proposition 3.1.

Proof of Proposition 3.3: i) This result is implicitly shown in the proof of Proposition 3.1.

ii) Part i) and Proposition 3.1 immediately imply that the optimal purchase quantity under the fixed-amount discount is identical to that under the all-unit discount when i) $s < \sigma_j(\theta, D^F)$, and ii) $s \geq \sigma_j(\theta, D^F)$ and $\theta < \bar{\theta}_j(D^F)$. Now, consider the remaining case where $s \geq \sigma_j(\theta, D^F)$ and $\theta \geq \bar{\theta}_j(D^F)$. If $\bar{\theta}_j(D^F) \leq \theta < K$, we have $U_{ij}(K, D^F) = U_{ij}(K, D^A)$ and $U_{ij}(\theta, D^F) = U_{ij}(\theta, D^A)$. Hence, the optimal purchase quantity is the same under the two discount schemes. If $\theta \geq K$, we have $U_{ij}(\theta, D^F) = s\theta - p\theta + m + T_j(D^F)$ and $U_{ij}(K, D^F) = sK - pK + m + T_j(D^F) = (s - \sigma_j(\theta, D^F))K$ under the fixed-amount discount, so $U_{ij}(\theta, D^F) - U_{ij}(K, D^F) = (s-p)(\theta - K)$. Notice that since $s \geq \sigma_j(\theta, D^F)$, $U_{ij}(K, D^F) \geq 0$. Note also that $U_{ij}(\theta, D^F) \geq 0$ if and only if $s \geq p$. Thus, the optimal purchase quantity under the fixed-amount discount is θ if $s \geq p$ or K if $s < p$. But on the other hand, for any $s \geq \sigma_j(\theta, D^A)$ and $\theta \geq K$, the optimal purchase quantity under the all-unit discount is θ if $s \geq p(1-r)$ or K if $s < p(1-r)$. The result follows from comparing the optimal purchase quantity under the two discount schemes.

Proof of Proposition 3.4: i) Suppose $\beta > 0$. We consider three cases: 1) $s_h < p$, 2) $s_l < p \leq s_h$, and 3) $s_l \geq p$. In each case, we prove the result by providing an all-unit discount $D^A = (r, K), r > 0, K > 0$, which yields a strictly higher profit than the profit from selling at no discount. Let $\Pi(0)$ be the seller's profit from offering no discount. We summarize $\Pi(0)$ and $\Pi(D^A)$ for each case in Table B.2.

Case	$D^A = (r, K)$	$\bar{\theta}_j(D^A)$	$\sigma_j(\theta_h, D^A)$	$\Pi(D^A)$	$\Pi(0)$
$s_h < p$	$(1 - \frac{s_h}{p}, \theta_h)$	$\frac{\theta_h s_h - T_j(D^A)}{p} \leq \theta_h$	$s_h - \frac{T_j(D^A)}{\theta_h} \leq s_h$	$\gamma s_h \theta_h +$ $(1-\gamma)\beta\Pi_{ld}(D^A)$	0
$s_l < p \leq s_h$	$(R, \frac{t+p\theta_h}{p(1-R)})$	$\frac{t+p\theta_h - T_j(D^A)}{p} \geq \theta_h$	p	$\gamma p\theta_h + \beta\gamma t$	$\gamma p \theta_h$
$s_l \geq p$	$(R, \frac{t+p\theta_h}{p(1-R)})$	$\frac{t+p\theta_h - T_j(D^A)}{p} \geq \theta_h$	p	$\gamma p\theta_h + \gamma\beta t$ $+(1-\gamma)p\theta_l$	$\gamma p\theta_h +$ $(1-\gamma)p\theta_l$

Table B.2: Seller's profit from offering no discount and D^A when $\beta > 0$

The seller's profits $\Pi(D^A)$ are derived using the result from Proposition 3.1 and equation (3.3). It is straightforward to see that $\Pi(D^A) > \Pi(0)$ for all three cases.

ii) Suppose $\beta = 0$. We will first show that no discount is optimal when condition a) or b) holds. Notice that $s_h \geq p$ in both condition a) and b). From Proposition 3.5, we know that when $s_h \geq p$, fixed-amount discount is the most profitable type of discounts. Hence, to show that no discount is optimal, it suffices to show that the optimal fixed-amount discount cannot yield a strictly greater profit than no discount.

a) Suppose $s_l \geq p$. From Figure 3.2, a type-i consumer buys either θ_i or $K > \theta_i$. Notice that it is not optimal to offer a discount and induce the consumer to buy θ_i since the consumer is already willing to buy θ_i at the full price. Suppose it is optimal to offer a fixed-amount discount $D^F = (m, K)$ to induce a type-i consumer to buy K. Then, from Proposition 3.1 and equation (B.1), the discount needs to satisfy $\theta_i \geq \bar{\theta}_v(D^F) = \frac{pK-m}{p}$. This implies that the seller's profit from selling to the type-i consumer is $pK - m \leq p\theta_i$. Notice however that the profit from selling to the type-i consumer under no discount is $p\theta_i$. Hence, no discount is optimal.

b) Suppose $s_h \geq p > s_l$ and $\gamma \geq \frac{s_l}{p}$. Under no discount, the seller's profit is $\gamma p \theta_h$. If a fixed-amount discount $D^F = (m, K)$ is offered, the following outcomes of (q_l, q_h) can be induced: $(0, 0), (0, \theta_h), (0, K), (K, K)$, and (K, θ_h). It is immediate to see that $(0,0)$ and $(0,\theta_h)$ cannot be more profitable than no discount. Now, consider $(0,K)$. To induce the high-type to buy K, the discount needs to satisfy $\theta_h \geq \bar{\theta}_v(D^F) = \frac{pK-m}{p}$ from Proposition 3.1 and equation (B.1(. This implies that the seller's profit from inducing $(0,K)$ is $\gamma(pK-m) \leq \gamma p\theta_h$, which cannot be strictly greater than the no-discount profit. Next, consider (K,K). To induce the low-type consumer to buy K, the discount needs to satisfy $s_l \geq \sigma_v(\theta_l, D^F)$, which from equation (B.1) implies that the seller's profit from inducing (K,K) is $pK-m \leq s_l\theta_l$. Notice that $s_l\theta_l \leq \gamma p\theta_h$ since $\gamma \geq \frac{s_l}{p}$ and $\theta_h \geq \theta_l$. Hence, the outcome (K,K) cannot be optimal. Finally, consider (K,θ_h), where the seller's profit is given by $\gamma(p\theta_h - m) + (1-\gamma)(pK-m)$. Solving the seller's problem, we obtain that the seller's profit from inducing (K,θ_h) is maximized at $K=0$ if $\gamma \geq \frac{s_l}{p}$, and is maximized at $K=\theta_l$ otherwise. Since $\gamma \geq \frac{s_l}{p}$ in this case, we have that the seller's profit is maximized at $K=0$ and $m=0$, which is equivalent to offering no discount.

To complete the proof, we will show that no discount is not optimal in the other cases outside of condition a) and b). That is, we will show there exists a discount which yields a strictly greater profit than offering no discount when c) $s_h < p$, or d) $s_l < p \leq s_h$ and $\gamma < \frac{s_l}{p}$.

c) Suppose $s_h < p$. In this case, the no-discount profit is zero. The seller can do better by offering a price markdown with a discount $r = \frac{p-s_l}{p}$. Under this price markdown, the high-type consumer buys θ_h and the low-type consumer buys θ_l. The seller receives a profit of $\gamma s_l\theta_h + (1-\gamma)s_l\theta_l > 0$. Hence, no discount is not optimal.

d) Suppose $s_l < p \leq s_h$ and $\gamma < \frac{s_l}{p}$. The seller can offer a fixed-amount discount with $K=\theta_l$ and $m=(p-s_l)\theta_l$. Under this discount, the high-type consumer buys θ_h and the low-type consumer buys $K=\theta_l$. The seller's profit is $\gamma p(\theta_h -$

$\theta_l) + s_l\theta_l$, which is greater than the no-discount profit of $\gamma p\theta_h$ since $\gamma < \frac{s_l}{p}$.

Proof of Proposition 3.5: i) We first show that the optimal all-unit discount dominates the optimal fixed-amount discount when $s_h < p$. For this, we prove that for an optimal fixed-amount discount $D^{F*} = (m^*, K)$, there exists an all-unit discount that results in at least as much profit for the seller. Consider an all-unit discount $D^A = (r = \frac{m^*}{pK}, K^A = K)$. We compare the profit that the seller obtains from a consumer type $ij, i \in \{l, h\}, j \in \{v, d\}$, under D^A and D^{F*}. From Proposition 3.3, if $K \geq \theta_i$ or $s_i < p(1-r)$, the consumer will buy the same quantity under both the all-unit and fixed-amount discount. Notice that under this situation, the consumer will never buy more than K since $s_h < p$. Hence, the seller's profits under the two discount schemes are the same. Now, consider the case where $K < \theta_i$ and $s_i \geq p(1-r)$. From Proposition 3.3, the consumer will buy exactly K under the fixed-amount discount, but will buy $\theta_i > K$ under the all-unit discount. Hence, the seller earns $pK - m^*$ under the fixed-amount discount, but earns $p(1-r)\theta_i = p(1 - \frac{m^*}{pK})\theta_i > pK - m^*$ under the all-unit discount. Thus, the all-unit discount weakly outperforms the fixed-amount discount.

Next, we will show the existence of $\bar{\beta} \in [0, 1]$ by contradiction. To represent a price markdown, we let $\epsilon > 0$ denote the smallest sellable unit of the product, where ϵ is arbitrarily small. Then, a price markdown is given by $D^M = (r, K = \epsilon)$. Suppose that there exist $\beta_1 < \beta_2$ such that when the proportion of deal-prone consumers is β_1, price markdown is not an optimal all-unit discount; but at β_2, price markdown is an optimal all-unit discount. We will show that there exists an all-unit discount with $K > \epsilon$ which is strictly more profitable than the optimal price markdown at β_2. Let $D^A = (r_1, K_1 > \epsilon)$ be an optimal all-unit discount at β_1, and $D^M = (r_2, \epsilon)$ be the optimal price markdown at β_2. Also, let $\Pi(\beta, D)$ be the seller's profit from offering an all-unit discount D at β. Then, $\Pi(\beta_1, D^A) > \Pi(\beta_1, D^M)$ and $\Pi(\beta_2, D^M) \geq \Pi(\beta_2, D^A)$. Since $\beta_2 > \beta_1$, $\Pi(\beta_2, D^A) \geq \Pi(\beta_1, D^A)$ from Proposition 3.2. Hence, it follows that $\Pi(\beta_2, D^M) > \Pi(\beta_1, D^M)$, which implies the seller's profit from the deal-prone consumer is strictly greater than that from the value-conscious consumer with D^M (i.e., D^M results in an overspending). From Proposition 3.2, overspending under D^M occurs when $q_{di} = K = \epsilon > q_{vi} = 0$ for some $i \in \{l, h\}$, resulting in the seller's expected profit from the type-i consumer of $\beta_2(p(1 - r_2))\epsilon$. Now, consider another all-unit discount $D^{A\prime} = (r_2, K = \min\{\frac{t}{p(1-r^*(\beta_2))-s_i}, \theta_i\})$. Under this discount, $\sigma_d(\theta_i, D^{A\prime}(\beta_2)) = p(1 - r^*(\beta_2)) - \frac{T_d(D^{A*}(\beta_2))}{K} \leq s_i$. Hence, from Proposition 3.1, the type-i deal-prone consumer buys K; other types buy the same quantity as under D^M. The seller's expected profit from the type-i consumer is $\beta_2(p(1 - r_2))K$, which is strictly greater than that under D^M since $K > \epsilon$. Thus, $D^{A\prime}$ is more profitable than D^M at β_2, which contradicts the optimality of price markdown at β_2.

ii) There are two cases to consider here: iia) $s_l \geq p$, and iib) $s_l < p$. We will show in each case that for any optimal all-unit discount $D^{A*} = (r^*, K)$, there exists a fixed-amount discount $D^F = (m, K^F)$ which yields at least as much profit to the seller. For this, we will compare the seller's profit from a consumer type $ij, i \in \{l, h\}, j \in \{v, d\}$, under D^F and D^{A*}.

iia) Suppose $s_l \geq p$. Consider a fixed-amount discount $D^F = (m = pr^*K, K^F = K)$. From Proposition 3.3, since $s_i \geq p$, the consumer always purchases the same quantity q under the two discount policies, which is either $\theta_i < K, K$, or $\theta_i > K$. If $q \leq K$, the seller's profit under both discount policies are the same. However, if $q = \theta_i > K$, the seller earns $p(1 - r^*)\theta_i$ under the all-unit discount, but $p(\theta_i - r^*K) > p(1 - r^*)\theta_i$ under the fixed-amount discount. Thus, the seller's profit from any consumer type ij under the fixed-amount discount weakly dominates that under the optimal all-unit discount.

iib) Suppose $s_l < p$. Consider a fixed-amount discount $D^F = (m = pr^*K, K^F = K)$. From Proposition 3.3, if $s_l < p(1 - r^*)$ or $\theta_l \leq K$, the consumer purchases the same quantity under D^F and D^{A*}. By the same argument as in iia), we can show that the seller's profit under D^F is greater than or equal to that under D^{A*}. Now, for the remaining case where $s_l \geq p(1-r^*)$ and $\theta_l > K$, consider another fixed-amount discount $D^{F'} = (m' = pr^*K, K^{F'} = \theta_l)$. From Proposition 3.3, under D^{A*}, the low-type consumer buys θ_l and the high-type consumer buys θ_h; under $D^{F'}$, the low-type consumer buys $K^{F'} = \theta_l$ and the high-type consumer buys θ_h as well. The seller's profits from selling to the low-type consumer under both discount policies are the same at $p(1-r^*)\theta_l$. However, the seller's profit from the high-type consumer under the optimal all-unit discount is $p(1-r^*)\theta_h$, which is less than that under the fixed amount discount of $p\theta_h - pr^*K$. Hence, the fixed-amount discount weakly dominates the all-unit discount.

The results from iia) and iib) together complete the proof of ii).

Proof of Proposition 3.6: First, we consider the seller's profit difference between no discount and an optimal conditional discount. Under no discount, deal-prone consumers never receive transaction utility. Hence, the purchase quantity of deal-prone and value-conscious consumers are always the same, implying the seller's profit under no discount is independent of β and t. On the other hand, the seller's profit under an optimal conditional discount weakly increases in β because the profit from selling to deal-prone consumers is always greater than or equal to the profit from selling to value-conscious consumers under any conditional discount. The profit under an optimal conditional discount also weakly increases in t because an increase in t weakly increases the transaction utility realized by deal-prone consumers under any given conditional discount, resulting in a weak increase in the purchase quantity of deal-prone consumers, and subsequently, in the seller's profit at optimality.

Next, we consider the seller's profit differences between price markdown and conditional discount. To represent a price markdown, we let $\epsilon > 0$ denote the smallest sellable unit of the product, where ϵ is arbitrarily small. Then, a price markdown is given by $D^M = (r, K = \epsilon)$. We will show the result for the three cases: i) $s_l \geq p$, ii)

$s_h \geq p > s_l$, and iii) $s_h < p$. For each case, Table B.3 summarizes the possible consumer purchase quantities $Q = (q_{hd}, q_{hv}, q_{ld}, q_{lv})$ under a price markdown, the optimal markdown $r^{M*}(Q)$ that induces the purchase quantity Q, and the seller's profit from the optimal price markdown $\Pi^{M*}(Q)$, derived using Figure 3.3 and (3.3).

Case	Possible Purchase Quantities	Optimal Markdown $r^{M*}(Q)$	Seller's Profit $\Pi^{M*}(Q)$
$s_l \geq p$	Q_1	0	$\gamma p\theta_h + (1-\gamma)p\theta_l$
$s_h \geq p > s_l$	Q_1	$1 - \frac{s_l}{p}$	$\gamma s_l\theta_h + (1-\gamma)s_l\theta_l$
	Q_2 if $R < 1 - \frac{s_l}{p}$	R	$\gamma p(1-R)\theta_h + (1-\gamma)\beta p(1-R)\epsilon$
	Q_3	0	$\gamma p\theta_h$
$s_h < p$	Q_1	$1 - \frac{s_l}{p}$	$\gamma s_l\theta_h + (1-\gamma)s_l\theta_l$
	Q_2 if $\max\{R, 1 - \frac{s_h}{p}\} < 1 - \frac{s_l}{p}$	$\max\{R, 1 - \frac{s_h}{p}\}$	$\begin{cases} \gamma s_h\theta_h + (1-\gamma)\beta s_h\epsilon & \text{if } R \leq 1 - \frac{s_h}{p} \\ \gamma p(1-R)\theta_h + (1-\gamma)\beta p(1-R)\epsilon & \text{if } R > 1 - \frac{s_h}{p} \end{cases}$
	Q_3 if $R > 1 - \frac{s_h}{p}$	$1 - \frac{s_h}{p}$	$\gamma s_h\theta_h$
	Q_4 if $R \leq 1 - \frac{s_h}{p}$	R	$\beta p(1-R)\epsilon$
	Q_5 if $R \leq 1 - \frac{s_h}{p}$	R	$\beta\gamma p(1-R)\epsilon$
	Q_6	0	0

where $Q_1 = (\theta_h, \theta_h, \theta_l, \theta_l), Q_2 = (\theta_h, \theta_h, \epsilon, 0), Q_3 = (\theta_h, \theta_h, 0, 0), Q_4 = (\epsilon, 0, \epsilon, 0), Q_5 = (\epsilon, 0, 0, 0), Q_6 = (0, 0, 0, 0)$.

Table B.3: Possible outcomes under a price markdown

1. $s_l \geq p$

   From Table B.3, no discount is optimal under price markdown. Hence, the result follows.

2. $s_h \geq p > s_l$

   From Table B.3, notice that $\Pi^{M*}(Q_2) < \Pi^{M*}(Q_3)$ for any $\gamma > 0$, and $\Pi^{M*}(Q_2) < \Pi^{M*}(Q_1)$ if $\gamma = 0$ since ϵ is arbitrarily small. Thus, Q_2 cannot be an optimal outcome under a price markdown. Now, consider Q_1 and Q_3, and note that $\Pi^{M*}(Q_1)$ and $\Pi^{M*}(Q_3)$ are independent of β and t. Hence, the result follows.

3. $s_h < p$

   First, notice from Table B.3 that $\Pi^{M*}(Q_6) \leq \Pi^{M*}(Q_5) \leq \Pi^{M*}(Q_4) < \Pi^{M*}(Q_1)$ since ϵ is arbitrarily small. Hence, Q_4, Q_5, and Q_6 cannot be an optimal outcome under a price markdown. If Q_1 or Q_3 is optimal, the result follows since the seller's profit is independent of β and t. Now, consider Q_2. Notice that

$\Pi^{M*}(Q_2)$ is independent of t. Hence, the result regarding an increase in t follows immediately. It remains to show the result regarding an increase in β when Q_2 is optimal. If $R > 1 - \frac{s_h}{p}$, then $\Pi^{M*}(Q_2) < \Pi^{M*}(Q_3)$ for any $\gamma > 0$, and $\Pi^{M*}(Q_2) < \Pi^{M*}(Q_1)$ if $\gamma = 0$ since ϵ is arbitrarily small. Hence, Q_2 cannot be optimal. If $R \leq 1 - \frac{s_h}{p}$, Q_2 is optimal if $\Pi^{M*}(Q_2) > \Pi^{M*}(Q_1)$, which requires $\gamma > \frac{s_l\theta_l}{s_h\theta_h - s_l\theta_h + s_l\theta_l}$. However, in this situation, the optimal conditional discount, derived from solving the seller's problem, is an all-unit discount which induces at least one type of deal-prone consumers to overspend by $K > \epsilon$. Thus, $\Pi^{A*} = \beta Pr(i)p(1 - r^{A*})K + (1 - Pr(i))p(1 - r^{A*})q_{-i}$, where $i \in \{l, h\}$ is the type of deal-prone consumers who overspends. Then, $\frac{d[\Pi^{A*} - \Pi^{M*}(Q_2)]}{d\beta} = Pr(i)p(1 - r^{A*})K - (1 - \gamma)s_h\epsilon > 0$ since ϵ is arbitrarily small. Hence, the result follows.

Proof of Lemma 3.1: We will first show the result for the all-unit discount. Notice from Proposition 3.4 part i) that since $\beta > 0$, no discount cannot be optimal. Thus, an optimal discount must increase purchase quantity of at least one type of consumers. Let $D^{A*}(\gamma) = (r^*(\gamma), K^*(\gamma))$ be the terms of the optimal all-unit discount when the proportion of high-type consumers is γ. Let $\Pi_H(D^{k*}(\gamma))$ and $\Pi_L(D^{k*}(\gamma))$ be the seller's profits earned from the high-type and the low-type consumers under the optimal discount, respectively.

We will show the result by contradiction. Suppose at $\gamma = \gamma_1$, the optimal discount $D^{A*}(\gamma_1)$ induces only the high-type consumers to increase their purchase quantity. Suppose also that there exists $\gamma_2 > \gamma_1$ such that the optimal discount $D^{A*}(\gamma_2)$ induces the low-type consumers to increase their purchase quantity. From the optimality of $D^{A*}(\gamma_1)$ at $\gamma = \gamma_1$, and $D^{A*}(\gamma_2)$ at $\gamma = \gamma_2$, we have

$$\gamma_1\Pi_H(D^{A*}(\gamma_1)) + (1-\gamma_1)\Pi_L(D^{A*}(\gamma_1)) > \gamma_1\Pi_H(D^{A*}(\gamma_2)) + (1-\gamma_1)\Pi_L(D^{A*}(\gamma_2)), \text{ and} \quad \text{(B.2)}$$

$$\gamma_2\Pi_H(D^{A*}(\gamma_2)) + (1-\gamma_2)\Pi_L(D^{A*}(\gamma_2)) \geq \gamma_2\Pi_H(D^{A*}(\gamma_1)) + (1-\gamma_2)\Pi_L(D^{A*}(\gamma_1)). \qquad \text{(B.3)}$$

We will now show that $\Pi_H(D^{A*}(\gamma_1)) \geq \Pi_H(D^{A*}(\gamma_2))$. To see this, suppose $\Pi_H(D^{A*}(\gamma_2)) > \Pi_H(D^{A*}(\gamma_1))$. Since $s_h > s_l$ and $\theta_h > \theta_l$, the high-type purchase quantity is always no less than the low-type purchase quantity. Given the fact that the low-type consumer increases purchase quantity under $D^{A*}(\gamma_2)$, we must have $q_h(D^{A*}(\gamma_2)) \geq q_l(D^{A*}(\gamma_2)) > 0$. Then, using the results from Proposition 3.1 and equation (B.1), we can construct an all-unit discount $D^{A\prime} = (r', K^*(\gamma_2))$, with $r' < r^*(\gamma_2)$, such that $q_l(D^{A\prime}) = q_l(r = 0)$ and $q_h(D^{A\prime}) = q_h(D^{A*}(\gamma_2))$. Since $r' < r^*(\gamma_2)$ and $q_h(D^{A\prime}) = q_h(D^{A*}(\gamma_2))$, it follows that $\Pi_H(D^{A\prime}) \geq \Pi_H(D^{A*}(\gamma_2)) > \Pi_H(D^{A*}(\gamma_1))$. Notice also that under the discount $D^{A\prime}$, the low-type does not increase purchase quantity. Hence, the low-type purchases the same quantity at no discount under both $D^{A\prime}$ and $D^{A*}(\gamma_1)$, implying $\Pi_L(D^{A\prime}) = \Pi_L(D^{A*}(\gamma_1))$. Hence, when $\gamma = \gamma_1$, we must have $\Pi(D^{A\prime}, \gamma_1) > \Pi(D^{A*}(\gamma_1), \gamma_1)$. But this contradicts the optimality of $D^{A*}(\gamma_1)$. Thus, it must be that $\Pi_H(D^{A*}(\gamma_1)) \geq \Pi_H(D^{A*}(\gamma_2))$. Applying this inequality to (B.3), we have $\Pi_L(D^{A*}(\gamma_2)) \geq \Pi_L(D^{A*}(\gamma_1))$. Now, since $\Pi_H(D^{A*}(\gamma_1)) \geq \Pi_H(D^{A*}(\gamma_2)), \Pi_L(D^{A*}(\gamma_2)) \geq \Pi_L(D^{A*}(\gamma_1))$, and $\gamma_1 < \gamma_2$, it follows that

$$\gamma_2[\Pi_H(D^{A*}(\gamma_1)) - \Pi_H(D^{A*}(\gamma_2))] - (1-\gamma_2)[\Pi_L(D^{A*}(\gamma_2)) - \Pi_L(D^{A*}(\gamma_1))] \geq$$
$$\gamma_1[\Pi_H(D^{A*}(\gamma_1)) - \Pi_H(D^{A*}(\gamma_2))] - (1-\gamma_1)[\Pi_L(D^{A*}(\gamma_2)) - \Pi_L(D^{A*}(\gamma_1))].$$

Then, from (B.2), we must have $\gamma_2\Pi_H(D^{A*}(\gamma_1)) + (1-\gamma_2)\Pi_L(D^{A*}(\gamma_1)) >$ $\gamma_2\Pi_H(D^{A*}(\gamma_2)) + (1-\gamma_2)\Pi_L(D^{A*}(\gamma_2))$, which contradicts the optimality of $D^{A*}(\gamma_2)$ at γ_2.

The result for the fixed-amount discount can be shown in the same way.

Proof of Proposition 3.7: We will first show the results for the all-unit discount. In preparation, we employ the results from Proposition 3.1 and enumerate possible purchase quantities for three cases: no discount, discount increases the high-type purchase quantity, and discount increases the low-type purchase quantity. Table B.4 summarizes these results.

<table>
<tr><th rowspan="2">Case</th><th rowspan="2">No Discount
(q_l, q_h)</th><th colspan="2">Only High-type Increases Purchase Quantity</th><th colspan="2">Low-type Increases Purchase Quantity</th></tr>
<tr><th>(q_l, q_h)</th><th>$D^{A*} = (r, K)$ satisfies</th><th>(q_l, q_h)</th><th>$D^{A*} = (r, K)$ satisfies</th></tr>
<tr><td rowspan="3">$s_h < p$</td><td rowspan="3">$(0, 0)$</td><td>$(0, \theta_h)$</td><td>$K \leq \theta_h$;
$s_h \geq p(1-r)$</td><td>(θ_l, θ_h)</td><td>$K \leq \theta_l; s_l \geq p(1-r)$</td></tr>
<tr><td rowspan="2">$(0, K)$</td><td rowspan="2">$\theta_h \geq \bar{\theta}_d(D^A)$;
$s_h \geq \sigma_d(\theta_h, D^A)$</td><td>$(K, \theta_h)$</td><td>$K \leq \theta_h; \theta_l \geq \bar{\theta}_d(D^A)$;
$s_l \geq \sigma_d(\theta_l, D^A)$</td></tr>
<tr><td>(K, K)</td><td>$\theta_l \geq \bar{\theta}_d(D^A)$;
$s_l \geq \sigma_d(\theta_l, D^A)$</td></tr>
<tr><td rowspan="3">$s_l < p \leq s_h$</td><td rowspan="3">$(0, \theta_h)$</td><td rowspan="3">$(0, K)$</td><td rowspan="3">$K > \theta_h \geq \bar{\theta}_d(D^A)$</td><td>$(\theta_l, \theta_h)$</td><td>$K \leq \theta_l; s_l \geq p(1-r)$</td></tr>
<tr><td>(K, θ_h)</td><td>$K \leq \theta_h; \theta_l \geq \bar{\theta}_d(D^A)$;
$s_l \geq \sigma_d(\theta_l, D^A)$</td></tr>
<tr><td>(K, K)</td><td>$K \geq \theta_h; \theta_l \geq \bar{\theta}_d(D^A)$;
$s_l \geq \sigma_d(\theta_l, D^A)$</td></tr>
<tr><td rowspan="2">$s_l \geq p$</td><td rowspan="2">(θ_l, θ_h)</td><td rowspan="2">(θ_l, K)</td><td rowspan="2">$K > \theta_h \geq \bar{\theta}_d(D^A)$</td><td>$(K, \theta_h)$</td><td>$\bar{\theta}_d(D^A) \leq \theta_l < K \leq \theta_h$</td></tr>
<tr><td>(K, K)</td><td>$K > \theta_l \geq \bar{\theta}_d(D^A)$</td></tr>
</table>

Table B.4: Possible outcomes under an optimal all-unit discount

i) Suppose $\gamma > \Gamma^A(t)$. From Lemma 3.1, the optimal all-unit discount increases the purchase quantity of only the high-type consumer.

ia) We will show that $K^{A*} \geq \theta_h$. Notice from Table B.4 that if $s_h \geq p$, the optimal discount must satisfy $K > \theta_h \geq \bar{\theta}_d(D^A)$; thus, the result holds. Now consider the case of $s_h < p$. From Table B.4, the possible outcomes in this case are $(0, \theta_h)$ and $(0, K)$. Suppose that $K < \theta_h$ under the optimal discount $D^A = (r, K)$. We will show that there exists another all-unit discount which yields a greater profit than D^A.

a.1) $(0, K)$: From the proof of Proposition 3.1 part iii), when $K < \theta_h$, the high-type consumer buys K if $\sigma_d(\theta_h, D^A) \leq s_h < p(1-r)$. Applying the definition of $\sigma_d(\theta_h, D^A)$ given in equation (B.1), we have $p(1-r) - \frac{T_d(D^A)}{K} \leq$

$s_h < p(1-r)$. This implies $T_d(D^A) = t > 0$. Thus, $p(1-r)K \leq s_h K + t$, implying $\Pi(D^A) = \gamma p(1-r)K \leq \gamma(s_h K + t)$.

Now, consider another all-unit discount $D^{A\prime} = (r' = \max\{\frac{pK' - s_h\theta_h - t}{pK'}, R\}, K' = \max\{\theta_h, \frac{t+s_h\theta_h}{p(1-R)}\})$. This discount has $K' \geq \theta_h$. Applying the definitions from (B.1), we have $\bar{\theta}_d(D^{A\prime}) = \frac{s_h\theta_h}{p} < \theta_h \leq K'$ and $s_h = \sigma_d(\theta_h, D^{A\prime})$ for the high-type consumer. Hence, from Proposition 3.1, the high-type purchases K'. For the low-type consumer, we have $s_l < \sigma_d(\theta_l, D^{A\prime})$, implying the low-type does not purchase. Applying the expressions of r' and K', we have $\Pi(D^{A\prime}) = \gamma p(1-r')K' = \gamma(s_h\theta_h + t) > \Pi(D^A)$ since $\theta_h > K$.

a.2) $(0, \theta_h)$: From the proof of Proposition 3.1 part iii), when $K < \theta_h$, the high-type consumer buys θ_h if $s_h \geq p(1-r)$. This implies $\Pi(D^A) = \gamma p(1-r)\theta_h \leq \gamma s_h\theta_h$. Notice that this profit is less than the profit from offering $D^{A\prime}$ given in a.1). since $t > 0$.

Hence, when $\gamma > \Gamma^A(t)$, an optimal all-unit discount has $K^{A*} \geq \theta_h$.

ib) We will show that $r^* \geq R$. From ia), $K^{A*} \geq \theta_h$. This implies the optimal discount must induce either $(q_l, q_h) = (0, K)$ when $s_l < p$, or (θ_l, K) when $s_l \geq p$. Suppose the optimal all-unit discount is $D^A = (r, K)$, with $K \geq \theta_h$ and $r < R$. Hence, the deal-prone consumers do not receive transaction utility; i.e., $T_d(D^A) = 0$. From equation (B.1), this results in $\sigma_d(\theta, D^A) = \sigma_v(\theta, D^A)$ and $\bar{\theta}_d(D^A) = \bar{\theta}_v(D^A)$. We will show that there exists another all-unit discount which yields a greater profit than D^A.

b.1) $(0, K)$: First, consider the case of $s_h \geq p$. From Proposition 3.1 and equation (B.1), since the high-type buys K, we must have $\theta_h \geq \bar{\theta}_d(D^A) = K(1-r)$. This implies $\Pi(D^A) = \gamma p(1-r)K \leq \gamma p\theta_h$. Notice that $D^{A\prime}$ with $r' \geq R$ given in a.1) results in $\Pi(D^{A\prime}) = \gamma(s_h\theta_h + t) > \Pi(D^A)$ since $s_h \geq p$ and $t > 0$. Now, consider the case of $s_h < p$. From Proposition 3.1, the high-type buys K if $\theta_h \geq \bar{\theta}_d(D^A)$ and $s_h \geq \sigma_d(\theta_h, D^A)$. Applying the definition of $\sigma_d(\theta_h, D^A)$ from equation (B.1), we have $s_h \geq \frac{p(1-r)K}{\theta_h}$. This implies $\Pi(D^A) = \gamma p(1-r)K \leq \gamma s_h \theta_h$, which is also dominated by the seller's profit from $D^{A\prime}$ given in a.1) since $t > 0$.

b.2) (θ_l, K): From Proposition 3.1 and equation (B.1), since the high-type buys K, we must have $\theta_h \geq \bar{\theta}_d(D^A) = K(1-r)$. This implies $\Pi(D^A) = \gamma p(1-r)K + (1-\gamma)p\theta_l$ is no greater than $\gamma p\theta_h + (1-\gamma)p\theta_l$, which is the seller's no-discount profit. However, we know from Proposition 3.4 part i) that no discount cannot be optimal since $\beta > 0$. Hence, D^A cannot be optimal for this outcome.

We have shown that when $\gamma > \Gamma^A(t)$, an optimal all-unit discount has $r^* \geq R$. This completes the proof of part i).

ii) Suppose $\gamma \leq \Gamma^A(t)$. From Lemma 3.1, the optimal all-unit discount increases the purchase quantity of the low-type consumer. We will show the result for $\hat{t} = (p\theta_h - s_l\theta_l)^+$.

iia) We will show that if $t > \hat{t}$, then $K^{A*} \geq \theta_l$. Notice from Table B.4 that if $s_l \geq p$, the optimal discount must satisfy $K > \theta_l \geq \bar{\theta}_d(D^A)$; thus, the result holds. Furthermore, if $s_l < p \leq s_h$ and both types buy K, the result also holds since $K \geq \theta_h > \theta_l$. It remains to show the result for the outcomes of $(q_l, q_h) = (\theta_l, \theta_h)$ and (K, θ_h) when $s_l < p$, and $(q_l, q_h) = (K, K)$ when $s_h < p$. Suppose that $K < \theta_l$ under the optimal discount $D^A = (r, K)$. We will show that there exists another all-unit discount resulting in a greater profit than D^A.

a.1) (K, K): From the proof of Proposition 3.1 part iii), when $K < \theta_l$, the low-type consumer buys K if $\sigma_d(\theta_l, D^A) \leq s_l < p(1-r)$. Applying the definition of $\sigma_d(\theta_l, D^A)$ given in equation (B.1), we have $p(1-r) - \frac{T_d(D^A)}{K} \leq s_l < p(1-r)$. This implies $T_d(D^A) = t > 0$. Thus, $p(1-r)K \leq s_hK + t$, implying $\Pi(D^A) = p(1-r)K \leq s_l K + t$. Now, consider another all-unit discount $D^{A\prime} = (r' = \max\{\frac{pK' - s_l\theta_l - t}{pK'}, R\}, K' = \frac{t + s_l\theta_l}{p(1-R)})$. Notice that since $t > \hat{t} = (p\theta_h - s_l\theta_l)^+$, this discount has $K' > \theta_h > \theta_l$. Applying the definition from equation (B.1), we have $\bar{\theta}_d(D^{A\prime}) = \frac{s_l\theta_l}{p} < \theta_l < K'$ and $\sigma_d(\theta_l, D^{A\prime}) = s_l$ for the low-type consumer. Hence, from Proposition 3.1, the low-type purchases K'. For the high-type consumer, we have $\bar{\theta}_d(D^{A\prime}) < \theta_h < K'$ and $\sigma_d(\theta_h, D^{A\prime}) = \frac{s_l\theta_l}{\theta_h} < s_h$. Hence, the high-type also purchases K'. Applying the expressions of r' and K', we have $\Pi(D^{A\prime}) = p(1-r')K' = s_l\theta_l + t$. Notice that $\Pi(D^{A\prime}) > \Pi(D^A)$ since $K < \theta_l$. Thus, D^A cannot be optimal.

a.2) (K, θ_h): Similar to a.1), for the low-type consumer to buy $K < \theta_l$, D^A needs to trigger transaction utility, i.e., $T_d(D^A) = t > 0$. This implies $r \geq R$. Hence, $\Pi(D^A) = p(1-r)(\gamma\theta_h + (1-\gamma)K) \leq p(1-R)(\gamma\theta_h + (1-\gamma)K) < p\theta_h$, where the last inequality comes from $K < \theta_l < \theta_h$, and $R > 0$. Notice that $\Pi(D^A)$ is less than the profit from $D^{A\prime}$ given in a.1) since $\Pi(D^{A\prime}) = s_l\theta_l + t > p\theta_h$ from $t > \hat{t} = (p\theta_h - s_l\theta_l)^+$.

a.3) (θ_l, θ_h): From the proof of Proposition 3.1 part iii), when $K < \theta_l$, the low-type buys θ_l if $s_l \geq p(1-r)$. This implies $\Pi(D^A) = p(1-r)(\gamma\theta_h + (1-\gamma)\theta_l) \leq s_l(\gamma\theta_h + (1-\gamma)\theta_l)$. Notice that this is less than $p\theta_h$ since $s_l < p$ and $\theta_l < \theta_h$. Hence, it is also dominated by $\Pi(D^{A\prime})$ as shown in a.2).

We have shown that when $\gamma \leq \Gamma^A(t)$ and $t > \hat{t}$, an optimal all-unit discount has $K^{A*} \geq \theta_l$.

iib) We will show that if $t > \hat{t}$, then $r^* \geq R$. From iia), $K^{A*} \geq \theta_l$. Hence, from Table B.4, the optimal discount must induce either $(q_l, q_h) = (K, K)$ or (K, θ_h). Suppose that the optimal all-unit discount is $D^A = (r, K)$ with $K \geq \theta_l$ and $r < R$. Hence, the deal-prone consumers do not receive transaction utility, i.e.,

$T_d(D^A) = 0$. From equation (B.1), this results in $\sigma_d(\theta, D^A) = \sigma_v(\theta, D^A)$ and $\bar{\theta}_d(D^A) = \bar{\theta}_v(D^A)$. We will show that there exists another all-unit discount which yields a greater profit than D^A.

b.1) (K, K): From Proposition 3.1, the low-type buys K if $\bar{\theta}_d(D^A) \leq \theta_l$ and $s_l \geq \sigma_d(\theta_l, D^A)$. Applying the definition of $\sigma_d(\theta_l, D^A)$ from equation (B.1), we have $s_l \geq \frac{p(1-r)K}{\theta_l}$. This implies $\Pi(D^A) = p(1-r)K \leq s_l\theta_l$. Notice that this profit is dominated by the profit from offering $D^{A\prime}$ with $r \geq R$, given in a.1), since $\Pi(D^{A\prime}) = s_l\theta_l + t$ and $t > 0$.

b.2) (K, θ_h): From Proposition 3.1, the low-type buys K if $\bar{\theta}_d(D^A) \leq \theta_l$. Applying the definition of $\bar{\theta}_d(D^A)$ from equation (B.1), we have $K(1-r) \leq \theta_l$. This implies $\Pi(D^A) = \gamma p(1-r)\theta_h + (1-\gamma)p(1-r)K \leq \gamma p(1-r)\theta_h + (1-\gamma)p\theta_l < p\theta_h$ since $r \geq 0$ and $\theta_h > \theta_l$. We have seen from a.1) that $\Pi(D^{A\prime}) > p\theta_h$. Hence, D^A is dominated by $D^{A\prime}$.

We have shown that when $\gamma \leq \Gamma^A(t)$ and $t > \hat{t}$, an optimal all-unit discount has $r^* \geq R$. This completes the proof of part ii).

Now, consider the fixed-amount discount. Given the results in Proposition 3.3, it follows that all possible purchase quantities under the fixed-amount discount $D^{F*} = (m^* = pr^*K^{F*}, K^{F*} = K^{A*})$ have already been included in Table B.4. Hence, the results for the fixed-amount discount can be shown in the same way.

Proof of Lemma 3.2: i) This part of the lemma is analogous to Lemma 3.1. Notice from Proposition 3.4 that since $\beta > 0$, no discount is never optimal. Furthermore, from Proposition 3.2, it is not possible to increase the value-conscious purchase quantity alone. Hence, an optimal discount must either increase the deal-prone purchase quantity alone, or both the deal-prone and the value-conscious purchase quantity. The result can be shown in the same way as the proof of Lemma 3.1 by replacing γ with β, H with d, and L with v, and noting that $q_d(D^{k*}(\beta)) \geq q_v(D^{k*}(\beta))$ from

Proposition 3.2.

ii) This part of the lemma is shown in the proof of Proposition 3.8.

Proof of Proposition 3.8: We will characterize the optimal all-unit and fixed-amount discount in the two cases: i) $s \geq p$ and ii) $s < p$, and show that the optimal all-unit and fixed-amount discount result in the same seller's profit. Notice from Lemma 3.2 part i) that an optimal discount must either increase the deal-prone purchase quantity alone, or increase both the deal-prone and value-conscious purchase quantity.

i) Suppose $s \geq p$. Then, at no discount, both types of consumers buy θ.

ia) **All-unit discount:** From Proposition 3.1, more specifically as laid out in Figure 3.2 and Figure 3.3, the only possible outcome where only the deal-prone consumer increases purchase quantity is $(q_v, q_d) = (\theta, K)$. The only possible outcome where both deal-prone and value-conscious consumer increase purchase quantity is $(q_v, q_d) = (K, K)$. We will compare the seller's profit under the two outcomes and show that for any $\beta > 0$, it is optimal to induce (θ, K) by offering $K > \theta$ and $r = R$. That is, $\bar{\beta}^A(t) = 0$ in this case.

First, consider the all-unit discount which results in $(q_v, q_d) = (\theta, K)$, and the seller's profit of $\beta p(1-r)K + (1-\beta)p\theta$. We will show that the best all-unit discount in this case is $D^A = (R, \frac{t+p\theta}{p(1-R)})$. Applying r and K in D^A to equation (B.1), we have $\theta = \bar{\theta}_d(D^A) < \bar{\theta}_v(D^A)$. Hence, by Proposition 3.1, this discount results in $(q_v, q_d) = (\theta, K)$ and gives a profit of $\Pi(D^A) = \beta t + p\theta$. Suppose there exists another all-unit discount $D^{A'} = (r', K')$ which results in a strictly greater profit than D^A does. To induce (θ, K'), $D^{A'}$ must satisfy $\theta \geq \bar{\theta}_d(D^{A'}) \geq K'(1-r') - \frac{t}{p}$ from equation (B.1). However, this implies $\Pi(D^{A'}) \leq \beta t + p\theta = \Pi(D^A)$. Hence, D^A is the best all-unit discount that results in $(q_v, q_d) = (\theta, K)$ and the seller's profit of $\Pi^*(\theta, K) = \beta t + p\theta$.

Now, consider the all-unit discount which results in $(q_v, q_d) = (K, K)$, and the seller's profit of $p(1-r)K$. Suppose $D^A = (r, K)$ is the best discount. To induce this outcome, D^A needs to satisfy $\theta \geq \bar{\theta}_v(D^A) = K(1-r)$ from equation (B.1). This implies $\Pi(D^A) \leq p\theta$. Notice that for any $\beta > 0$, this profit is less than $\Pi^*(\theta, K) = \beta t + p\theta$, the maximum profit from inducing (θ, K). Hence, the outcome (K, K) is never optimal. It is always optimal to induce the outcome (θ, K) with $D^{A*} = (R, \frac{t+p\theta}{p(1-R)})$, where $K^* > \theta$ and $r^* = R$.

ib) **Fixed-amount discount:** From Proposition 3.1, the only possible outcome where only the deal-prone consumer increases purchase quantity is $(q_v, q_d) = (\theta, K)$. The only possible outcome where both deal-prone and value-conscious consumer increase purchase quantity is $(q_v, q_d) = (K, K)$. We will compare the seller's profit under the two outcomes and show that for any $\beta > 0$, it is optimal to induce (θ, K) by offering $K > \theta$ and $m = M = pRK$. That is, $\bar{\beta}^F(t) = 0$ in this case.

First, consider the fixed-amount discount which results in $(q_v, q_d) = (\theta, K)$, and the seller's profit of $\beta(pK - m) + (1-\beta)p\theta$. We will show that the best fixed-amount discount in this case is $D^F = (m = pRK, \frac{t+p\theta}{p(1-R)})$. Notice that D^F has the same K and the same discount depth of $\frac{m}{pK} = R$ as D^{A*} in ia). Hence, from Proposition 3.3 part i), the switching curves under D^F are identical to those under D^{A*}. It then follows from Proposition 3.1 that under D^F, consumers buy $(q_v, q_d) = (\theta, K)$ and the seller receives a profit of $\Pi(D^F) = \beta t + p\theta$. Now, suppose there exists another fixed-amount discount $D^{F'} = (m', K')$ which results in a strictly greater profit than D^F does. To induce (θ, K'), $D^{F'}$ must satisfy $\theta \geq \bar{\theta}_d(D^{F'}) \geq K'(1 - \frac{m'}{pK'}) - \frac{t}{p}$ from equation (B.1). However, this implies $\Pi(D^{F'}) \leq \beta t + p\theta = \Pi(D^F)$. Hence, D^F is the best fixed-amount discount that results in $(q_v, q_d) = (\theta, K)$ and the seller's profit of $\Pi^*(\theta, K) = \beta t + p\theta$.

Now, consider the fixed-amount discount which results in $(q_v, q_d) = (K, K)$, and the seller's profit of $pK - m$. Suppose $D^F = (m, K)$ is the best discount. To induce this outcome, D^F needs to satisfy $\theta \geq \bar{\theta}_v(D^F) = K(1-\frac{m}{pK})$ from equation (B.1). This implies $\Pi(D^F) \leq p\theta$. Notice that for any $\beta > 0$, this profit is less than $\Pi^*(\theta, K) = \beta t + p\theta$, the maximum profit from inducing (θ, K). Hence, the outcome (K, K) is never optimal. It is always optimal to induce the outcome (θ, K) with $D^{F*} = (m = pRK, \frac{t+p\theta}{p(1-R)})$.

From ia) and ib), we have that the optimal all-unit and fixed-amount discount have the same K and the same discount depth, and result in the same seller's profit.

ii) Suppose $s < p$. Then, at no discount, both types of consumers do not buy.

iia) **All-unit discount:** From Proposition 3.1, the only possible outcome where only the deal-prone consumer increases purchase quantity is $(q_v, q_d) = (0, K)$. There are two possible outcomes where both deal-prone and value-conscious consumer increase purchase quantity: $(q_v, q_d) = (K, K)$ and $(q_v, q_d) = (\theta, \theta)$. Following the same procedure as in part i), we can characterize the best all-unit discount which increases only the deal-prone purchase quantity, and the best discount which increases both deal-prone and value-conscious purchase quantity. The best all-unit discount in each case is provided below.

If only the deal-prone purchase quantity is increased, it is most profitable to offer

$$D^A = (r, K) = \begin{cases} (\frac{\theta(p-s)-t}{p\theta}, \theta) & \text{if } \frac{\theta(p-s)-t}{p\theta} \geq R \\ (R, \frac{\theta s+t}{p(1-R)}) & \text{if } \frac{\theta(p-s)-t}{p\theta} < R \end{cases}$$

Notice that $r \geq R$ and $K \geq \theta$. The resulting outcome is $(q_v, q_d) = (0, K)$, and the seller's profit is $\Pi^*(0, K) = \beta(t + s\theta)$.

If both the deal-prone and value-conscious purchase quantity are increased, it is most profitable to offer $D^A = (1 - \frac{s}{p}, \theta)$. The resulting outcome is $(q_v, q_d) =$

(θ, θ), and the seller's profit is $\Pi^*(\theta, \theta) = s\theta$.

Now, notice that $\Pi^*(0, K) > \Pi^*(\theta, \theta)$ if and only if $\beta > \frac{s\theta}{t+s\theta}$. Hence, it is optimal to increase only the deal-prone purchase quantity if $\beta > \frac{s\theta}{t+s\theta}$, and it is optimal to increase both deal-prone and value-conscious purchase quantity if $\beta \leq \frac{s\theta}{t+s\theta}$. This shows that $\bar{\beta}^A(t) = \frac{s\theta}{t+s\theta}$ from Lemma 3.2 part i).

iib) **Fixed-amount discount:** We know from Proposition 3.5 that when $s < p$, the all-unit discount weakly dominates the fixed-amount discount. Hence, to show that the optimal fixed-amount discount results in the same profit as the optimal all-unit discount does, it suffices to show that there exists a fixed-amount discount which yields the same seller's profit as $D^{A*} = (r^*, K^*)$ does. For this, consider $D^F = (m = pr^*K^*, K^*)$. It is easy to check that D^F results in the same outcomes and the same seller's profit as D^{A*}, characterized above. Then, it also follows that $\bar{\beta}^F(t) = \bar{\beta}^A(t) = \frac{s\theta}{t+s\theta}$.

iia) and iib) complete the proof of part ii).

Proof for Optimal Endogenous Prices

We will first show the result for the all-unit discount by showing that the seller's profit from setting $p = s_l, s_h$, or $\frac{s_h\theta_h+t}{\theta_h(1-R)}$ weakly dominates the seller's profit obtained from all other prices.

i) $p < s_l$ is dominated by $p = s_l$

Suppose $D^A = (p, r, K)$ is optimal for some $p < s_l$. Now, consider $D^{A'} = (p' = s_l, r' = \frac{s_l - p(1-r)}{s_l}, K' = K)$. Since $p'(1-r') = p(1-r), K' = K$, and $r' > r$, from equation (B.1) and Proposition 3.1, we have that a consumer purchases at least as much under $D^{A'}$ as under D^A. Notice that if the consumer purchases at discount, the seller's margin is the same at $p(1-r)$ under both D^A and $D^{A'}$. However, if the consumer purchases at no discount, the seller's margin is greater $(p' = s_l > p)$ under $D^{A'}$. Hence, $D^{A'}$ weakly dominates D^A.

ii) $s_l < p < s_h$ is dominated by $p = s_h$

Suppose $D^A = (p, r, K)$ is optimal for some $s_l < p < s_h$. Now, consider $D^{A'} = (p' = s_h, r' = \frac{s_h - p(1-r)}{s_h}, K' = K)$. Applying the same logic as i), we can show that $D^{A'}$ weakly dominates D^A.

iii) $s_h < p < \frac{s_h\theta_h + t}{\theta_h(1-R)}$ is dominated by $p = \frac{s_h\theta_h + t}{\theta_h(1-R)}$

Suppose $D^A = (p, r, K)$ is optimal for some $s_h < p < \frac{s_h\theta_h + t}{\theta_h(1-R)}$. Now, consider $D^{A'} = (p' = \frac{s_h\theta_h + t}{\theta_h(1-R)}, r' = 1 - \frac{p(1-r)(1-R)\theta_h}{s_h\theta_h + t}, K' = K)$. Applying the same logic as i), we can show that $D^{A'}$ weakly dominates D^A.

iv) $p > \frac{s_h\theta_h + t}{\theta_h(1-R)}$ can be replicated by $p = \frac{s_h\theta_h + t}{\theta_h(1-R)}$. Suppose $D^A = (p, r, K)$ is optimal for some $p > \frac{s_h\theta_h + t}{\theta_h(1-R)}$. Notice that since $p > \frac{s_h\theta_h + t}{\theta_h(1-R)} > s_h$, the seller can make a positive profit only when he offers $r \geq R$ such that $\sigma_d^A(r, \theta_h) \geq s_h$ (so that at least high-type deal-prone consumers buy). Now, consider $D^{A'} = (p' = \frac{s_h\theta_h + t}{\theta_h(1-R)}, r' = 1 - \frac{p(1-r)(1-R)\theta_h}{s_h\theta_h + t}, K' = K)$. Note that $r' \geq R$ since $\sigma_d^A(r, \theta_h) \geq s_h$. Then, from equation (B.1) and Proposition 3.1, consumers always purchase the same quantity at discount under both D^A and $D^{A'}$. Since $p(1-r) = p'(1-r')$, the seller always makes the same profit under D^A and $D^{A'}$.

The result for the fixed-amount discount can be shown in the same way.

B.2 Optimal Discount Strategies for Deal-Prone Market

All-unit discount:

i) Suppose $s_h < p$.

ia) If $\gamma > \Gamma^A(t)$, $K^* \geq \max\{\frac{t+\theta_h s_h}{p(1-R)}, \theta_h\}, r^*(K) = \max\{\frac{pK - \theta_h s_h - t}{pK}, R\}, (q_l, q_h) = (0, K^*)$, and $\Pi^{A*} = \gamma(\theta_h s_h + t)$.

ib) If $\gamma \leq \Gamma^A(t)$, there exist $t_i^A(R), i \in \{1, 2, 3\}$ where

- For $t \leq t_1^A(R)$, $\Gamma^A(t) = \Gamma_1^A(t, R)$, $K^* \leq \theta_l$ and $(q_l, q_h) = (\theta_l, \theta_h)$.
- For $t_1^A(R) < t \leq t_2^A(R)$, $\Gamma^A(t) = \Gamma_2^A(t, R)$, $(q_l, q_h) = (\theta_l, \theta_h)$ or (K^*, θ_h).
- For $t_2^A(R) < t \leq t_3^A(R)$, $\Gamma^A(t) = \Gamma_3^A(t, R)$, $K^* \geq \theta_l$ and $(q_l, q_h) = (\theta_l, \theta_h)$ or (K^*, θ_h) or (K^*, K^*).
- For $t > t_3^A(R)$, $\Gamma^A(t) = \Gamma_4^A(t, R)$, $K^* \geq \theta_l$ and $(q_l, q_h) = (K^*, K^*)$.

The closed-form expressions of $t_i^A(R)$ and $\Gamma_i^A(t, R)$ are summarized in Table B.5.

ii) Suppose $s_l < p \leq s_h$. In this case, $\Gamma^A(t)$ and the optimal all-unit discount in this case are as characterized in i) by replacing s_h with p.

iii) Suppose $s_l \geq p$.

iiia) If $\gamma > \Gamma^A(t), K^* = \frac{t+p\theta_h}{p(1-R)}, r^* = R, (q_l, q_h) = (\theta_l, K^*)$, and $\Pi^* = \gamma(t + p(\theta_h - \theta_l)) + p\theta_l$.

iiib) If $\gamma \leq \Gamma^A(t)$:

- For $t \leq \theta_h p(1-R) - p\theta_l$, $\Gamma^A(t) = \frac{t}{2t+pR\theta_h}, K^* = \frac{t+p\theta_l}{p(1-R)}, r^* = R, (q_l, q_h) = (K^*, \theta_h)$, and $\Pi^* = (1-\gamma)t + p(\theta_l + \gamma(\theta_h(1-R) - \theta_l))$.
- For $t > \theta_h p(1-R) - p\theta_l$, $\Gamma^A(t) = \frac{t}{t+p(\theta_h-\theta_l)}, K^* = \frac{t+p\theta_l}{p(1-R)}, r^* = R, (q_l, q_h) = (K^*, K^*)$, and $\Pi^* = p\theta_l + t$.

Fixed-amount discount:

i) $s_h < p$

For all $t \geq 0$, $\Gamma^F(t) = \frac{t+s_l\theta_l}{t+s_h\theta_h}$.

ia) If $\gamma > \Gamma^F(t), K^* \geq \max\{\theta_h, \frac{t+\theta_h s_h}{p(1-R)}\}, m^* \geq \max\{\theta_h(p - s_h) - t, M\}, (q_l, q_h) = (0, K^*)$ and $\Pi^* = \gamma(s_h\theta_h + t)$.

ib) If $\gamma \leq \Gamma^F(t), K^* \geq \max\{\theta_h, \frac{\theta_l s_l+t}{p(1-R)}\}, m^* \geq \max\{p\theta_h - \theta_l s_l - t, M\}, (q_l, q_h) = (K^*, K^*)$, and $\Pi^* = s_l\theta_l + t$.

ii) $s_l < p \leq s_h$

iia) If $\gamma > \Gamma^F(t)$, $K^* = \frac{p\theta_h+t}{p(1-R)}, m^* = \frac{R(p\theta_h+t)}{1-R}, (q_l, q_h) = (0, K^*)$, and $\Pi^* = \gamma(p\theta_h + t)$.

iib) If $\gamma < \Gamma^F(t)$, there exist $t_i^F(R), i \in \{1, 2, 3\}$ where

- For $t \leq t_1^F(R)$, $\Gamma^F(t) = \Gamma_1^F(t, R)$, $K^* \leq \theta_l$ and $(q_l, q_h) = (\theta_l, \theta_h)$.
- For $t_1^F(R) < t \leq t_2^F(R)$, $\Gamma^F(t) = \Gamma_2^F(t, R)$, $(q_l, q_h) = (\theta_l, \theta_h)$ or (K^*, θ_h).
- For $t_2^F(R) < t \leq t_3^F(R)$, $\Gamma^F(t) = \Gamma_3^F(t, R)$, $K^* \geq \theta_l$ and $(q_l, q_h) = (\theta_l, \theta_h)$ or (K^*, θ_h) or (K^*, K^*).
- For $t > t_3^F(R)$, $\Gamma^F(t) = \Gamma_4^F(t, R)$, $K^* \geq \theta_h$ and $(q_l, q_h) = (K^*, K^*)$.

The closed-form expressions of $t_i^F(R)$ and $\Gamma_i^F(t, R)$ are summarized in Table B.6.

iii) $s_l \geq p$

iiia) If $\gamma > \Gamma^F(t), K^* = \frac{t+p\theta_h}{p(1-R)}, m^* = \frac{R(p\theta_h+t)}{1-R}, (q_l, q_h) = (\theta_l, K^*)$, and $\Pi^* = \gamma(t + p(\theta_h - \theta_l)) + p\theta_l$.

iiib) If $\gamma \leq \Gamma^F(t)$:

- For $t \leq \theta_h p(1-R) - p\theta_l$, $\Gamma^F(t) = \frac{t(1-R)}{2t-Rt+pR\theta_l}, K^* = \frac{t+p\theta_l}{p(1-R)}, m^* = \frac{R(p\theta_l+t)}{p(1-R)}$, $(q_l, q_h) = (K^*, \theta_h)$, and $\Pi^* = \gamma p\theta_h + \frac{(1-\gamma-R)(p\theta_l+t)}{1-R}$.
- For $t > \theta_h p(1-R) - p\theta_l$, $\Gamma^F(t) = \frac{t}{t+p(\theta_h-\theta_l)}, K^* = \frac{t+p\theta_l}{p(1-R)}, m^* = \frac{R(p\theta_l+t)}{1-R}$, $(q_l, q_h) = (K^*, K^*)$, and $\Pi^* = p\theta_l + t$.

i	$t_i^A(R)$	$\Gamma_i^A(t,R)$
1	$\begin{cases} 0 & \text{if } R < \bar{R}(s_h) \text{ or } R \geq \bar{R}(s_l) \\ \frac{s_l\theta_h\theta_l(s_h-p(1-R))}{p(1-R)\theta_h - s_l\theta_l} & \text{if } \bar{R}(s_h) \leq R < \bar{R}(s_l) \end{cases}$	$\begin{cases} \frac{s_l\theta_l}{s_h\theta_h - s_l(\theta_h-\theta_l)} & \text{if } R < \bar{R}(s_h) \text{ or } R \geq \bar{R}(s_l) \\ \frac{\theta_l(t+s_l\theta_l)}{\theta_l(s_h\theta_h - s_l(\theta_h-\theta_l)) - t(\theta_h - 2\theta_l)} & \text{if } \bar{R}(s_h) \leq R < \bar{R}(s_l) \end{cases}$
2	$\begin{cases} \theta_l(p(1-\tilde{r})-s_l) & \text{if } R < \bar{R}(s_l) \\ \frac{1}{2}[\sqrt{\theta_h(\theta_h(s_h-s_l)^2 + 4s_l\theta_l(s_l - p(1-R)))} - \theta_h(s_h-s_l)] & \text{if } \bar{R}(s_l) \leq R < \bar{R}(\hat{s}^A) \\ \theta_h p(1-R) - s_l\theta_l & \text{if } R \geq \bar{R}(\hat{s}^A) \end{cases}$	$\begin{cases} \frac{p(1-\tilde{r})t}{(2p(1-\tilde{r})-s_l)t + \theta_h(s_h - p(1-\tilde{r}))(p(1-\tilde{r})-s_l)} & \text{if } R < \bar{R}(s_l) \\ \frac{s_l\theta_l}{t + s_h\theta_h - s_l(\theta_h - \theta_l)} & \text{if } R \geq \bar{R}(s_l) \end{cases}$
3	$\begin{cases} \theta_h p(1-\tilde{r}) - s_l\theta_l & \text{if } R < \bar{R}(\hat{s}^A) \\ \theta_h\hat{s}^A - s_l\theta_l & \text{if } R \geq \bar{R}(\hat{s}^A) \end{cases}$	$\begin{cases} \frac{t+s_l\theta_l}{2t+\theta_h(s_h-p(1-\tilde{r}))+s_l\theta_l} & \text{if } R < \bar{R}(\hat{s}^A) \\ \Gamma_2^A(t,R) & \text{if } R \geq \bar{R}(\hat{s}^A) \end{cases}$
4	N/A	$\frac{t+s_l\theta_l}{t+s_h\theta_h}$

Table B.5: Closed-form expressions of $t_i^A(R)$ and $\Gamma_i^A(t,R)$ for all-unit discount when $s_h < p$

$$
\begin{aligned}
\text{where } \bar{R}(x) &:= \frac{p-x}{p} \\
\hat{s}^A &:= \frac{1}{2\theta_h}[\sqrt{4s_l^2\theta_l(\theta_h-\theta_l) + (s_h\theta_h - s_l(\theta_h-\theta_l))^2} - s_h\theta_h + s_l(\theta_h+\theta_l)] \\
\tilde{r} &:= \max\{1 - \frac{s_h}{p}, R\} \\
\hat{s}^F &:= \frac{1}{2\theta_h}[\sqrt{p\theta_l(4s_l(\theta_h-\theta_l)+p\theta_l)} + \theta_l(2s_l - p)]
\end{aligned}
$$

i	$t_i^F(R)$	$\Gamma_i^F(t,R)$
1	0	$\dfrac{s_l}{p}$
2	$\begin{cases} \theta_l(p(1-R)-s_l) & \text{if } R < \bar{R}(s_l) \\ \frac{\theta_l}{2(1-R)}[s_l - p(1-R) + \sqrt{(s_l-p(1-R))(5s_l-4Rs_l-p(1-R))}] & \text{if } \bar{R}(s_l) \le R < \bar{R}(\hat{s}^F) \\ \theta_h p(1-R) - s_l\theta_l & \text{if } R \ge \bar{R}(\hat{s}^F) \end{cases}$	$\begin{cases} \frac{p(1-R)}{2p-pR-s_l} & \text{if } R < \bar{R}(s_l) \\ \frac{s_l\theta_l}{p\theta_l+t} & \text{if } R \ge \bar{R}(s_l) \end{cases}$
3	$\begin{cases} \theta_h p(1-R) - s_l\theta_l & \text{if } R < \bar{R}(\hat{s}^F) \\ \theta_h \hat{s}^F - s_l\theta_l & \text{if } R \ge \bar{R}(\hat{s}^F) \end{cases}$	$\begin{cases} \frac{(1-R)(t+s_l\theta_l)}{2t-Rt+s_l\theta_l} & \text{if } R < \bar{R}(\hat{s}^F) \\ \Gamma_2^F(t,R) & \text{if } R \ge \bar{R}(\hat{s}^F) \end{cases}$
4	N/A	$\dfrac{t+s_l\theta_l}{t+p\theta_h}$

Table B.6: Closed-form expressions of $t_i^F(R)$ and $\Gamma_i^F(t,R)$ for fixed-amount discount when $s_l < p \le s_h$

APPENDIX C

Proofs for Chapter 4

Proof of Proposition 4.1: 1. Suppose $I_1^t > 0$ and $I_2^t > 0$, so $A(\mathbf{I^t}) = \{1, 2\}$. We will first show that the retailer's profit-to-go in period t is strictly unimodal in the uniform price p^t. Notice from equation (4.3) that only $J^t(p^t, \mathbf{I^t})$ depends on p^t, where $J^t(p^t, \mathbf{I^t}) = \mu_1 \left[p^t - c_1^t - \beta\Delta_1^t(\mathbf{I^t})\right] + \mu_2 \left[p^t - c_2^t - \beta\Delta_2^t(\mathbf{I^t})\right]$. Hence, to show the unimodality of the profit function in p^t, it suffices to show that $\frac{\partial^2 J^t(p^t,\mathbf{I^t})}{\partial (p^t)^2} < 0$ whenever $\frac{\partial J^t(p^t,\mathbf{I^t})}{\partial p^t} = 0$. We derive the first- and second-order derivative of $J^t(p^t, \mathbf{I^t})$ with respect to p^t as follows:

$$\begin{aligned}
\frac{\partial J^t(p^t, \mathbf{Y^t})}{\partial p^t} &= \frac{\partial \mu_1}{\partial p^t}[p^t - c_1^t - \beta\Delta_1^t(Y_1^t, Y_2^t)] + \mu_1 + \\
&\quad \frac{\partial \mu_2}{\partial p^t}[p^t - c_2^t - \beta\Delta_2^t(Y_1^t, Y_2^t)] + \mu_2 \\
\frac{\partial^2 J^t(p^t, \mathbf{Y^t})}{\partial (p^t)^2} &= \frac{\partial^2 \mu_1}{\partial (p^t)^2}[p^t - c_1^t - \beta\Delta_1^t(Y_1^t, Y_2^t)] + 2\frac{\partial \mu_1}{\partial p^t} + \\
&\quad \frac{\partial^2 \mu_2}{\partial (p^t)^2}[p^t - c_2^t - \beta\Delta_2^t(Y_1^t, Y_2^t)] + 2\frac{\partial \mu_2}{\partial p^t},
\end{aligned} \tag{C.1}$$

where

$$
\begin{aligned}
\frac{\partial \mu_i}{\partial p^t} &= -\mu_i \mu_0 \\
\frac{\partial^2 \mu_i}{\partial (p^t)^2} &= \mu_i \mu_0 (2\mu_0 - 1), \ i \in A(\mathbf{I^t}).
\end{aligned} \tag{C.2}
$$

From equation (C.1), if $\frac{\partial J^t(p^t,\mathbf{I^t})}{\partial p^t} = 0$, then

$$
p^t - c_1^t - \beta \Delta_1^t(\mathbf{I^t}) = \frac{\left(-\frac{\partial \mu_2}{\partial p^t}[p^t - c_2^t - \beta \Delta_2^t(\mathbf{I^t})] - \mu_1 - \mu_2\right)}{\frac{\partial \mu_1}{\partial p^t}} \tag{C.3}
$$

Substituting equations (C.3) and (C.2) into (**??**), we obtain

$$
\frac{\partial^2 J^t(p^t,\mathbf{I^t})}{\partial (p^t)^2}\Big|_{\frac{\partial J^t(p^t,\mathbf{I^t})}{\partial p^t}=0} = -\mu_1 - \mu_2 < 0.
$$

This shows that the retailer's profit-to-go in period t is strictly unimodal in the uniform price p^t, so the optimal p^t is unique. Notice from equation (C.2) that $\frac{\partial \mu_i}{\partial p^t} < 0, i \in A(\mathbf{I^t})$. Hence, it is easy to see from equation (C.1) that $\frac{\partial J^t(p^t,\mathbf{I^t})}{\partial p^t} > 0$ for sufficiently small p^t, and $\frac{\partial J^t(p^t,\mathbf{I^t})}{\partial p^t} < 0$ for sufficiently large p^t. This implies that the optimal uniform price p_{un}^{t*} is an interior solution, characterized by the first-order condition, $\frac{\partial J^t(p^t,\mathbf{I^t})}{\partial p^t} = 0$. We can rearrange the first-order condition and apply the definition of $\mu_i, i \in \{0,1,2\}$, from equation (4.1) to obtain: $p_{un}^{t*} = \frac{\exp(v_1)(c_1^t+\beta\Delta_1^t(\mathbf{I^t}))+\exp(v_2)(c_2^t+\beta\Delta_2^t(\mathbf{I^t}))}{\exp(v_1)+\exp(v_2)} + \frac{1}{\mu_0(p_{un}^{t*})}$. Applying this to the definition of J^t, we have $J^t(p_{un}^{t*},\mathbf{I^t}) = \exp(v_1 - p_{un}^{t*}) + \exp(v_2 - p_{un}^{t*}) = R(p_{un}^{t*}, A(\mathbf{I^t}))$. Hence, from (4.3), the retailer's profit-to-go is $V_{un}^t(\mathbf{I^t}) = R_{un}^{t*}\lambda^t + \beta V_{un}^{t-1}(\mathbf{I^t})$.

2. Suppose $I_1^t = 0$ and $I_2^t > 0$, so $A^t(\mathbf{I^t}) = \{2\}$. Then, $\mu_1 = 0$, and $J^t(p^t,\mathbf{I^t}) = \mu_2[p^t - c_2^t - \beta\Delta_2^t(0, I_2^t)]$. Following the same procedure as part 1., we will show that the optimal uniform price is unique and is characterized by the first-order condition by showing $\frac{\partial^2 J^t(p^t,\mathbf{I^t})}{\partial (p^t)^2} < 0$ whenever $\frac{\partial J^t(p^t,\mathbf{I^t})}{\partial p^t} = 0$. The first- and second-order derivative

are given by:

$$\frac{\partial J^t(p^t,\mathbf{I^t})}{\partial p^t} = \frac{\partial \mu_2}{\partial p^t}[p^t - c_2^t - \beta\Delta_2^t(\mathbf{I^t})] + \mu_2 \tag{C.4}$$

$$\frac{\partial^2 J^t(p^t,\mathbf{I^t})}{\partial (p^t)^2} = \frac{\partial^2 \mu_2}{\partial (p^t)^2}[p^t - c_2^t - \beta\Delta_2^t(\mathbf{I^t})] + 2\frac{\partial \mu_2}{\partial p^t}, \tag{C.5}$$

where

$$\frac{\partial \mu_2}{\partial p^t} = -\mu_2\mu_0 \tag{C.6}$$

$$\frac{\partial^2 \mu_2}{\partial (p^t)^2} = \mu_2\mu_0(2\mu_0 - 1)$$

From equation (C.4), if $\frac{\partial J^t(p^t,\mathbf{I^t})}{\partial p^t} = 0$, then

$$p^t - c_2^t - \beta\Delta_2^t(\mathbf{I^t}) = \frac{-\mu_2}{\frac{\partial \mu_2}{\partial p^t}} \tag{C.7}$$

Substituting equations (C.7) and (C.6) into (C.5), we obtain: $\frac{\partial^2 J^t(p^t,\mathbf{I^t})}{\partial (p^t)^2}|_{\frac{\partial J^t(p^t,\mathbf{I^t})}{\partial p^t}=0} = -\mu_2 < 0$. Hence, the optimal uniform price p_{un}^{t*} is characterized by the first-order condition, $\frac{\partial J^t(p^t,\mathbf{I^t})}{\partial p^t} = 0$. Rearranging the condition, we obtain: $p_{un}^{t*} = c_2^t + \beta\Delta_2^t(\mathbf{I^t}) + \frac{1}{\mu_0(p_{un}^{t*})}$. Applying this to equation (4.3) results in $J^t(p_{un}^{t*},\mathbf{I^t}) = \exp(v_2 - p_{un}^{t*}) = R(p_{un}^{t*}, A(\mathbf{I^t}))$, and the retailer's profit-to-go of $V_{un}^t(\mathbf{I^t}) = R_{un}^{t*}\lambda^t + \beta V_{un}^{t-1}(\mathbf{I^t})$.

Proof of Proposition 4.2: 1. Suppose $I_1^t > 0$ and $I_2^t > 0$, so $A^t(\mathbf{I^t}) = \{1,2\}$. We will show that the retailer's profit-to-go in period t is strictly unimodal in the price vector $\mathbf{p^t} = (p_1^t, p_2^t)$. Notice from equation (4.4) that only $J^t(\mathbf{p^t},\mathbf{I^t})$ depends on $\mathbf{p^t}$, where $J^t(\mathbf{p^t},\mathbf{I^t}) = \mu_1[p_1^t - c_1^t - \beta\Delta_1^t(\mathbf{I^t})] + \mu_2[p_2^t - c_2^t - \beta\Delta_2^t(\mathbf{I^t})]$. Hence, we will show the result by showing that J^t is unimodal in $\mathbf{p^t}$. We derive the first-order derivative of $J^t(\mathbf{p^t},\mathbf{I^t})$ with respect to $p_i^t, i \in \{1,2\}$, as follows:

$$\frac{\partial J^t(\mathbf{p^t},\mathbf{I^t})}{\partial p_i^t} = \frac{\partial \mu_1}{\partial p_i^t}[p_1^t - c_1^t - \beta\Delta_1^t(\mathbf{I^t})] + \frac{\partial \mu_2}{\partial p_i^t}[p_2^t - c_2^t - \beta\Delta_2^t(\mathbf{I^t})] + \mu_i \tag{C.8}$$

where

$$\frac{\partial \mu_i}{\partial p_i^t} = -\mu_i(1-\mu_i) \tag{C.9}$$
$$\frac{\partial \mu_j}{\partial p_i^t} = \mu_j \mu_i,\ i,j \in A(\mathbf{I^t}), i \neq j.$$

Notice from equation (C.8) that $\frac{\partial J^t(\mathbf{p^t},\mathbf{I^t})}{\partial p_i^t} > 0$ for sufficiently small p_i^t since $\frac{\partial \mu_i}{\partial p_i^t} < 0$. Likewise, $\frac{\partial J^t(\mathbf{p^t},\mathbf{I^t})}{\partial p_i^t} < 0$ for sufficiently large p_i^t. This implies that an optimal solution must be an interior solution and satisfy the first-order conditions with respect to p_i^t: $\frac{\partial J^t(\mathbf{p^t},\mathbf{I^t})}{\partial p_i^t} = 0, i \in A(\mathbf{I^t})$.

Substituting (C.9) into (C.8), we obtain the first-order conditions:

$$\begin{aligned}\frac{\partial J^t(\mathbf{p^t},\mathbf{I^t})}{\partial p_i^t} &= \mu_i\left[-(1-\mu_i)\left(p_i^t - c_i^t - \beta\Delta_i^t(\mathbf{I^t})\right) + \mu_j\left(p_j^t - c_j^t - \beta\Delta_j^t(\mathbf{I^t})\right) + 1\right] \\ &= \mu_i\left[J^t(\mathbf{p^t},\mathbf{I^t}) + 1 - \left(p_i^t - c_i^t - \beta\Delta_i^t(\mathbf{I^t})\right)\right] = 0,\ i,j \in A^t(\mathbf{I^t}), i \neq j.\end{aligned}$$

Since $\mu_i > 0$, a p_i^t which satisfies the first-order conditions must be such that $p_i^t - c_i^t - \beta\Delta_i^t(\mathbf{I^t}) = J^t(\mathbf{p^t},\mathbf{I^t}) + 1$ for $i \in A^t(\mathbf{I^t})$. Hence, an optimal $\mathbf{p^t} = (p_1^t, p_2^t)$ must satisfy $p_1^t - c_1^t - \beta\Delta_1^t(\mathbf{I^t}) = p_2^t - c_2^t - \beta\Delta_2^t(\mathbf{I^t})$. Let $m := p_i^t - c_i^t - \beta\Delta_i^t(\mathbf{I^t})$. Then, substituting m in the definition of J^t, we have $J^t(\mathbf{p^t},\mathbf{I^t}) = (\mu_1+\mu_2)m$. From the first-order conditions, $J^t(\mathbf{p^t},\mathbf{I^t}) = m - 1$. Hence, we have $J^t(\mathbf{p^t},\mathbf{I^t}) = m - 1 = (\mu_1+\mu_2)m$. Solving this yields $m = \frac{1}{1-\mu_1-\mu_2} = \frac{1}{\mu_0}$. Thus, an optimal solution must satisfy $p_i^{t*} = m + c_i^t + \beta\Delta_i^t(\mathbf{I^t}) = \frac{1}{\mu_0(p_1^{t*}, p_2^{t*})} + c_i^t + \beta\Delta_i^t(\mathbf{I^t}), i \in A^t(\mathbf{I^t})$.

Next, we will show that the optimal price $\mathbf{p^t}$ is unique. For notational simplicity, we will drop $\mathbf{I^t}$ from the expressions below. Since $p_1^t - c_1^t - \beta\Delta_1^t = p_2^t - c_2^t - \beta\Delta_2^t = m$, J^t can be written as a function of m as follows:

$$J^t(m) = [\mu_1(\mathbf{p^t}(m)) + \mu_2(\mathbf{p^t}(m))]m = [1 - \mu_0(\mathbf{p^t}(m))]m \tag{C.10}$$

where $\mathbf{p^t}(m) = (m + c_1^t + \beta\Delta_1^t, m + c_2^t + \beta\Delta_2^t)$, and

$\mu_0(\mathbf{p^t}(m)) = \frac{1}{1+\exp(v_1-m-c_1^t-\beta\Delta_1^t)+\exp(v_2-m-c_2^t-\beta\Delta_2^t)}$.

The first-order derivative of J^t is given by $\frac{\partial J^t(m)}{\partial m} = 1 - \mu_0(\mathbf{p^t}(m)) - m\frac{\partial \mu_0(\mathbf{p^t}(m))}{\partial m}$. Notice that

$$\begin{aligned}\frac{\partial \mu_0(\mathbf{p^t}(m))}{\partial m} &= \frac{\exp(v_1 - m - c_1^t - \beta\Delta_1^t) + \exp(v_2 - m - c_2^t - \beta\Delta_2^t)}{(1 + \exp(v_1 - m - c_1^t - \beta\Delta_1^t) + \exp(v_2 - m - c_2^t - \beta\Delta_2^t))^2} \\ &= \mu_0(\mathbf{p^t}(m))[1 - \mu_0(\mathbf{p^t}(m))]\end{aligned} \tag{C.11}$$

Applying (C.11) and (C.10) to $\frac{\partial J^t(m)}{\partial m}$, we obtain

$$\frac{\partial J^t(m)}{\partial m} = 1 - \mu_0(\mathbf{p^t}(m))[1 + J^t(m)] \tag{C.12}$$

Then, we can derive the second-order derivative of J^t as:

$$\frac{\partial^2 J^t(m)}{\partial m^2} = -\frac{\partial \mu_0(\mathbf{p^t}(m))}{\partial m}[1 + J^t(m)] - \mu_0(\mathbf{p^t}(m))\frac{\partial J^t(m)}{\partial m} \tag{C.13}$$

Notice from (C.12) and (C.13) that whenever $\frac{\partial J^t(m)}{\partial m} = 0$, $\frac{\partial^2 J^t(m)}{\partial m^2} = -\frac{\partial \mu_0(\mathbf{p^t}(m))}{\partial m}[1 + J^t(m)] = -\mu_0(\mathbf{p^t}(m))(1 - \mu_0(\mathbf{p^t}(m)))[1 + J^t(m)] < 0$. This implies the unimodality of J^t with respect to m, which in turn implies the uniqueness of the optimal solution $\mathbf{p^t}$. Hence, the unique optimal prices are characterized by: $p_i^{t*} = c_i^t + \frac{1}{\mu_0(p_1^{t*}, p_2^{t*})} + \beta\Delta_i^t(I^t)$. Applying this to equation (4.4), we have $J^t(\mathbf{p^{t*}}, \mathbf{I^t}) = \exp(v_1 - p_1^{t*}) + \exp(v_2 - p_2^{t*}) = R(\mathbf{p^{t*}}, A(\mathbf{I^t}))$, resulting in the retailer's profit-to-go of $V_{dn}^t(\mathbf{I^t}) = R_{ut}^{t*}\lambda^t + \beta V_{dn}^{t-1}(\mathbf{I^t})$.

2. Suppose $I_1^t = 0$ and $I_2^t > 0$, so $A^t(\mathbf{I^t}) = \{2\}$. Since transshipment is not allowed, we have $I_1^t = 0$ and $\mu_1 = 0$ for $t, t+1, ..., T$. Then, the retailer's pricing problem in period t and on is the same as that under the uniform pricing policy without transshipment. It follows immediately that p_2^{t*} is the same as p_{un}^{t*} and the retailer receives the same profit as characterized in Proposition 4.1 part 2.

Proof of Corollary 4.1: First, notice from Proposition 4.2 part 2 that if $I_1^t = 0$ and $I_2^t > 0$, then the optimal price under both pricing policies are the same, resulting in the same sale ratio. Hence, the result trivially holds for this case. Next, we will show the result for $I_1^t > 0$ and $I_2^t > 0$. From Proposition 4.1 and 4.2, we show that the retailer's expected profit under the optimal uniform pricing policy is given by $V_{un}^t(\mathbf{I^t}) = R_{un}^{t*}\lambda^t + \beta V_{un}^{t-1}(\mathbf{I^t})$, and the retailer's expected profit under the optimal price differentiating policy is given by $V_{dn}^t(\mathbf{I^t}) = R_{dn}^{t*}\lambda^t + \beta V_{dn}^{t-1}(\mathbf{I^t})$. If the retailer adopts the optimal price differentiating policy in period t only, then $V_{dn}^{t-1}(\mathbf{I^t}) = V_{un}^{t-1}(\mathbf{I^t})$. Let $\hat{V}_{dn}^t$ denote the retailer's expected profit under the optimal price differentiating policy in period t only. Then, we have $\hat{V}_{dn}^t(\mathbf{I^t}) = R_{dn}^{t*}\lambda^t + \beta V_{un}^{t-1}(\mathbf{I^t})$. Since the uniform pricing policy is a special case of the price differentiating policy, the retailer's expected profit under the optimal price differentiating policy in period t must weakly dominate the profit under the optimal uniform pricing policy. That is, the profit difference is nonnegative: $\hat{V}_{dn}^t - V_{un}^t = R_{dn}^{t*}\lambda^t - R_{un}^{t*}\lambda^t \geq 0$. Since $\lambda^t \geq 0$, this implies $R_{dn}^{t*} \geq R_{un}^{t*}$.

Proof of Proposition 4.3: From the proof of Corollary 4.1, we have shown that the benefit of adopting the optimal price differentiation policy in period t only is given by $\hat{V}_{dn}^t - V_{un}^t = (R_{dn}^{t*} - R_{un}^{t*})\lambda^t$. Since $\lambda^t \geq 0$, it follows immediately that the benefit from price differentiation is monotonically increasing with the sale ratio difference: $R_{dn}^{t*} - R_{un}^{t*}$.

Proof of Proposition 4.4: 1. We will consider the two cases of i) $I_1^t = 0$ and $I_2^t > 1$, and ii) $I_1^t = 0$ and $I_2^t = 1$.

i) Suppose $I_1^t = 0$ and $I_2^t > 1$. To show when a transshipment from channel 2 to 1 in period t is optimal or not, we compare the retailer's profit with and without such a transshipment. If the retailer does not make a transshipment, then his profit under the optimal uniform pricing strategy is as given in Proposition 4.1 part 2: $V_{un}^t(\mathbf{I^t}) = R_{un}^{t*}\lambda^t + \beta V_{un}^{t-1}(\mathbf{I^t})$, where $\mathbf{I^t} = (0, I_2^t)$. If he makes a transshipment in period t, the resulting inventory level is $\mathbf{Y^t} = (Y_1^t = 1, Y_2^t = I_2^t - 1)$, and he incurs a transshipment cost of m_{21}. Notice that since $Y_1^t > 0$ and $Y_2^t > 0$, the optimal uniform price after the transshipment, p_{ut}^{t*}, can be characterized by Proposition 4.1 part 1 by replacing $\mathbf{I^t}$ with $\mathbf{Y^t}$. Hence, the retailer's profit with the transshipment in period t only is given by $\hat{V}^t_{ut}(\mathbf{I^t}) = R^{t*}_{ut}\lambda^t + \beta V^{t-1}_{un}(\mathbf{Y^t}) - m_{21}$, where $\mathbf{Y^t} = (1, I^t_2 - 1)$. The profit difference is $\hat{V}^t_{ut}(\mathbf{I^t}) - V^t_{un}(\mathbf{I^t}) = (R^{t*}_{ut} - R^{t*}_{un})\lambda^t + \beta\left[V^{t-1}_{un}(1, I^t_2 - 1) - V^{t-1}_{un}(0, I^t_2)\right] - m_{21}$. Notice that $V^{t-1}_{un}(1, I^t_2 - 1) - V^{t-1}_{un}(0, I^t_2) = V^{t-1}_{un}(1, I^t_2 - 1) - V^{t-1}_{un}(0, I^t_2 - 1) - V^{t-1}_{un}(0, I^t_2) + V^{t-1}_{un}(0, I^t_2 - 1) = \Delta^t_1(1, I^t_2 - 1) - \Delta^t_2(0, I^t_2) = \Delta^t_{12}(\mathbf{I^t})$, by definition. Hence, it is optimal to make the transshipment if and only if $m_{21} \leq (R^{t*}_{ut} - R^{t*}_{un})\lambda^t + \beta\Delta^t_{12}(\mathbf{I^t})$.

ii) Suppose $I^t_1 = 0$ and $I^t_2 = 1$. The only difference in this case from i) is that after the transshipment is made from channel 2 to 1, the inventory level at channel 2 becomes 0. Since $\mathbf{Y^t} = (1, 0)$, the optimal uniform price after transshipment, p^{t*}_{ut}, is characterized by Proposition 4.1 part 2 by replacing $\mathbf{I^t}$ with $\mathbf{Y^t}$, c^t_2 with c^t_1, and v_2 with v_1. Hence, the retailer's profit after the transshipment is made in period t only is given by $\hat{V}^t_{ut}(\mathbf{I^t}) = \exp(v_1 - p^{t*}_{ut})\lambda^t + \beta V^{t-1}_{un}(\mathbf{Y^t}) - m_{21} = R^{t*}_{ut}\lambda^t + \beta V^{t-1}_{un}(\mathbf{Y^t}) - m_{21}$, where $\mathbf{Y^t} = (1, 0)$. This results in the profit difference of $\hat{V}^t_{ut}(\mathbf{I^t}) - V^t_{un}(\mathbf{I^t}) = (R^{t*}_{ut} - R^{t*}_{un})\lambda^t + \beta\left[V^{t-1}_{un}(1, 0) - V^{t-1}_{un}(0, 1)\right] - m_{21} = (R^{t*}_{ut} - R^{t*}_{un})\lambda^t + \beta\left[V^{t-1}_{un}(1, I^t_2 - 1) - V^{t-1}_{un}(0, I^t_2)\right] - m_{21}$. Notice that $V^{t-1}_{un}(1, I^t_2 - 1) - V^{t-1}_{un}(0, I^t_2) = \Delta^t_{12}(\mathbf{I^t})$ as shown in part i). Hence, the result follows.

2. We will show $p^{t*}_{ut} \geq p^{t*}_{un}$ for the two sets of conditions, stated in the proposition: a) $I^t_2 > 1$, $\beta\Delta^t_{12}(\mathbf{I^t}) \geq c^t_2 - c^t_1$, and $\Delta^t_2(1, I^t_2 - 1) \geq \Delta^t_2(0, I^t_2)$, and b) $I^t_2 = 1$, $v_1 \geq v_2$, and $\beta\Delta^t_{12}(\mathbf{I^t}) \geq c^t_2 - c^t_1$.

a) Suppose $I^t_2 > 1$, $\beta\Delta^t_{12}(\mathbf{I^t}) \geq c^t_2 - c^t_1$, and $\Delta^t_2(1, I^t_2 - 1) \geq \Delta^t_2(0, I^t_2)$. First, recall from the proof of Proposition 4.1 part 2 that under the uniform pricing policy without transshipment, and for $I^t_1 = 0$ and $I^t_2 > 0$, we have $J^t_{un}(p^t, \mathbf{I^t})$ is unimodal in p^t, and p^{t*}_{un} is such that $\frac{\partial J^t_{un}(p^t,\mathbf{I^t})}{\partial p^t}|_{p^t = p^{t*}_{un}} = 0$. Hence, to show $p^{t*}_{ut} \geq p^{t*}_{un}$, it suffices to show $\frac{\partial J^t_{un}(p^t,\mathbf{I^t})}{\partial p^t}|_{p^t = p^{t*}_{ut}} \leq 0$. From equation (C.4) and (C.6), we obtain $\frac{\partial J^t_{un}(p^t,\mathbf{I^t})}{\partial p^t} = \mu_2(p^t, A(\mathbf{Y^t}))\left[-\mu_0(p^t, A(\mathbf{Y^t}))\left(p^t - c^t_2 - \beta\Delta^t_2(\mathbf{I^t})\right) + 1\right] = \frac{\exp(v_2+p^t)\left[c^t_2 + \beta\Delta^t_2(\mathbf{I^t}) + 1 + \exp(v_2 - p^t) - p^t\right]}{(\exp(p^t) + \exp(v_2))^2}$ since $A(\mathbf{Y^t}) = \{2\}$. Let $C(p^t)$ denote $\frac{\exp(v_2+p^t)}{(\exp(p^t)+\exp(v_2))^2}$, which is positive for any p^t. Now, consider p^{t*}_{ut}. Since $I^t_2 > 1$, we have $Y^t_1 = 1$ and $Y^t_2 > 0$ after the transshipment is made. Hence, p^{t*}_{ut} satisfies the condition in Proposition 4.1 part 1 for $\mathbf{I^t} = \mathbf{Y^t} = (1, I_2^t - 1)$, which is $p_{ut}^{t*} = \frac{\exp(v_1)(c_1^t+\beta\Delta_1^t(\mathbf{Y^t}))+\exp(v_2)(c_2^t+\beta\Delta_2^t(\mathbf{Y^t}))}{\exp(v_1)+\exp(v_2)} + \frac{1}{\mu_0(p_{ut}^{t*})}$, where $\frac{1}{\mu_0(p_{ut}^{t*})} = 1 + \exp(v_1 - p_{ut}^{t*}) + \exp(v_2 - p_{ut}^{t*})$ since $A(\mathbf{Y^t}) = \{1, 2\}$. Applying this to $\frac{\partial J_{un}^t(p^t,\mathbf{I^t})}{\partial p^t}$, we obtain

$$\begin{aligned}\frac{\partial J_{un}^t(p^t,\mathbf{I^t})}{\partial p^t}|_{p^t=p_{ut}^{t*}} &= C(p_{ut}^{t*})\left[c_2^t + \beta\Delta_2^t(0, I_2^t) - \exp(v_1 - p_{ut}^{t*}) - \right.\\ &\left.\left[\frac{exp(v_1)(c_1^t + \beta\Delta_1^t(1, I_2^t - 1)) + exp(v_2)(c_2^t + \beta\Delta_2^t(1, I_2^t - 1))}{exp(v_1) + exp(v_2)}\right]\right].\end{aligned}$$

Given that $\beta\Delta_{12}^t \geq c_2^t - c_1^t$ and $\Delta_2^t(1, I_2^t - 1) \geq \Delta_2^t(0, I_2^t)$, we have $c_1^t + \beta\Delta_1^t(1, I_2^t - 1) \geq c_2^t + \beta\Delta_2^t(0, I_2^t)$. Hence, $\left[\frac{exp(v_1)(c_1^t+\beta\Delta_1^t(1,I_2^t-1))+exp(v_2)(c_2^t+\beta\Delta_2^t(1,I_2^t-1))}{exp(v_1)+exp(v_2)}\right] \geq c_2^t + \beta\Delta_2^t(0, I_2^t)$. Note also that $\exp(v_1 - p_{ut}^{t*}) \geq 0$. Thus, it follows that $\frac{\partial J_{un}^t(p^t,\mathbf{I^t})}{\partial p^t}|_{p^t=p_{ut}^{t*}} \leq 0$, which implies $p_{ut}^{t*} \geq p_{un}^{t*}$.

b) Suppose $I_2^t = 1$, $v_1 \geq v_2$, and $\beta\Delta_{12}^t(\mathbf{I^t}) \geq c_2^t - c_1^t$. Following the same procedure as in a), we will show $p_{ut}^{t*} \geq p_{un}^{t*}$ by showing $\frac{\partial J_{un}^t(p^t,\mathbf{I^t})}{\partial p^t}|_{p^t=p_{ut}^{t*}} \leq 0$. Notice that the only difference in this case from a) is that since $I_2^t = 1$, after the transshipment is made from channel 2 to channel 1, the inventory level at channel 2 becomes zero. Hence, p_{ut}^{t*} in this case is characterized by Proposition 4.1 part 2, with $\mathbf{I^t}$ replaced by $\mathbf{Y^t} = (1, 0)$, c_2^t replaced by c_1^t, v_2 replaced by v_1, and $\Delta_2^t(\mathbf{I^t})$ replaced by $\Delta_1^t(\mathbf{Y^t})$. That is, $p_{ut}^{t*} = c_1^t + \beta\Delta_1^t(\mathbf{Y^t}) + \frac{1}{\mu_0(p_{ut}^{t*})}$. Applying this to $\frac{\partial J_{un}^t(p^t,\mathbf{I^t})}{\partial p^t}$, we obtain $\frac{\partial J_{un}^t(p^t,\mathbf{I^t})}{\partial p^t}|_{p^t=p_{ut}^{t*}} = C(p_{ut}^{t*})\left[c_2^t + \beta\Delta_2^t(0, I_2^t) + \exp(v_2 - p_{ut}^{t*}) - c_1^t - \beta\Delta_1^t(1, I_2^t - 1) - \exp(v_1 - p_{ut}^{t*})\right] = C(p_{ut}^{t*})\left[c_2^t - c_1^t - \beta\Delta_{12}^t + \exp(v_2 - p_{ut}^{t*}) - \exp(v_1 - p_{ut}^{t*})\right]$. Given that $\beta\Delta_{12}^t(\mathbf{I^t}) \geq c_2^t - c_1^t$, we have $c_2^t - c_1^t - \beta\Delta_{12}^t \leq 0$. From $v_1 \geq v_2$, we have $\exp(v_2 - p_{ut}^{t*}) - \exp(v_1 - p_{ut}^{t*}) \leq 0$. Hence, it follows that $\frac{\partial J_{un}^t(p^t,\mathbf{I^t})}{\partial p^t}|_{p^t=p_{ut}^{t*}} \leq 0$, which implies $p_{ut}^{t*} \geq p_{un}^{t*}$.

3. From part 1, we have shown that the retailer's benefit from transshipment in period t is given by $\hat{V}_{ut}^t(\mathbf{I^t}) - V_{un}^t(\mathbf{I^t}) = (R_{ut}^{t*} - R_{un}^{t*})\lambda^t + \beta\Delta_{12}^t(\mathbf{I^t}) - m_{21}$. Since $\lambda^t \geq 0$, it is immediate to see that the benefit of transshipment is monotonically increasing in

the sale ratio $R_{ut}^{t*} - R_{un}^{t*}$. For m_{21}, notice that p_{un}^{t*}, p_{ut}^{t*}, and $\Delta_{12}^{t}(\mathbf{I^t})$ are independent of the transshipment cost m_{21}. Hence, R_{ut}^{t*} and R_{un}^{t*} are also independent of m_{21}. Then, it is easy to see that the benefit of transshipment is monotonically decreasing in m_{21}.

Now, consider the effect of c_1^t on the benefit of transshipment. Let $\Delta\Pi_{ut}^t = \hat{V}_{ut}^t(\mathbf{I^t}) - V_{un}^t(\mathbf{I^t})$ denote the benefit of transshipment. Then, by total differentiation, we have $\frac{d\Delta\Pi_{ut}^t}{dc_1^t} = \frac{\partial\Delta\Pi_{ut}^t}{\partial c_1^t} + \frac{\partial\Delta\Pi_{ut}^t}{\partial p_{un}^{t*}}\frac{\partial p_{un}^{t*}}{\partial c_1^t} + \frac{\partial\Delta\Pi_{ut}^t}{\partial p_{ut}^{t*}}\frac{\partial p_{ut}^{t*}}{\partial c_1^t}$. Notice that p_{un}^{t*} is independent of c_1^t since without a transshipment, the customer can never buy from channel 1 in period t. Hence, $\frac{\partial p_{un}^{t*}}{\partial c_1^t} = 0$. On the other hand, we have seen from part 2 that p_{ut}^{t*} is a function of c_1^t. By the implicit function theorem, $\frac{\partial p_{ut}^{t*}}{\partial c_1^t} = -(\frac{\partial G_{ut}}{\partial c_1^t})/(\frac{\partial G_{ut}}{\partial p_{ut}^{t*}})$, where G_{ut} is the first-order condition for p_{ut}^{t*}. Applying these to $\frac{d\Delta\Pi_{ut}^t}{dc_1^t}$, we obtain $\frac{d\Delta\Pi_{ut}^t}{dc_1^t} = \frac{-\exp(v_1)\lambda^t}{\exp(p_{ut}^{t*})+\exp(v_1)+\exp(v_2)} \leq 0$ if $I_2^t > 1$; and $\frac{d\Delta\Pi_{ut}^t}{dc_1^t} = \frac{-\exp(v_1)\lambda^t}{\exp(p_{ut}^{t*})+\exp(v_1)} \leq 0$ if $I_2^t = 1$. Hence, the benefit from the transshipment from channel 2 to channel 1 in period t is monotonically decreasing in the unit transaction cost c_1^t.

Proof of Theorem 4.1: 1. Suppose $\Delta_2^t(1, I_2^t - 1) \geq \Delta_2^t(0, I_2^t)$. Since $I_1^t = 0$ and $I_2^t > 0$, we know from Proposition 4.2 part 2 that without transshipment, the optimal price at channel 2, $p_{2,10}^{t*}$, is the same as the optimal uniform price without transshipment. Hence, the retailer's profit is unimodal in p_2^t, and $p_{2,10}^{t*}$ satisfies the first-order condition $\frac{\partial J_{dn}^t(p_2^t,\mathbf{I^t})}{\partial p_2^t}|_{p_2^t=p_{2,10}^{t*}} = 0$. Let $p_{2,11}^{t*}$ denote the optimal price at channel 2 after a transshipment is made from channel 2 to channel 1. We will show $p_{2,11}^{t*} \geq p_{2,10}^{t*}$ by showing $\frac{\partial J_{dn}^t(p_2^t,\mathbf{I^t})}{\partial p_2^t}|_{p_2^t=p_{2,11}^{t*}} \leq 0$. From equation (C.4) and (C.6), we derive:

$$\frac{\partial J_{dn}^t(p_2^t,\mathbf{I^t})}{\partial p_2^t} = \frac{\exp(v_2+p_2^t)\,[c_2^t+\beta\Delta_2^t(\mathbf{I^t})+1+\exp(v_2-p_2^t)-p_2^t]}{(\exp(p_2^t)+\exp(v_2))^2}. \tag{C.14}$$

Let $C(p_2^t)$ denote $\frac{\exp(v_2+p_2^t)}{(\exp(p_2^t)+\exp(v_2))^2}$, which is positive for any p_2^t. Now, consider $p_{2,11}^{t*}$. After a transshipment is made, $Y_1^t = 1$ and $Y_2^t = I_2^t - 1 > 0$. Hence, we have that $p_{2,11}^{t*}$ is as characterized in Proposition 4.2 part 1 by replacing $\mathbf{I^t}$ with $\mathbf{Y^t} = (1, I_2^t - 1)$: $p_{2,11}^{t*} =$

$c_2^t+1+\exp(v_1-p_{1,11}^{t*})+\exp(v_2-p_{2,11}^{t*})+\beta\Delta_2^t(1,I_2^t-1)$. Applying this to (C.14), we obtain $\frac{\partial J_{dn}^t(p_2^t,\mathbf{I^t})}{\partial p_2^t}|_{p_2^t=p_{2,11}^{t*}} = C(p_{2,11}^{t*})\left[-\exp(v_1-p_{1,11}^{t*})-\beta(\Delta_2^t(1,I_2^t-1)-\Delta_2^t(0,I_2^t))\right] \leq 0$ since $\Delta_2^t(1,I_2^t-1) \geq \Delta_2^t(0,I_2^t)$.

2. Let m_{21}^t denote the transshipment cost from channel 2 to channel 1 in period t. We will first characterize the optimal transshipment decision in period t under the uniform pricing and price differentiating policies by showing that there exist m_L and m_H, $m_L \leq m_H$, such that:

i) If $m_{21}^t \leq m_L$, it is optimal to transship a unit from channel 2 to channel 1 under both uniform pricing and price differentiating policies.

ii) If $m_L < m_{21}^t \leq m_H$, it is optimal to transship a unit from channel 2 to channel 1 under the price differentiating policy, but it is optimal to not transship under the uniform pricing policy.

iii) If $m_{21}^t > m_H$, it is optimal to not transship under both uniform pricing and price differentiating policies.

First, consider the transshipment decision under the uniform pricing policy. From Proposition 4.4 part 1 and 3, it is implied that there exists m_{ut} such that it is optimal to transship a unit from channel 2 to channel 1 in period t if $m_{21}^t \leq m_{ut}$, and it is optimal not to transship if $m_{21}^t > m_{ut}$. That is, $\hat{V}_{ut}^t(\mathbf{I^t}) \geq V_{un}^t(\mathbf{I^t})$ if and only if $m_{21}^t \leq m_{ut}$, where $\hat{V}_{ut}^t(\mathbf{I^t})$ denotes the retailer's expected profit when making a transshipment from channel 2 to channel 1 and charge the optimal uniform price in period t. Now, consider the transshipment decision under the price differentiating policy. Let $\hat{V}_{dt}^t(\mathbf{I^t})$ denote the retailer's expected profit when making a transshipment from channel 2 to channel 1 and charge the optimal differentiating prices in period t. Notice that $\hat{V}_{dt}^t(\mathbf{I^t})$ is monotonically decreasing in m_{21}^t since the retailer needs to pay a fixed cost of m_{21}^t when making the transshipment. On the other hand, if the retailer does not make a transshipment, he can only sell through channel 2 and will receive the profit of $V_{un}^t(\mathbf{I^t})$, independent of m_{21}^t, which is the same as that under the uniform pricing without transshipment. Hence, the profit difference under the price differentiating policy, $\hat{V}_{dt}^t(\mathbf{I^t}) - V_{un}^t(\mathbf{I^t})$, is monotonically decreasing in m_{21}^t, implying the existence of m_{dt} such that $\hat{V}_{dt}^t(\mathbf{I^t}) \geq V_{un}^t(\mathbf{I^t})$ if and only if $m_{21}^t \leq m_{dt}$. Since the uniform pricing policy is a special case of the price differentiating policy, we must have $\hat{V}_{dt}^t(\mathbf{I^t}) \geq \hat{V}_{ut}^t(\mathbf{I^t})$. This implies whenever it is optimal to transship under the uniform pricing policy, it must also be optimal to transship under the price differentiating policy, which in turn implies $m_{dt} \geq m_{ut}$. Now, let $m_L = m_{ut}$ and $m_H = m_{dt}$. It is easy to check that they satisfy the properties stated above.

Next, we will show that when the transshipment decision is made optimally, the benefit from price differentiation monotonically decreases in m_{21}^t. Let $\bar{V}_{dt}^t(\mathbf{I^t})$ denote the retailer's expected profit under the optimal transshipment and price differentiation policy in period t. Likewise, let $\bar{V}_{ut}^t(\mathbf{I^t})$ denote the retailer's expected profit under the optimal transshipment and uniform pricing policy in period t. Let $\Delta V_{dt}^t := \bar{V}_{dt}^t(\mathbf{I^t}) - \bar{V}_{ut}^t(\mathbf{I^t})$ denote the benefit from price differentiation. We will consider the three cases of optimal transshipment decisions stated above.

i) Suppose $m_{21}^t \leq m_L$. Then, $\bar{V}_{dt}^t(\mathbf{I^t}) = \hat{V}_{dt}^t(\mathbf{I^t})$ and $\bar{V}_{ut}^t(\mathbf{I^t}) = \hat{V}_{ut}^t(\mathbf{I^t})$. The benefit from price differentiation is given by $\Delta V_{dt}^t = \hat{V}_{dt}^t(\mathbf{I^t}) - \hat{V}_{ut}^t(\mathbf{I^t})$. Notice that whether the retailer uses the uniform pricing or price differentiating policy, he incurs the same transshipment cost of m_{21}^t. Furthermore, since m_{21}^t is a fixed cost, it does not affect the retailer's pricing decisions. Hence, ΔV_{dt}^t is independent of m_{21}^t.

ii) Suppose $m_L < m_{21}^t \leq m_H$. Then, $\bar{V}_{dt}^t(\mathbf{I^t}) = \hat{V}_{dt}^t(\mathbf{I^t})$ and $\bar{V}_{ut}^t(\mathbf{I^t}) = V_{un}^t(\mathbf{I^t})$. The benefit from price differentiation is given by $\Delta V_{dt}^t = \hat{V}_{dt}^t(\mathbf{I^t}) - V_{un}^t(\mathbf{I^t})$. Since $\hat{V}_{dt}^t(\mathbf{I^t})$ is monotonically decreasing in m_{21}^t while $V_{un}^t(\mathbf{I^t})$ is independent of m_{21}^t, we have that ΔV_{dt}^t monotonically decreases in m_{21}^t.

iii) Suppose $m_{21}^t > m_H$. Then, $\bar{V}_{dt}^t(\mathbf{I^t}) = V_{un}^t(\mathbf{I^t})$ and $\bar{V}_{ut}^t(\mathbf{I^t}) = V_{un}^t(\mathbf{I^t})$. The benefit from price differentiation is given by $\Delta V_{dt}^t = V_{un}^t(\mathbf{I^t}) - V_{un}^t(\mathbf{I^t}) = 0$. Hence, ΔV_{dt}^t is independent of m_{21}^t.

We have shown that in each interval of m_{21}^t, the benefit from price differentiation is (weakly) monotonically decreasing in m_{21}^t. To complete the proof, it suffices to show that ΔV_{dt}^t is continuous at $m_{21}^t = m_L$ and $m_{21}^t = m_H$. To see this, notice that at $m_{21}^t = m_L$, $\hat{V}_{ut}^t(\mathbf{I^t}) = V_{un}^t(\mathbf{I^t})$, and at $m_{21}^t = m_H$, $\hat{V}_{dt}^t(\mathbf{I^t}) = V_{un}^t(\mathbf{I^t})$. Then, the continuity follows immediately from the monotonically of $\hat{V}_{ut}^t(\mathbf{I^t})$ and $\hat{V}_{dt}^t(\mathbf{I^t})$ in m_{21}^t.

BIBLIOGRAPHY

BIBLIOGRAPHY

Achabal, D. D., S. McIntyre, and S. A. Smith (1990), Maximizing profits from periodic department store promotions, *J. Retailing*, *66*(4), 383–407.

Armstrong, M. (1996), Multiproduct nonlinear pricing, *Econometrica*, *64*(1), 51–75.

Barnes-Schuster, D., Y. Bassok, and R. Anupindi (2002), Coordination and flexibility in supply contracts with options, *Manufacturing Service Oper. Management*, *4*(3), 171–207.

Bell, D. R., J. Chiang, and V. Padmanabhan (1999), The decomposition of promotional response: An empirical generalization, *Marketing Sci.*, *18*(4), 504–526.

Bitran, G., R. Caldentey, and S. Mondschein (1998), Coordinating clearance markdown sales of seasonal products in retail chains, *Oper. Res.*, *46*(5), 609–624.

Bitran, G. R., and S. V. Mondschein (1997), Periodic pricing of seasonal products in retailing, *Management Sci.*, *43*(1), 64–79.

Bitta, A. J. D., K. B. Monroe, and J. M. McGinnis (1981), Consumer perceptions of comparative price advertisements, *J. Marketing Res.*, *18*(4), 416–427.

Blair, A. (Jul. 1, 2013), Best buy enables online fulfillment from all dcs, pilots ship-from-store, *Retail Info Systems News*.

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

Bolton, P., and M. Dewatripont (2005), *Contract theory*, MIT Press.

Borrell Associates, I. (2010), Proximity-based marketing: Mobile devices untether advertising from media.

Brown, A., and H. Lee (1997), Optimal pay-to-delay capacity reservation with application to the semiconductor industry, *Working paper, Stanford University*.

Brownell, M. (May 31, 2013), Which stores price-match their own websites - and which ones won't, *Daily Finance*.

Cachon, G., and M. Lariviere (2001), Contracting to assure supply: How to share demand forecasts in a supply chain, *Management Sci.*, *47*(5), 629–646.

Campbell, M. C. (1999), Perceptions of price unfairness: Antecedents and consequences, *J. Marketing Res.*, *36*(2), 187–199.

Ceryan, O., O. Sahin, and I. Duenyas (2013), Dynamic pricing of substitutable products in the presence of capacity flexibility, *Manufacturing Service Oper. Management*, *15*(1), 86–101.

Chen, S. S., K. B. Monroe, and Y. Lou (1998), The effects of framing price promotion messages on consumers' perceptions and purchase intentions, *J. Retailing*, *74*(3), 353–372.

Chu, J., P. K. Chintagunta, and N. J. Vilcassim (2007), Assessing the economic value of distribution channels: An application to the personal computer industry, *J. Marketing Res.*, *44*(1), 29–41.

Chuang, J. C., and M. A. Sirbu (1999), Optimal bundling strategy for digital information goods: network delivery of articles and subscriptions, *Inf. Econ. Policy*, *11*, 147–176.

Corbett, C. J., and X. de Groote (2000), A supplier's optimal quantity discount policy under asymmetric information, *Management Sci.*, *46*(3), 444–450.

Courty, P., and L. Hao (2000), Sequential screening, *Rev. Econ. Stud.*, *67*(4), 697–717.

Dada, M., and K. N. Srikanth (1987), Pricing policies for quantity discounts, *Management Sci.*, *33*(10), 1247–1252.

DelVecchio, D. (2005), Deal-prone consumers' response to promotion: The effects of relative and absolute promotion value, *Psychol. Market.*, *22*(5), 373–391.

DelVecchio, D., H. S. Krishnan, and D. C. Smith (2007), Cents or percent? the effects of promotion framing on price expectations and choice, *J. Marketing*, *71*(3), 158–170.

DeNicola, L. (May 13, 2013), Best price-matching policies, *Cheapism.com.*

Dhar, S. K., and S. J. Hoch (1996), Price discrimination using in-store merchandising, *J. Marketing*, *60*(1), 17–30.

Dong, L., P. Kouvelis, and Z. Tian (2009), Dynamic pricing and inventory control of substitute products, *Manufacturing Service Oper. Management*, *11*(2), 317–339.

Donohue, K. (2000), Efficient supply contracts for fashion goods with forecast updating and two production modes, *Management Sci.*, *46*(11), 1397–1411.

Erhun, F., P. Keskinocak, and S. Tayur (2008), Dynamic procurement, quantity discounts, and supply chain efficiency, *Prod. Oper. Manag.*, *17*(5), 543–550.

Federgruen, A., and A. Heching (2002), Multilocation combined pricing and inventory control, *Manufacturing Service Oper. Management*, *4*(4), 275–295.

Ferguson, M. (2003), When to commit in a serial supply chain with forecast updating, *Nav. Res. Log.*, *50*(8), 917–936.

Ferguson, M., G. Decroix, and P. Zipkin (2005), Commitment decisions with partial information updating, *Nav. Res. Log.*, *52*(8), 780–795.

Financial-Post (Sep. 24, 2013), Retailers creating a seamless shopping experience, Financial Post.

Fontinelle, A. (Aug. 17, 2011), 6 sneaky ways coupons make you spend more, *Investopedia*.

Foubert, B., and E. Gijsbrechts (2007), Shopper response to bundle promotions for packaged goods, *J. Marketing Res.*, *44*(4), 647–662.

Fudenberg, D., B. Holmstrom, and P. Milgrom (1990), Short-term contracts and long-term agency relationships, *J. Econ. Theory*, *51*, 1–31.

Gallego, G., and G. van Ryzin (1994), Optimal dynamic pricing of inventories with stochastic demand over finite horizons, *Management Sci.*, *40*(8), 999–1020.

Gendall, P., J. Hoek, T. Pope, and K. Young (2006), Message framing effects on price discounting, *J. of Product & Brand Management*, *15*(7), 458–465.

Gurnani, H., and C. Tang (1999), Optimal ordering decisions with uncertain cost and demand forecast updating, *Management Sci.*, *45*(10), 1456–1462.

Hackleman, E. C., and J. M. Duker (1980), Deal proneness and heavy usage: Merging two market segmentation criteria, *J. Acad. Market. Sci.*, *8*(4), 332–344.

Hammond, J. H., and A. Raman (1994), Sport obermeyer, ltd.

Hanson, W., and R. K. Martin (1990), Optimal bundle pricing, *Management Sci.*, *36*(2), 155–174.

Hardesty, D. M., and W. O. Bearden (2003), Consumer evaluations of different promotion types and price presentations: the moderating role of promotional benefit level, *J. Retailing*, *79*(1), 17–25.

Heath, T. B., S. Chatterjee, and K. R. France (1995), Mental accounting and changes in price: The frame dependence of reference dependence, *J. Consum. Res.*, *22*(1), 90–97.

Herrmann, A., F. Huber, and R. H. Coulter (1997), Product and service bundling decisions and their effects on purchase intention, *Pricing Strategy and Practice*, *5*(3), 99–107.

Hitt, L. M., and P. Chen (2005), Bundling with customer self-selection: A simple approach to bundling low-marginal-cost goods, *Management Sci.*, *51*(10), 1481–1493.

Hu, J., E. Watson, and H. Schneider (2005), Approximate solutions for multi-location inventory systems with transshipments, *Int. J. Prod. Econ.*, *97*(1), 31–43.

Huang, H., S. Sethi, and H. Yan (2005), Purchase contract management with demand forecast updates, *IIE Trans.*, *37*(8), 775–785.

Hui, W., B. Yoo, and K. Y. Tam (2008), The optimal number of versions: Why does goldilocks pricing work for information goods?, *J. Manage. Inform. Syst.*, *24*(3), 167–191.

Iyer, A., and M. Bergen (1997), Quick response in manufacturer-retailer channels, *Management Sci.*, *43*(4), 559–570.

Jeuland, A. P., and C. Narasimhan (1985), Dealing-temporary price cuts-by seller as a buyer discrimination mechanism, *J. Bus.*, *58*(3), 295–308.

Kacen, J. J., J. D. Hess, and W. K. Chiang (2003), Bricks or clicks? consumer attitudes toward traditional stores and online stores, *Working paper*.

Kivetz, R., O. Urminsky, and Y. Zheng (2006), The goal-gradient hypothesis resurrected: Purchase acceleration, illusionary goal progress, and customer retention, *J. Marketing Res.*, *43*(1), 39–58.

Klaft, H. (Sep. 3, 2011), Coupon fever: Using deals to save on groceries a way of life for local woman, *Jackson Citizen Patriot*.

Kukar-Kinney, M., N. M. Ridgway, and K. B. Monroe (2012), The role of price in the behavior and purchase decisions of compulsive buyers, *J. Retailing*, *88*(1), 63–71.

Laffont, J. J., and D. Martimort (2002), *The theory of incentives: The principal-agent model*, Princeton University Press.

Lariviere, M. A. (2002), Inducing forecast revelation through restricted returns, *Working paper, Northwestern University*.

Laroche, M., F. Pons, N. Zgolli, and C. Kim (2001), Consumers use of price promotions: a model and its potential moderators, *J. Retailing Consum. Services*, *8*, 251–260.

Laroche, M., F. Pons, N. Zgolli, M. Cervellon, and C. Kim (2003), A model of consumer response to two retail sales promotion techniques, *J. Bus. Res.*, *56*, 513–522.

Lazear, E. P. (1986), Retail pricing and clearance sales, *Am. Econ. Rev.*, *76*(1), 14–32.

Lee, H. L., and M. J. Rosenblatt (1986), A generalized quantity discount pricing model to increase supplier's profits, *Management Sci.*, *32*(9), 1177–1185.

Lee, L., and D. Ariely (2006), Shopping goals, goal concreteness, and conditional promotions, *J. Consum. Res.*, *33*(1), 60–70.

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

Lichtenstein, D. R., R. G. Netemeyer, and S. Burton (1990), Distinguishing coupon proneness from value consciousness: An acquisition-transaction utility theory perspective, *J. Marketing*, *54*(3), 54–67.

Lichtenstein, D. R., S. Burton, and R. G. Netemeyer (1997), An examination of deal proneness across sales promotion types: A consumer segmentation perspective, *J. Retailing*, *73*(2), 283–297.

Lobel, I., and W. Xiao (2013), Optimal long-term supply contracts with asymmetric demand information, *Working paper, New York University.*

Luce, R. D. (1959), *Individual Choice Behavior: A Theoretical Analysis*, Wiley, New York.

Lynch, E. (Jul. 18, 2013), Why omnichannel fulfillment matters, *Multichannel Merchant.*

McFadden, D. L. (1974), *Conditional logit analysis of qualitative choice behavior*, 105-142 pp., Academic Press, New York.

Miyaoka, J., and W. Hausman (2008), How improved forecasts can degrade decentralized supply chains, *Manufacturing Service Oper. Management*, *10*(3), 547–562.

Monahan, J. P. (1984), A quantity discount pricing model to increase vendor profits, *Management Sci.*, *30*(6), 720–726.

Moon, Y., T. Yao, and T. L. Friesz (2010), Dynamic pricing and inventory policies: A strategic analysis of dual channel supply chain design, *Service Science*, *2*(3), 196–215.

Munson, C. L., and M. J. Rosenblatt (1998), Theories and realities of quantity discounts: An exploratory study, *Prod. Oper. Manag.*, *7*(4), 352–369.

Narayanan, V. G., and A. Raman (2004), Aligning incentives in supply chains, *Harvard Bus. Rev.*, *82*(11), 94–102.

NCH Marketing Services, I. (2011), Coupon facts report.

Neslin, S. A., D. Grewal, R. Leghorn, V. Shankar, M. L. Teerling, J. S. Thomas, and P. C. Verhoef (2006), Challenges and opportunities in multichannel customer management, *J. Serv. Res.*, *9*(2), 95–112.

Oh, S., and Ö. Özer (2012), Mechanism design for capacity planning under dynamic evolutions of asymmetric demand forecasts, *Management Sci.*, *59*(4), 987–1007.

Özer, Ö., and W. Wei (2006), Strategic commitment for optimal capacity decision under asymmetric forecast information, *Management Sci.*, *52*(8), 1238–1257.

Paterson, C., G. Kiesmller, R. Teunter, and K. Glazebrook (2011), Inventory models with lateral transshipments: A review, *Eur. J. Oper. Res.*, *210*(2), 125–136.

Plambeck, E., and S. Zenios (2000), Performance-based incentives in a dynamic principal-agent model, *Manufacturing Service Oper. Management*, *2*(3), 240–263.

Rochet, J. C., and P. Choné (1998), Ironing, sweeping, and multidimensional screening, *Econometrica*, *66*(4), 783–826.

Rudi, N., S. Kapur, and D. F. Pyke (2001), A two-location inventory model with transshipment and local decision making, *Management Sci.*, *47*(12), 1668–1680.

Schindler, R. M. (1989), The excitement of getting a bargain: Some hypotheses concerning the origins and effects of smart-shopper feelings, *Adv. Consum. Res.*, *16*, 447–453.

Schindler, R. M. (1992), A coupon is more than a low price: Evidence from a shopping-simulation study, *Psychol. Market.*, *9*(6), 431–451.

Schindler, R. M. (1998), Consequences of perceiving oneself as responsible for obtaining a discount: Evidence for smart-shopper feelings, *J. Consum. Psychol.*, *7*(4), 371–392.

Schinkel, M. (Jun. 25, 2013), Out of stock online while best buys are on the shelf!, *Wall St. Cheat Sheet*, pp. –.

Schmalensee, R. (1984), Gaussian demand and commodity bundling, *J. Bus.*, *57*(1), S211–S230.

Seifert, R. W., U. W. Thonemann, and W. H. Hausman (2004), Optimal procurement strategies for online spot markets, *Eur. J. Oper. Res.*, *152*(3), 781–799.

Shen, C., and X. Zhang (2012), Retailer's pricing and channel selection under e-commerce environment, *Soft Computing in Information Communication Technology*, *161*, 355–363.

Sherman, L. (Feb. 5, 2009), Eight reasons why we overspend, *Forbes*.

Smith, S. A., and D. D. Achabal (1998), Clearance pricing and inventory policies for retail chains, *Management Sci.*, *44*(3), 285–300.

Sotgiu, F., and F. Ancarani (2004), Exploiting the opportunities of internet and multichannel pricing: an exploratory research, *J. of Product & Brand Management*, *13*(2), 125–136.

Suh, M., and G. Aydin (2011), Dynamic pricing of substitutable products with limited inventories under logit demand, *IIE Trans.*, *43*(5), 323–331.

Taylor, T., and W. Xiao (2009), Incentives for retailer forecasting: Rebates versus returns, *Management Sci.*, *55*(10), 1654–1669.

Taylor, T., and W. Xiao (2010), Does a manufacturer benefit from selling to a better-forecasting retailer?, *Management Sci.*, *56*(9), 1584–1598.

Taylor, T. A. (2006), Sale timing in a supply chain: When to sell to the retailer, *Manufacturing Service Oper. Management*, *8*(1), 23–42.

Thaler, R. (1985), Mental accounting and consumer choice, *Marketing Sci.*, *4*(3), 199–214.

Thaler, R. H. (1999), Mental accounting matters, *J. Behav. Decis. Making*, *12*(3), 183–206.

Tsai, D. (2007), Cash rewards effect on unplanned buying behaviors–watch your wallet carefully!, *Consumer Interests Annual, Nbr. 53*.

Tuttle, B. (Jul. 23, 2010), Bound to buy: The 10 types of consumers who inevitably overspend, *Time Magazine*.

Varian, H. R. (1980), A model of sales, *Am. Econ. Rev.*, *70*(4), 651–659.

Venkatesh, R., and V. Mahajan (1993), A probabilistic approach to pricing a bundle of products or services, *J. Marketing Res.*, *30*(4), 494–508.

Wansink, B., R. J. Kent, and S. J. Hoch (1998), An anchoring and adjustment model of purchase quantity decisions, *J. Marketing Res.*, *35*(1), 71–81.

Weng, Z. K. (1995), Modeling quantity discounts under general price-sensitive demand functions: Optimal policies and relationships, *Eur. J. Oper. Res.*, *86*, 300–314.

Wilcox, J. B., R. D. Howell, P. Kuzdrall, and R. Britney (1987), Price quantity discounts: Some implications for buyers and sellers, *J. Marketing*, *51*(3), 60–70.

Wolk, A., and C. Ebling (2010), Multi-channel price differentiation: An empirical investigation of existence and causes, *Int. J. Res. Mark.*, *27*(2), 142–150.

Wu, S., L. M. Hitt, P. Chen, and G. Anandalingam (2008), Customized bundle pricing for information goods: A nonlinear mixed-integer programming approach, *Management Sci.*, *54*(3), 608–622.

Yan, R. (2008), Pricing strategy for companies with mixed online and traditional retailing distribution markets, *J. of Product & Brand Management*, *17*(1), 48–56.

Zhang, H., and S. Zenios (2008), A dynamic principal-agent model with hidden information: Sequential optimality through truthful state revelation, *Oper. Res.*, *56*(3), 681–696.