Pricing for Multinomial Logit Choice Models with Network Effects

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL OF THE

UNIVERSITY OF MINNESOTA

BY

Chenhao Du

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

Advisors: William L. Cooper and Zizhuo Wang

JULY, 2017

Acknowledgments

I am blessed to have two brilliant professors as my advisers. Professor William L. Cooper, with his rigorous scholarship and insightful vision, has consistently supported me and patiently guided me throughout my graduate studies. I wish I could carry his good characteristics into my future path. Professor Zizhuo Wang, with his full knowledge and genius analytical skills, has not only guided me in the academic directions but also assisted me in the technical details I encountered in writing. His entrepreneurial initiatives further inspire me with spirit of adventure and enthusiasm in regard to life. Whether I receive excitement or frustration during the research experience, both of them have always made efforts to guide and mentor me, and I would like to express my deepest appreciation to them. This dissertation would not have been possible without their persistent caring, encouragement, and support.

I would also like to thank my committee members, Professor Saif Benjaafar, Professor Tony Cui, and Professor Guangwen Kong, for their valuable comments and insightful suggestions on my dissertation. I also feel grateful for their generous advice on my research and career development.

It is my luck to meet so many great friends here at Minnesota in the past six years. My life has been enriched by their surroundings and accompanies. Specifically, I want to thank: Yibin Chen, Lu Yuan, Rui Chen, Yu Wang, Yan Liu, Xiao Chen and Xiang Li, for their help and support.

My final thanks go to my parents for the care and sacrifices they have made for me since I was born. It is my fortune to have parents like you.

Dedication

To my parents for their selfless love and care.

Abstract

We are experiencing a time in which advances in information technology change everything around us. Almost every aspect of our work and life, from daily travel to grocery shopping, is now or will soon be connected to the internet. This drastic change has led to the emergence of many new products that display network effects, meaning that individual consumers value a product more when more other consumers use the same product. At the same time, development of data storage/processing tools offers a great chance for many firms to collect large amounts of data and understand their customers' purchasing behavior. For these firms, a significant challenge — as well as opportunity — is to develop models of customer choice behavior that incorporate network effects and to use such models to make better pricing decisions. This dissertation addresses the development and analysis of some such models.

We consider a seller's problem of determining revenue-maximizing prices for an assortment of products that exhibit network effects. Customers make purchase decisions according to a multinomial logit (MNL) choice model, modified — to incorporate network effects — so that the utility each individual customer gains from purchasing a particular product depends on the market's total consumption of that product. In the setting of homogeneous products, we show that if the network effect is comparatively weak, then the optimal pricing decision of the seller is to set identical prices for all products. However, if the network effect is strong, then the optimal pricing decision is to set the price of one product low and

to set the prices of all other products to a single high value. This pricing scheme boosts the sales of the single low-price product in comparison to the sales of all other products. The analysis is also extended to settings with heterogeneous products, and we show that optimal solutions have a structure similar to that found in the homogeneous case: either maintain a semblance of balance among all products, or boost the sales of just one product. Based on this structure, we propose an effective computational algorithm for such general heterogeneous settings.

Subsequently, we study the preceding pricing problem from a robust optimization perspective. Unlike the classical MNL model where products' prices and sales have a one-to-one correspondence, in the MNL model with network effects a fixed set of prices may not uniquely determine sales. This occurs because, for given prices, sales arise as the solution to an equilibrium condition. In some cases, there may be multiple sales levels that satisfy the equilibrium condition. Among those sales equilibria corresponding to a given set of prices, we call the one with the highest revenue the "optimistic" equilibrium, and the one with the lowest revenue the "pessimistic" equilibrium. In our initial study mentioned in the previous paragraph, we implicitly took an "optimistic" approach. We next take the pessimistic attitude and study the revenue-maximizing problem in the pessimistic setting. In the case that there is only one product to sell, the problem has the same pessimistic optimal price as the optimistic one, when the network effect is relatively weak. However, when the network effect is strong, the optimal policy requires the seller to become more conservative and the pessimistic optimal price is lower than the optimistic one. In the case that there are two products, the structure of equations at equilibrium becomes more complicated and we are not able to derive an analytical solution. To numerically solve this problem, two directions for finding the pessimistic optimal solution are proposed: divide and search, and linear relaxation. In addition, our numerical studies show that when the network effect is strong, the revenue from offering exactly one product is almost as good as that from selling an assortment of multiple products. This suggests that a company selling multiple products with strong network effects may be wise to simply offer just one product.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Revenue Management (RM) practices in the airline industry originated largely in response to the Airline Deregulation Act in 1978. With this Act, the U.S. Civil Aviation Board (CAB) relaxed its control of airline prices and schedules, and therefore gave airline carriers unprecedented freedom to make their own pricing decisions. Since then, all carriers, ranging from long-established legacy carriers to new low-cost startups, have been actively developing new selling and marketing strategies, in hope to entice more customers and bring in higher revenues. At the same time, advances in information technology have enabled carriers to track and manage demands and sales, it further encouraging them to build new revenue management systems.

The central idea of revenue management is to sell the right product for the right price to the right customer at the right time, with the goal to maximize the total revenue. In the airline industry, the primary controllable decisions are prices, products, and inventory. On a single flight, it seems that any two seats in the same cabin have almost the same value for a customer. However, by imposing various purchase restrictions, cancellation policies and length of stay requirements, a carrier is able to essentially create ticket products with different values from the same resource. These differentiated products are labeled as different fare classes in the airline industry. Furthermore, even if product characteristics and prices are determined on a flight, the carrier has the flexibility to control the availability of each product at any time by adjusting their inventories.

In practice, there are two main aspects of revenue management: forecasting and optimization. In the forecasting stage, the work involves collecting data on customers' behavior (e.g., purchases) and using that data to predict future demand, or more generally, future responses by customers to carriers' decisions. In the optimization stage, RM practitioners need to determine the pricing decisions and inventory levels to maximize the total revenue. An excellent implementation of RM requires integration of both aspects. In addition to the airline industry, many other industries are taking the power of RM and applying it to their businesses, including hotels, retailing, car rentals, cruise lines, etc.

The success of implementing RM in a business depends on understanding consumer behavior and building a model that reflects reality. One of the earliest models in capacity control for a single-leg problem was proposed by Littlewood (1972). This two-fare-class problem can be elegantly solved by an explicit expression of the protection levels for the high-fare class. The key idea is, when remaining inventory drops to the optimal protection level, revenue from selling one seat to a certain low-fare request is equivalent to that of reserving that seat for a potential high-fare request. Although the multi-fare-class version of this problem has no explicit optimal solution, Belobaba (1987a, b) used the same idea and designed two simple but effective heuristics (EMSR-a and EMSR-b) that are still widely used in business. Inspired by this classical model, much research has focused on variations and extensions of this model. Some examples include Brumelle and McGill (1993), Wollmer (1992), Robinson (1995) and van Ryzin and McGill (2000).

"All models are wrong but some are useful." This famous quote by statistician George E. P. Box has been mentioned in several OR/MS works, for example Besbes and Zeevi (2015). Inevitably, it also applies to the Littlewood model as well as most of its extensions. The classical model makes an oversimplified assumption that consumer demands for each fare-class are independent of each other and also independent of the controls being applied on them. In addition, it assumes that low-fare customers always arrive first and high-fare customers arrive later. This assumption may not hold in reality, especially today when product price and information are largely transparent to customers over the internet selling channel. Buy-up (switching to a high-fare when a low-fare is unavailable) and buy-down (substituting a low-fare for a high-fare when discounts are open) effects have attracted research attention and there is much work focusing on this direction. Please see Brumelle et al. (1990), Belobaba and Hopperstad (1999), Cooper et al. (2006), and Cooper and Li (2012).

To address the oversimplification in early models, Talluri and van Ryzin (1994) applied the multinomial logit (MNL) choice model to the realm of revenue management. The MNL model was proposed and formulated by McFadden (1974) and has been actively studied and implemented in many areas including economics, operations research, marketing, etc. It derives from an underlying random utility model, in which each individual customer has a random utility for each product. In the model, consumer choice behavior is explicitly represented with a purchasing probability for each product, which depends on all the available products. Talluri and van Ryzin took up this model to study the single-leg multi-fare problem, and showed that it can directly address the issue of independent demand over substitutable products. Based on the MNL model, they further found a surprisingly simple structure for the optimal solution to the assortment optimization problem.

Despite the MNL model's many advantages in modeling multi-product consumer choice behavior, there still exist some restrictions for its own. One of these restrictions is the independence of irrelevant alternatives (IIA) property. According to the model, the ratio of any two alternatives contained in the available set remains unchanged whenever a new alternative is added. A classical example to demonstrate IIA is the well-known red-bus/blue-bus paradox. Because of this, IIA is often considered one of the most unfavorable properties for the MNL model.

There have been many studies of multi-product pricing problems in which customer choices are governed by the MNL model or variations thereof. One of the central conclusions that emerges from this work is that if consumer price sensitivities are identical across products, then an optimal pricing strategy involves a constant markup for all products; that is, it is optimal to set prices so that the difference between a product's price and its unit cost is the same for all products. This result was first obtained by Anderson and de Palma (1992). In the interim, a number of alternative proofs and extensions have appeared. For a summary of this work, we refer to the literature review provided in the paper by Gallego and Wang (2014). Early research on computational approaches for MNL pricing problems includes that of Hanson and Martin (1996), who show that in such problems the objective function (total revenue) is not a concave function of the vector of prices. Subsequent work has established, however, that if the problem is re-formulated with the vector of sales quantities as the decision variable, then the objective function is, in fact, concave and thus the problem can be solved efficiently; see Xue and Song (2007), Dong et al. (2009), and Li and Huh (2011). One notable computational approach for MNL pricing problems involves a reduction of the multivariate optimization problem to a suitable single-variable optimization problem that can be solved with a one-dimensional search.

In this dissertation, we study a multi-product pricing optimization problem based on the MNL model. However, we are particularly interested in a group of products with network effects. A product exhibits network effects if individual consumers value it more when more other consumers purchase it. The idea of network effects originated from telephone service, see Katz and Shapiro (1986). The more people who join a telephone network, the more people each individual user in the network could have contact with, therefore the more valuable this telephone network is to each user. Online social platforms, such as Facebook and Twitter, may be viewed as modern versions of telephone networks. As we have been increasingly relying on the internet for everything in our work and life, the world has seen the development and success of many other products or services with network effects, including computer software, online video games, and app-based ride sharing services.

There are already many works focusing on products with network effects, and most of them can be divided into two categories. One category addresses products with global network effects, whereby a customer's utility for a product depends on the total consumption of the product in question. The other addresses products with local network effects, whereby a consumer gains utility if his/her "neighbors" purchase the same product. A common starting point in studies involving local network effects is a graph of social connections. We refer to Candogan et al. (2012), Bloch and Quérou (2012), and references therein for studies of local network effects. In this dissertation, we consider only global network effects. The study of global network effects has a long history in the economics literature. For reviews, we refer to Farrell and Saloner (1985), Katz and Shapiro (1985), and Economides (1996a). The MNL model with network effects that serves as the input to our pricing optimization problem has been considered in previous studies by Anderson et al. (1992), Brock and Durlauf (2002), and Starkweather (2003). Anderson et al. (1992, Section 7.8) considers an oligopoly in which each firm sets the price of its own single product, and market shares are determined by the MNL model with network effects. The questions, conclusions, and methodology in this dissertation are quite distinct from those of Anderson et al.

This dissertation is organized as follows. Chapter 2 studies a seller's problem of determining revenue-maximizing prices for an assortment of products that exhibit network effects. Customers make purchase decisions according to a multinomial logit choice model, modified — to incorporate network effects — so that the utility each individual customer gains from purchasing a particular product depends on the market's total consumption of that product. The content of this chapter appeared in Du et al. (2016). Chapter 3 continues to address the same problem from a different perspective. The introduction of network effects into the MNL model causes a result that a fixed pricing decision may have multiple sales quantities in equilibrium. Put differently, the sales may not be uniquely determined by the price, even in a deterministic model. This poses an important question to the seller — which equilibrium to expect in face of multiple equilibria. To address this issue and to guard the seller against poor outcomes, we study the optimal pricing problem from a "pessimistic" perspective, wherein it is assumed that for a given pricing decision, the worst possible sales level (from the seller's perspective) arises. Chapter 4 summarizes this dissertation and points to several future research directions.

Chapter 2

Optimal Pricing for a Multinomial Logit Choice Model

2.1 Introduction

A product is said to exhibit network effects if individual consumers value it more when more other consumers purchase it. For example, video game industry has never been such booming like today and the most popular games fall into the category of multi-player online games. In these games, each individual player must interact with other players to finish a task, and therefore they would obtain more enjoyment (and thus greater value) if there are more other players to play with. In addition to video games, rise of app-based technology has led to the emergence of new products or services that display network effects (e.g., ride sharing services and group buying deals). At the same time, expansion of social media has served to strengthen the network effects for traditional products such as movies, television programs, and books by allowing people to easily participate in communities built around those products. Furthermore, development of data storage/processing tools offers a great chance for many firms to collect large amounts of data and understand their customers' purchasing behavior. For these firms, a significant challenge — as well as opportunity — is to develop models of customer choice behavior that incorporate network effects and to use such models to make better pricing decisions.

We take up this issue, and address a pricing problem faced by a seller that offers a given assortment of products that exhibit network effects. In particular, we consider a setting in which demands for various products are determined by a variant of the multinomial logit (MNL) model in which the expected utility a typical consumer derives from purchasing an individual product depends on its intrinsic quality, its price, and also a network effect term that is a linear function of the market's total consumption of that product. Given a product, the slope of this network effect term represents the strength of that product's network effect. We focus on the seller's optimal (i.e., revenue maximizing) pricing and sales decisions. More specifically, we seek to answer the following three questions in this chapter: (1) What is the structure of the optimal solution in this setting with network effects? (2) Does the presence of network effects yield solutions that are fundamentally different than those that arise in problems without network effects? (3) How do optimal decisions depend upon the strength of network effects?

To answer these questions, we first consider a homogeneous case in which model primitives, including the slope parameters that determine the strength of the network effects, are the same across all products. We establish that the optimal solution takes one of two different forms, depending upon the strength of the network effect. When the network effect is relatively weak, we show that it is optimal for the seller to price all the products identically. This result is consistent with the classical MNL pricing problem without network effects for which it is known that it is optimal to price homogeneous products identically (we will review the literature below). We note, however, that although optimal prices are identical when the network effect is weak, those prices are different from the price that is optimal when there are no network effects. When the network effect is strong enough, we show that the optimal pricing policy is such that exactly one product is priced low and that all other products are priced at a single higher price. As a result, the sales of the one low-price product will be higher than those of each of the other products. This strategy that "boosts" the sales quantity of one single product differs from the equal-pricing strategy that is optimal in classical MNL models and thus is a unique feature that arises in the presence of network effects. We also show that in such scenarios with a strong network effect, if the network effect becomes stronger, then the optimal prices will be such that the sales of the single low-price product will increase, while the sales of each of the other products will decrease. As the strength of the network effect increases, the solution becomes progressively more dissimilar to the solution of the classic MNL model. In the limit as the network effect becomes "very strong" (i.e., as the aforementioned slope becomes very large), sales of the low-price product grow to capture the entire market while the sales of all the high-priced products decrease to zero. This stands in particularly stark contrast to the equal prices and sales quantities that emerge in classical MNL pricing problems without network effects.

Next we study a general problem in which parameters are heterogeneous across products. In such cases we find that the optimal prices will generally all be distinct. Nevertheless, the optimal solution retains some of the structure seen in the homogeneous case. In particular, the sales of at most one product will be boosted by pricing it low, while all other products will be priced high. A precise meaning of "low" and "high" will be provided later. The main idea is that there are many candidate price vectors that may satisfy the first-order necessary optimality conditions for maximizing revenues. If we restrict attention to only these potential optimal price vectors (which is sufficient for finding an optimal solution) then at most a single product will be priced at its lowest potential value while all others will be priced at their highest potential values. We exploit this structure to obtain an computational algorithm that quickly solves the multi-product pricing problem — for an arbitrary number of products — with a simple two-dimensional search.

To better understand each parameter's impact on the pricing decisions, we examine settings in which only network effect parameters or price sensitivities differ across products. Such settings have a limited degree of heterogeneity. This allows us to obtain stronger results than in the fully heterogeneous case. For instance, we are able to establish that the products' prices will take the reverse (respectively, same) order as the products' network effect (resp., price sensitivity) parameters. Moreover, the seller will boost only the sales of the product with highest network effect (resp., lowest price sensitivity) parameter or else boost none at all. We also obtain further simplifications to our computational algorithm.

Finally, we discuss settings in which there are inter-product network effects (i.e., the sales of one product may affect the utility a customer gains from purchasing another product), settings in which network effects enter through more-general functional forms in individual customers' utility functions, and settings in which the no-purchase option experiences network effects. We show that most of our results hold in these more general cases. Overall, our results present a clear picture of the optimal pricing strategy for a wide class of multi-product pricing problems with network effects.

From a technical point of view, there are several novel aspects of the analysis in this chapter. First, unlike in the classical MNL model, demand cannot be written as an explicit function of the prices. Instead, for any vector of prices, sales quantities arise as a fixed point of a mapping that comes from inclusion of the network effects. Second, even after transforming the problem so that demand is the primary decision variable, the objective function in the revenue maximization problem is the sum of a convex function and a concave function, which in general is difficult to analyze. This too differs from the classical MNL case, for which it is well known that revenue is a concave function of demand. Nevertheless, by exploiting the special structure of our problem, we are able to show that the problem with network effects admits a remarkably simple solution structure. That structure allows us to solve the problem quickly, and also reveals key tradeoffs. Finally, in order to explore the comparative statics of optimal solutions with respect to the network effect parameters, we introduce a novel transformation of variables. After transformation, the objective function satisfies a supermodularity property that does not exist in the forms of the optimization problem that have either prices or sales as the decision variables. The key of the transformation is the selection of an orthogonal matrix that maps the sum of the original variables to a single new variable, and that maps the region described by certain constraints on the unit simplex in the original problem to a sublattice. Since simplex constraints and objective functions that depend on the sum of variables are typical in allocation problems, we believe such a transformation could be useful in other contexts.

In the introductory chapter, we have already reviewed much related literature. These works can be categorized into two lines: multi-product problems studies and network products studies. In the line of multi-product problem studies, in addition to pricing optimization problems, there are other recent work that addresses pricing, assortment planning, or availability problems for the MNL model or its variants includes, e.g., van Ryzin and Mahajan (1999), Talluri and van Ryzin (2004), Aydin and Porteus (2008), Suh and Aydin (2011), Wang (2012), and Davis et al. (2013). To draw an important distinction between this chapter and this line of work on operational decision making with MNL models and their relatives, we note that none of the papers mentioned above consider network effects. Among the other line of dealing with products with network effects, Starkweather (2003) also studies pricing and product compatibility problems based upon an MNL model with network effects. However, Starkweather does not find optimal prices and sales levels for a revenue maximizing seller or study how the strength of network effects influences those quantities. Thus this chapter is quite different from that of Starkweather. In a recent paper, Wang and Wang (2016) consider assortment planning problems in which sales are governed by the MNL model with network effects. However, they do not consider pricing decisions, which are the topic of this chapter.

The rest of this chapter is organized as follows. Section 2.2 describes the MNL model with network effects. Section 2.3 focuses on the homogeneous-products case and contains two theorems that describe properties of optimal solutions in that case. Section 2.4 focuses on settings with heterogeneous products. Section 2.6 describes results of numerical studies. Section 2.5 discusses three extensions to the MNL model with network effects. Section 2.7 concludes the chapter. Proofs and extensions are contained in appendices.

2.2 The General Model

We consider a variant of the MNL choice model that incorporates network effects in customer utilities. Suppose a single seller has a line of n products indexed by $i \in \mathcal{N} = \{1, \ldots, n\}$ to sell to a market of total size M. Each individual customer in the market buys at most one of the n products. Such a customer may also decide not to purchase, in which case we view this as selecting the "no purchase" product, which is indexed by 0. The market is comprised of "infinitesimal" customers, so the probability that an individual customer purchases product i is also the overall fraction of customers that purchase product i. The utility a customer obtains from purchasing product i is

$$u_i = v_i + \epsilon_i,$$

where v_i is the expected utility from consuming product i and ϵ_i is a random variable that represents customer-specific idiosyncracies. As in the standard MNL model, we assume $\epsilon_0, \epsilon_1, \ldots, \epsilon_n$ are i.i.d. Gumbel random variables. In our model, v_i is determined by the quality of product i, its price, and a network effect term that depends upon the market's overall consumption of that product. More precisely, for $i \in \mathcal{N}$ we have

$$v_i = y_i - \gamma_i p_i + \alpha_i x_i, \tag{2.1}$$

where y_i is the intrinsic utility of product i, $\gamma_i > 0$ is the price sensitivity parameter, p_i is the price, α_i is the network effect sensitivity parameter, and x_i is the market's overall consumption of product i. The parameter α_i represents the strength of network effects for product i. A larger value of α_i makes the utility of product i for an individual more sensitive to others' consumption of that product. In the basic model, we assume the network effect on v_i depends only on the market's consumption of product i itself. In Section 2.5.1, we will extend the model to consider settings in which the network effect also depends on the market's consumption of other products. Throughout, we shall assume that $\alpha_i \geq 0$. This means that each customer gains greater utility from product i if more other customers purchase product i. In Section 2.5.2 we allow a more general form of network effects by replacing $\alpha_i x_i$ in (2.1) by $f_i(x_i)$ where $f_i(\cdot)$ is a function that satisfies some conditions. These conditions are satisfied by, for instance, $f(x) = \log(x+1)$ and $f(x) = x^2$.

As is common in the literature, without loss of generality, we normalize v_0 to zero. (We also assume that there is no network effect for the no-purchase option. We relax this assumption in Appendix 2.5.3.) We assume that $y_i \geq 0$, which indicates that when offered any product for free, a customer is expected to obtain higher utility from accepting it than not. By standard results for the MNL model, if we are given consumption levels $x_1, \ldots, x_n$ and prices $p_1, \ldots, p_n$, then the probability that a customer purchases product $i \in \mathcal{N}$ is

$$\begin{aligned} q_i = P(u_i = \max_{j\in\{0,1,\dots,n\}} u_j) &= \frac{\exp(v_i)}{1+\sum_{j=1}^n \exp(v_j)} \\ &= \frac{\exp(y_i - \gamma_i p_i + \alpha_i x_i)}{1+\sum_{j=1}^n \exp(y_j - \gamma_j p_j + \alpha_j x_j)}. \end{aligned} \tag{2.2}$$

Because of the infinitesimal customer assumption mentioned above, q_i is also the fraction of the M customers that purchase product i. Thus we have

$$x_i = Mq_i. \tag{2.3}$$

Without loss of generality, we can further normalize the total market size to $M = 1$. (To see this, we can redefine $\tilde{\alpha}_i = M\alpha_i$ and the problem will be equivalent.) With this normalization in place, we will refer to q_i as the sales quantity (or simply sales) of product i. Conditions (2.2)–(2.3) now reduce to

$$q_i = F_i(\boldsymbol{q}) \text{ for all } i \in \mathcal{N} \quad \text{where} \quad F_i(\boldsymbol{q}) = \frac{\exp(y_i - \gamma_i p_i + \alpha_i q_i)}{1+\sum_{j=1}^n \exp(y_j - \gamma_j p_j + \alpha_j q_j)}. \tag{2.4}$$

Note that sales quantities affect choice probabilities, which themselves affect sales quantities. Hence, the preceding expression (2.4) may be viewed as an equilibrium condition. In equilibrium, a vector of sales quantities $\boldsymbol{q} = (q_1, \dots, q_n)$ must be such that for each product $i \in \mathcal{N}$, the sales of that product q_i equals the probability that a customer will purchase that product given the sales $\boldsymbol{q}$. We can obtain a slightly different justification of (2.4) by interpreting $x_1, \dots, x_n$ on the right side of (2.2) as customers' perceptions of sales levels, in which case (2.4) can be viewed as a rational expectations equilibrium of sorts in which customers' perceptions are consistent with reality.

The seller wishes to select prices that maximize its total revenue. (Without loss of generality we assume the cost to the seller of the products is zero.) Given any price vector $\boldsymbol{p} = (p_1, \ldots, p_n)$, the function $F(\boldsymbol{q}) = (F_1(\boldsymbol{q}), \ldots, F_n(\boldsymbol{q}))$ is a continuous function of $\boldsymbol{q} = (q_1, \ldots, q_n)$ from $[0,1]^n$ to $[0,1]^n$. By the Brouwer fixed point theorem, there exists at least one solution to (2.4). In general, given prices $\boldsymbol{p}$, there could be more than one $\boldsymbol{q}$ that satisfies (2.4). Nevertheless, given any $\boldsymbol{q}$ satisfying $q_1, \ldots, q_n > 0$ and $\sum_{i \in \mathcal{N}} q_i < 1$, there is a unique $\boldsymbol{p} = (p_1(\boldsymbol{q}), \ldots, p_n(\boldsymbol{q}))$ defined by

$$p_i(\boldsymbol{q}) = \frac{1}{\gamma_i}\left(\alpha_i q_i - \log q_i + \log\left(1 - \sum_{j=1}^{n} q_j\right) + y_i\right) \tag{2.5}$$

such that (2.4) holds for $\boldsymbol{q}$. For some $\boldsymbol{q}$, if the seller charges prices determined by (2.5), there may be solutions to (2.4) other than $\boldsymbol{q}$. We will discuss the issue of potential existence of multiple equilibria in Section 2.6.1 and Appendix 4.2. For now, we assume that the seller has the capability to choose the sales $\boldsymbol{q}$ by implementing prices $(p_1(\boldsymbol{q}), \ldots, p_n(\boldsymbol{q}))$. This allows us to use $\boldsymbol{q} = (q_1, \ldots, q_n)$ as the decision variables.

Now, with $\boldsymbol{q} = (q_1, \ldots, q_n)$ as the primary variables, the seller's total revenue is

$$\begin{aligned} \pi(\boldsymbol{q}) &= \sum_{j=1}^{n} q_j p_j(\boldsymbol{q}) \\ &= \sum_{j=1}^{n} \frac{\alpha_j}{\gamma_j} q_j^2 + \sum_{j=1}^{n} \frac{q_j}{\gamma_j} \log\left(1 - \sum_{j=1}^{n} q_j\right) + \sum_{j=1}^{n} \frac{q_j}{\gamma_j}\left(y_j - \log q_j\right). \end{aligned} \tag{2.6}$$

The seller's multi-product pricing problem can now be formulated as the following optimization problem:

$$\begin{aligned} \max \quad & \pi(q_1, \ldots, q_n) \\ \text{s.t.} \quad & \sum_{j=1}^{n} q_j \leq 1 \\ & q_i \geq 0,\ i = 1, \ldots, n. \end{aligned} \tag{P0}$$

In the following sections, we analyze the preceding maximization problem, and study optimal choices of both sales quantities $\boldsymbol{q}$ and prices $\boldsymbol{p}$. We will pay particular attention to how network effects influence optimal decisions.

2.3 The Homogeneous Case

In this section, we consider a special case of the general model, in which all n products have the same intrinsic utilities, price sensitivities, and network sensitivities. That is $y_i = y, \gamma_i = \gamma$, and $\alpha_i = \alpha$ for all $i \in \mathcal{N}$. In the classical multi-product price optimization problem (when $\alpha = 0$), it is well known that the optimal decision is to set all prices to be equal (and the sales of different products are equal); see, e.g., Gallego and Wang (2014) and references therein. As we will see shortly, this may not be the case in the presence of network effects, even in this homogeneous setting.

In the homogeneous setting, there is no loss of generality to assume $\gamma = 1$, and therefore we take $\gamma = 1$ in the remainder of this section. With this assumption, given $\boldsymbol{q} = (q_1, \ldots, q_n)$, the expression (2.5) for prices becomes

$$p_i(\boldsymbol{q}) = \alpha q_i - \log q_i + \log\left(1 - \sum_{j=1}^{n} q_j\right) + y, \tag{2.7}$$

and the total revenue (2.6) becomes

$$\pi(\boldsymbol{q}) = \alpha \sum_{j=1}^{n} q_j^2 + \sum_{j=1}^{n} q_j \left(y + \log\left(1 - \sum_{j=1}^{n} q_j\right)\right) - \sum_{j=1}^{n} q_j \log q_j. \tag{2.8}$$

Note that the objective function (2.8) is symmetric in $\boldsymbol{q}$. Thus we can assume

$q_1 \geq q_2 \geq \cdots \geq q_n$ in (P0) without loss of optimality. Then problem (P0) becomes

$$\begin{aligned} \max \quad & \pi(q_1, \ldots, q_n) \\ \text{s.t.} \quad & \sum_{j=1}^{n} q_j \leq 1 \\ & q_i \geq q_{i+1},\ i = 1, \ldots, n-1 \\ & q_n \geq 0. \end{aligned} \tag{P1}$$

In the following, we will study the structure of the optimal solution to problem (P1). In addition, we will identify how optimal sales quantities and optimal prices respond to different values of the network sensitivity parameter α.

Before we state our main result, we first specify a means of selecting a particular optimal solution to problem (P1) in case there are multiple optimal solutions. The situation of having multiple optimal solutions is not typical, therefore one can view the discussion below as mainly serving technical purposes. Nevertheless, the way we select a specific optimal solution is closely connected to our analysis, in which we perform a transformation of variables and analyze the problem using the transformed variables. Let $\mathbf{e}$ be an n-vector of ones. Define Q^* to be the set containing all optimal solutions to (P1) and define $\boldsymbol{q}^* = (q_1^*, \ldots, q_n^*)$ as

$$\boldsymbol{q}^* = \{\boldsymbol{q} | \boldsymbol{q} \in Q^* \text{ and } \mathbf{e}^T A\boldsymbol{q} \geq \mathbf{e}^T A\boldsymbol{q}' \text{ for all } \boldsymbol{q}' \in Q^*\}, \tag{2.9}$$

where $A = (A_{ij})_{n\times n}$ is an orthogonal matrix defined by

$$A_{ij} = \begin{cases} \frac{1}{\sqrt{n}} & i = 1 \\ \frac{1}{\sqrt{(i-1)i}} & i \geq 2, i > j \\ -\frac{i-1}{\sqrt{(i-1)i}} & i \geq 2, i = j \\ 0 & \text{otherwise.} \end{cases} \tag{2.10}$$

Later we will show that $\boldsymbol{q}^*$ defined above is unique. That is, the set on the right side of (2.9) contains only a single element. A key step in our analysis is to make a change of variables $\boldsymbol{s} = A\boldsymbol{q}$ in (P1). We will establish that after the change of variables, the objective function will be supermodular in $(\boldsymbol{s}, \alpha)$. Such supermodularity is not present when the variables are sales or prices. As will be verified later, such a transformation also has the following key properties: (i) s_1 corresponds to the scaled sum of the q_i, (ii) each s_i, $i \geq 2$ corresponds to a weighted sum of differences between each q_j $(j \leq i)$ and q_i; and (iii) an ordered simplex in the $\boldsymbol{q}$-space (the feasible region of (P1)) is still a sublattice in the $\boldsymbol{s}$-space. These properties of the transformation allow us to appeal to monotonicity results on the maximization of supermodular functions on a sublattice, which yield the comparative statics described in our main results below. Full details appear in our subsequent analysis.

To get some more intuition regarding the above definition, it is instructive to consider a problem with $n = 2$, in which case we have

$$A = \begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ 1/\sqrt{2} & -1/\sqrt{2} \end{bmatrix} \text{ and } \boldsymbol{s} = \begin{bmatrix} s_1 \\ s_2 \end{bmatrix} = A\boldsymbol{q} = \frac{1}{\sqrt{2}} \begin{bmatrix} q_1 + q_2 \\ q_1 - q_2 \end{bmatrix}.$$

Here, the variables become scaled versions of the total sales $(q_1 + q_2)$ and the difference of the sales quantities $(q_1 - q_2)$.

The following result describes the optimal solution to (P1) and shows how that solution depends upon the strength α of the network effects. The proof can be found in Appendix 4.2.

Theorem 2.3.1. *There exists $\hat{\alpha}$ (which depends upon y and n) such that*

(a) *if $\alpha \leq \hat{\alpha}$, then $q_1^* = q_2^* = \cdots = q_n^*$ and q_i^* increases in α for all $i \in \mathcal{N}$;*

(b) *if $\alpha > \hat{\alpha}$, then $q_1^* > q_2^* = \cdots = q_n^*$ and q_1^* increases in α, q_i^* decreases in α for $i \geq 2$, and $\sum_{i \in \mathcal{N}} q_i^*$ increases in α. Moreover, $\lim_{\alpha \to \infty} q_1^* = 1$ and $\lim_{\alpha \to \infty} q_i^* = 0$ for $i = 2, \ldots, n$.*

In addition, let α^R be the unique solution to $R(\alpha) = y + \log(2\alpha - n) - \frac{n}{2\alpha - n} = 0$. Then $1/2 < \hat{\alpha} \leq \alpha^R$. Furthermore, if $n = 2$, then $\hat{\alpha} = \alpha^R$.

For the intuition behind Theorem 2.3.1, consider the objective function (2.8), and suppose we fix the sum $\sum_{j=1}^n q_j$ to a constant. With this added constraint, the problem is equivalent to maximizing $\alpha \sum_{j=1}^n q_j^2 - \sum_{j=1}^n q_j \log q_j$. Note that $\alpha \sum_{j=1}^n q_j^2$ is convex in $\boldsymbol{q}$ while $-\sum_{j=1}^n q_j \log q_j$ is concave in $\boldsymbol{q}$. When α is small, the concave term dominates and the optimal solution is symmetric. This corresponds to part (a) of the theorem. On the other hand, when α is large enough, the convex term dominates the concave term. Because the maximal point of a symmetric convex function must be on the boundary, the optimal solution in this case is no longer symmetric. This corresponds to part (b) of the theorem. From the preceding discussion, it is apparent that the optimal solution must be symmetric if $\alpha < 0$.

In addition to providing structural insights, Theorem 2.3.1 greatly simplifies the process of calculating an optimal solution to (P1). The objective function in (P1) is in general neither concave nor convex, and thus finding a global maximum may not be easy, especially when n is large. However, Theorem 2.3.1 allows us to narrow the search dimension from n to 2 (and to 1 for small enough values of α), thereby obtaining a problem that is easily solvable by a brute force search. For any α, we can also further simplify computation of the optimal solution by using the approach described later in Section 2.4.1 that allows us to solve the problem using a one-dimensional search.

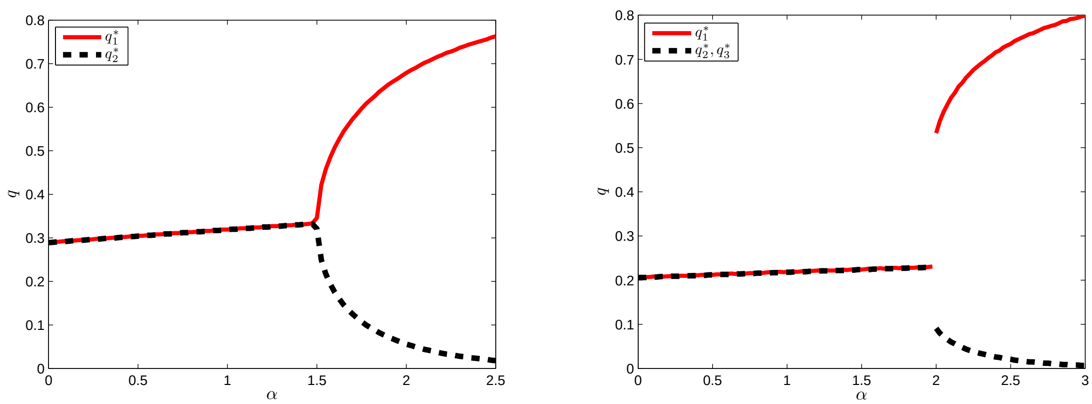

(a) $n = 2$ (b) $n = 3$

Figure 2.1: Optimal $\boldsymbol{q}^*$ as a function of α

Next, we will present two examples to illustrate the above results. In both examples we set $y = 2$. In Figure 2.1(a) we take $n = 2$, and in Figure 2.1(b) we take $n = 3$. Both figures show how the optimal $\boldsymbol{q}^*$ changes with α. As one can see, in both cases, when α is small, all the entries in $\boldsymbol{q}^*$ are equal and increase monotonically in α. When α passes a certain threshold, q_1^* becomes the single largest entry of $\boldsymbol{q}^*$ and all the other entries remain identical and smaller than q_1^*. In addition, the largest entry increases in α while all the other entries decrease in α. Theorem 2.3.1 describes this pattern. In Figure 2.1(a), the threshold is $\alpha = 1.5$, which is precisely the solution to $R(\alpha) = 0$ at $y = 2, n = 2$. In Figure 2.1(b), the threshold is around $\alpha = 1.99$ and the solution to $R(\alpha) = 0$ is $\alpha^R = 2.16$. Thus the thresholds are consistent with the result in Theorem 2.3.1.

Observe in Figure 2.1(b) that $\boldsymbol{q}^*$ is not continuous in α and that there is a jump at the threshold $\hat{\alpha}$. This discontinuity arises from the inherent non-concavity of the problem. When α is close to $\hat{\alpha}$, there are separate local optima of the form $q_1 = \cdots = q_n$ and of the form $q_1 > q_2 = \cdots = q_n$. When α is smaller than $\hat{\alpha}$, the former local optimal achieves a higher objective value, and when α is larger than $\hat{\alpha}$, the latter local optimal achieves a higher objective value. The optimal solution jumps from one local optimal to another as α passes $\hat{\alpha}$. To help better understand this phenomenon, let $\Pi(d)$ denote the highest objective value in (P1) with the added constraints $q_1 - q_2 = d$ and $q_2 = \cdots = q_n$. (So we can solve (P1) by maximizing $\Pi(d)$ over d.) Figures 2.2(a) and 2.2(b) show $\Pi(d)$ plotted against d for two distinct values of α in the example from Figure 2.1(b). In both Figure 2.2(a) and Figure 2.2(b), one can see two peaks corresponding to the aforementioned two local optima. In Figure 2.2(a) in which $\alpha = 1.98$, the left peak achieves a higher objective value, and thus the entries of $\boldsymbol{q}^*$ are all equal ($d = 0$). In Figure 2.2(b) in which $\alpha = 1.99$, the right peak achieves a higher objective value, and thus $q_1^* > q_2^*$ at optimality ($d > 0$). The discontinuity in Figure 2.1(b) arises as we move from a range in which the left peak is higher to a range in which the right peak is higher. (We can show that when $n \leq 2$, the optimal $\boldsymbol{q}^*$ is always continuous in α, and thus the discontinuity does not appear in Figure 2.1(a). The proof is given in Appendix 4.2.)

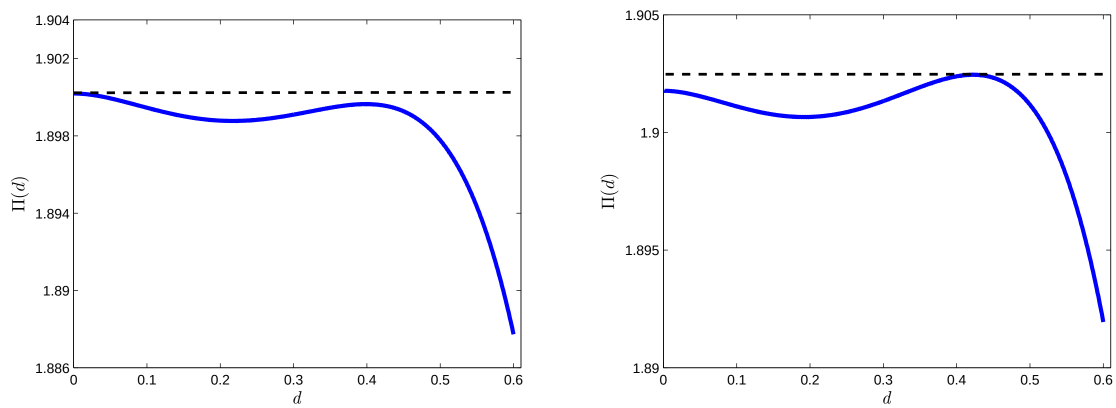

(a) $\alpha = 1.98$ (b) $\alpha = 1.99$

Figure 2.2: Two local optima of $\Pi(d)$

In addition to the sales $\boldsymbol{q}^*$, we are also interested in the optimal prices $\boldsymbol{p}^* = (p_1(\boldsymbol{q}^*), \ldots, p_n(\boldsymbol{q}^*))$ induced by (2.7). We have the following theorem describing the behavior of $\boldsymbol{p}^*$ for different values of α. The proof, which relies heavily on Theorem 2.3.1, is in Appendix 4.2.

Theorem 2.3.2. *Let $\hat{\alpha}$ be defined as in Theorem 2.3.1.*

(a) *If $\alpha \leq \hat{\alpha}$, then $p_1^* = p_2^* = \cdots = p_n^* = p^*$ and one of the following three scenarios holds.*

(i) *p^* increases monotonically in α, or*

(ii) *p^* decreases monotonically in α, or*

(iii) *p^* first decreases and then increases in α.*

(b) *If* $\alpha > \hat{\alpha}$*, then* $p_1^* < p_2^* = \cdots = p_n^*$ *and* p_2^* *increases in* α*. Moreover,* $\lim_{\alpha\to\infty} p_1^* = \infty$ *and* $\lim_{\alpha\to\infty} p_2^* = \infty$.

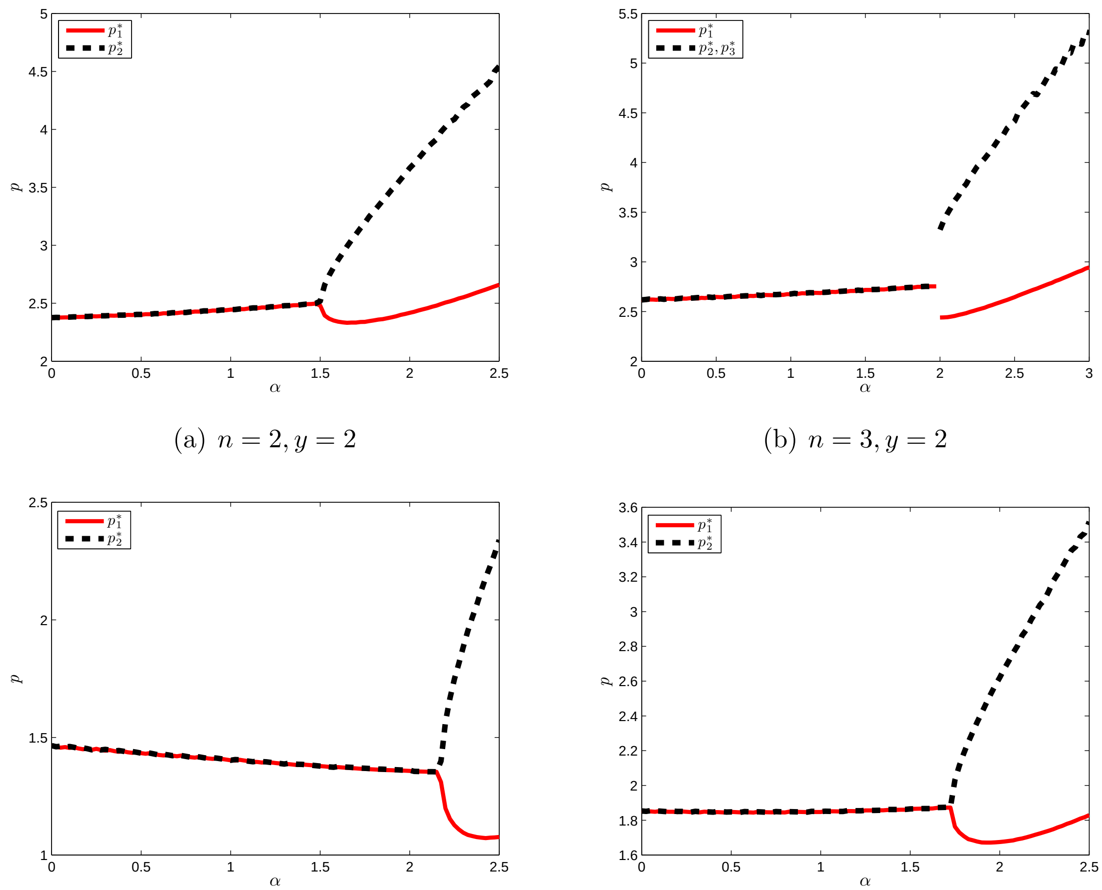

(a) $n = 2, y = 2$ (b) $n = 3, y = 2$

(c) $n = 2, y = 0$ (d) $n = 2, y = 1$

Figure 2.3: Optimal $\boldsymbol{p}^*$ as a function of α

In Figure 2.3, we use four examples to illustrate the behavior of $\boldsymbol{p}^*$ in Theorem 2.3.2. Figures 2.3(a) and 2.3(b) correspond to the same two cases as Figures 2.1(a) and 2.1(b). In both examples, the optimal prices increase in α when $\alpha \leq \hat{\alpha}$, which corresponds to the first scenario in part (a). In Figure 2.3(c), we take $n = 2$,

$y = 0$ and in this case, the optimal prices decrease in α when $\alpha \leq \hat{\alpha}$, which corresponds to the second scenario in part (a). In Figure 2.3(d), we take $n = 2$, $y = 1$ and now the optimal prices first decrease and then increase in α when $\alpha \leq \hat{\alpha}$, which corresponds to the third scenario in part (a). With only the eye, it may be difficult to discern the pattern of the prices to the left of $\hat{\alpha}$ in Figure 2.3(d); however, our numerical calculations confirm that the prices are indeed decreasing and then increasing in that range. In all four figures, when α exceeds $\hat{\alpha}$, the higher price(s) monotonically increase in α while the lower price may initially decrease but eventually will increase in α.

As the theorem shows, the behavior of the optimal prices $\boldsymbol{p}^*$ in response to network effects is more complicated than that of the optimal sales quantities $\boldsymbol{q}^*$. This can be explained by considering two forces driving the direction of the pricing strategy in response to an increase in α. One is to increase the prices to get more unit revenue; the other is to pull down the prices to generate more demand. The results in Theorem 2.3.2 reflect this tradeoff. When α is small, the combined force could be in either direction; when α grows sufficiently large, i.e., the products have already attracted a large share of demand, the first force is stronger, and thus it is more profitable for the seller to raise the prices in response to the increase in α.

In the remainder of this section, we discuss the main steps of the proof of Theorem 2.3.1. (The complete proof is provided in the appendix.) We start with a proposition. The conditions in (2.11) below are derived from the first-order optimality conditions for (P1).

Proposition 2.3.3. *There are at most two distinct entries in* $\boldsymbol{q}^* = (q_1^*, \ldots, q_n^*)$, *and*

$$2\alpha q_i^* - \log q_i^* = C\left(\sum\nolimits_{j \in \mathcal{N}} q_j^*\right) \qquad \textit{for all } i \in \mathcal{N}, \tag{2.11}$$

where $C(\sigma) = \frac{1}{1-\sigma} - \log(1-\sigma) - y$.

To provide some insight into Theorem 2.3.1, a key point is that the expression on the right side in (2.11) does not depend upon i, and hence the expression on the left, $2\alpha q_i^* - \log q_i^*$ must be the same for all $i \in \mathcal{N}$ as well. It follows from the strict convexity of the function $h(q) = 2\alpha q - \log q$ that there are at most two distinct entries in $\boldsymbol{q}^*$ as indicated in Theorem 2.3.1. In the appendix, we show that among the entries of $\boldsymbol{q}^*$, at most one takes the larger value, while all others must take the smaller value.

To prove the portions of the theorem that describe how the sales $\boldsymbol{q}^*$ vary with α, we apply the variable transformation mentioned earlier. Let $\boldsymbol{s} = A\boldsymbol{q}$, where A is given by (2.10). Then $\boldsymbol{s} = (s_1, \ldots, s_n)$ can be written in terms of $\boldsymbol{q}$ as follows

$$s_i = \begin{cases} \frac{1}{\sqrt{n}} \sum_{j=1}^{n} q_j & i = 1 \\ \frac{1}{\sqrt{(i-1)i}} \sum_{j=1}^{i-1} q_j - \frac{i-1}{\sqrt{(i-1)i}} q_i & i = 2, \ldots, n. \end{cases} \tag{2.12}$$

As explained earlier, the change of variables above allows us to appeal to monotonicity results about maximizers of supermodular functions. Since A is an orthogonal matrix, $\boldsymbol{q} = A^{-1}\boldsymbol{s} = A^T\boldsymbol{s}$, and thus q_i can be written as

$$q_i = \begin{cases} \frac{s_1}{\sqrt{n}} + \sum_{j=2}^{n} \frac{s_j}{\sqrt{(j-1)j}} & i = 1 \\ \frac{s_1}{\sqrt{n}} - \frac{(i-1)s_i}{\sqrt{(i-1)i}} + \sum_{j=i+1}^{n} \frac{s_j}{\sqrt{(j-1)j}} & i = 2, \ldots, n. \end{cases} \tag{2.13}$$

After substituting for $\boldsymbol{q}$ in terms of $\boldsymbol{s}$, the objective function of (P1) becomes

$$\tilde{\pi}(s_1, \ldots, s_n) = \alpha \sum_{i=1}^{n} s_i^2 + \sqrt{n}s_1\left(y + \log(1 - \sqrt{n}s_1)\right) - \sum_{i=1}^{n} q_i \log q_i, \tag{2.14}$$

where q_i is defined in (2.13). The constraints of (P1) are equivalent to

$$\frac{1}{\sqrt{n}} \geq s_1 \geq \sqrt{n-1}s_n, \qquad s_2 \geq 0, \qquad \text{and} \qquad s_{i+1} \geq \sqrt{\frac{i-1}{i+1}}s_i, \qquad i = 2, \dots, n-1.$$

Therefore, we can equivalently write (P1) as an optimization problem over $\boldsymbol{s}$:

$$\begin{aligned} \max \quad & \tilde{\pi}(s_1, \dots, s_n) \\ \text{s.t.} \quad & \frac{1}{\sqrt{n}} \geq s_1 \geq \sqrt{n-1}s_n \\ & s_2 \geq 0 \\ & s_{i+1} \geq \sqrt{\frac{i-1}{i+1}}s_i, \; i = 2, \dots, n-1. \end{aligned} \tag{P2}$$

It follows from the preceding developments that a vector $\boldsymbol{q}$ of sales quantities is optimal for (P1) if and only if $\boldsymbol{s} = A\boldsymbol{q}$ is optimal for (P2). We have the following result about problem (P2).

Proposition 2.3.4. *$\tilde{\pi}(s_1, \dots, s_n)$ is supermodular in $(\boldsymbol{s}, \alpha)$ and the feasible region of (P2) is a sublattice in $\mathbb{R}^n$. There exists an entry-wise maximal optimal solution $\boldsymbol{s}^* = (s_1^*, \dots, s_n^*)$ to (P2) such that $s_i^* \geq s_i$ for $i = 1, \dots, n$ for any other optimal $\boldsymbol{s} = (s_1, \dots, s_n)$. Moreover, $\boldsymbol{s}^*$ increases monotonically in α.*

The uniqueness of $\boldsymbol{q}^*$ in (2.9) follows from Proposition 2.3.4. To see this, observe that if there are multiple optimal $\boldsymbol{q}$ for (P1), then (2.9) picks among them by selecting those for which the sum of the entries of $\boldsymbol{s} = A\boldsymbol{q}$ is largest. The proposition establishes that there is an entry-wise maximal optimal solution $\boldsymbol{s}^*$ to (P2). Of course, $\boldsymbol{s}^*$ must also have the property that the sum of its entries is strictly greater than that of any other optimal $\boldsymbol{s}$. So $\boldsymbol{q}^*$ is unique and $\boldsymbol{q}^* = A^{-1}\boldsymbol{s}^*$.

Using the monotonicity properties of $\boldsymbol{s}^*$ from the proposition, we can prove the monotonicity properties of $\boldsymbol{q}^*$ that are stated in Theorem 2.3.1. Details appear in Appendix 4.2.

2.4 The Heterogeneous Case

In this section, we consider the general problem (P0) in which network sensitivities (the α_i), price sensitivities (the γ_i), and intrinsic utilities (the y_i) differ across products. In contrast to the homogeneous setting described in Section 2.3, here the entries of optimal quantity and price vectors will generally all be distinct. Nevertheless, we show below that the optimal solutions in the heterogeneous case have a structure similar to that which underlies Theorems 2.3.1 and 2.3.2 in the homogeneous setting. In addition, based on this structure, we show that (P0) remains amenable to solution for general problems with heterogeneous product parameters.

We will begin with a proposition that describes the first-order optimality conditions for problem (P0).

Lemma 2.4.1. *Any optimal solution $\boldsymbol{q}^\dagger = (q_1^\dagger, \ldots, q_n^\dagger)$ to (P0) must satisfy*

$$2\alpha_i q_i^\dagger - \log q_i^\dagger = C_i(\boldsymbol{q}^\dagger) \quad \text{for all } i \in \mathcal{N}, \tag{2.15}$$

where $C_i(\boldsymbol{q}) = \gamma_i \frac{\sum_{j=1}^n q_j/\gamma_j}{1-\sum_{j=1}^n q_j} - \log(1 - \sum_{j=1}^n q_j) + 1 - y_i$.

To further understand the form of the solution to (P0), we need to understand the behavior of the expression that appears on the left side of (2.15).

Lemma 2.4.2. *The function $h_\alpha(q) = 2\alpha q - \log q$ is convex in q. In addition, we have the following.*

1. *For $c > 1 + \log(2\alpha)$, $h_\alpha(q) = c$ has two solutions, $\underline{q}^c < \overline{q}^c$. Furthermore, $\underline{q}^c \in (0, 1/2\alpha)$ and $\underline{q}^c$ decreases in c, while $\overline{q}^c \in (1/2\alpha, \infty)$ and $\overline{q}^c$ increases in c.*

2. For $c = 1 + \log(2\alpha)$, there is only one solution $q = 1/2\alpha$ to $h_\alpha(q) = c$.

3. For $c < 1 + \log(2\alpha)$, there is no solution to $h_\alpha(q) = c$.

Proofs of the preceding lemmas can be found in Appendix 4.2. For $i \in \mathcal{N}$ and $c \geq 1 + \log(2\alpha_i)$, define $\underline{q}_i^c \leq \overline{q}_i^c$ to be the solutions to $h_{\alpha_i}(q) = c$. (For $c = 1 + \log(2\alpha_i)$, we have $\underline{q}_i^c = \overline{q}_i^c = 1/2\alpha_i$.) It is easy to find $\underline{q}^{c_i}$ and $\overline{q}^{c_i}$ by the convexity of $h_{\alpha_i}(q)$. We now have the following proposition, which describes an important property of the optimal solution to (P0).

Proposition 2.4.3. *Any optimal solution $\boldsymbol{q}^\dagger$ to (P0) must have one of the following two structures:*

(a) $q_i^\dagger = \underline{q}_i^{C_i}$ for all $i \in \mathcal{N}$; or

(b) $q_i^\dagger = \overline{q}_i^{C_i}$ and $q_j^\dagger = \underline{q}_j^{C_j}$ for all $j \in \mathcal{N} \backslash \{i\}$ for some single $i \in \mathcal{N}$

where $C_k = C_k(\boldsymbol{q}^\dagger)$ for $k \in \mathcal{N}$ and $C_k(\cdot)$ is defined in Lemma 2.4.1.

From the proposition, we see that for an optimal solution, there is at most one sales quantity that takes the "high value" $\overline{q}_i^{C_i}$, while all the other sales quantities must take their "low values" $\underline{q}_j^{C_j}$. Recall that we also found this form of solution in Theorem 2.3.1 for the homogeneous setting. Therefore, although we have distinct q_i with heterogeneous parameters across different i, the structure of the solution remains similar. Thus, a general optimal pricing strategy for the multi-product pricing problem with network effects is either to maintain a semblance of balance among all products (all q_i take the low value) or else to boost the sales of just one product (only one q_i takes the high value).

Proposition 2.4.3 essentially says that it is never optimal to promote more than one product simultaneously. To understand this, suppose we fix the sum $\sum_{j=1}^{n} q_j = K$. Then the objective function (2.6) can be rewritten as $\pi(\boldsymbol{q}) = \sum_{j=1}^{n} \pi_j(q_j)$ where $\pi_j(q_j) = \frac{1}{\gamma_j}\left(\alpha_j q_j^2 + (y_j + \log(1-K))q_j - q_j \log q_j\right)$. We have $\pi_j''(q_j) = \frac{1}{\gamma_j}\left(2\alpha_j - \frac{1}{q_j}\right)$, so $\pi_j(q_j)$ is concave when q_j is low ($q_j < 1/2\alpha_j$) and is convex when q_j is high ($q_j > 1/2\alpha_j$). Suppose at an optimal solution $\boldsymbol{q}$ there are two entries q_i and q_j that lie in the regions of convexity. Consider $\boldsymbol{q}^{\epsilon} = \boldsymbol{q} + \epsilon \boldsymbol{e}_{ij}$ and $\boldsymbol{q}^{-\epsilon} = \boldsymbol{q} - \epsilon \boldsymbol{e}_{ij}$ where $\boldsymbol{e}_{ij}$ is an n-vector of zeros except that its i-th entry is 1 and its j-th entry is -1. When ϵ is sufficiently small, due to the local convexity of $\pi_i(\cdot)$ and $\pi_j(\cdot)$ at q_i and q_j respectively, we have $\pi_i(q_i+\epsilon) + \pi_i(q_i-\epsilon) > 2\pi_i(q_i)$ and $\pi_j(q_j+\epsilon) + \pi_j(q_j-\epsilon) > 2\pi_j(q_j)$. Therefore $\pi(\boldsymbol{q}^{\epsilon}) + \pi(\boldsymbol{q}^{-\epsilon}) > 2\pi(\boldsymbol{q})$, implying at least one of $\boldsymbol{q}^{\epsilon}$ and $\boldsymbol{q}^{-\epsilon}$ performs better than $\boldsymbol{q}$, which is a contradiction with the optimality of $\boldsymbol{q}$. A more rigorous proof of Proposition 2.4.3 is provided in Appendix 4.2.

In view of the non-concavity of the objective function (2.6), it is important to develop a specialized computational approach for (P0). We take this up next. Given a $\boldsymbol{c} = (c_1, \ldots, c_n)$, consider the vectors $\boldsymbol{q}$ for which

$$2\alpha_i q_i - \log q_i = c_i \quad \text{for all } i \in \mathcal{N}. \tag{2.16}$$

Note that there are at most 2^n such $\boldsymbol{q}$. For each such $\boldsymbol{q}$, if $C_i(\boldsymbol{q}) = c_i$ for all $i \in \mathcal{N}$ then we have a $\boldsymbol{q}$ that satisfies (2.15) and that is an initial candidate to be an optimal solution to (P0). A key insight is that by Proposition 2.4.3, given a $\boldsymbol{c}$, we can a priori eliminate all but $(n+1)$ of the vectors $\boldsymbol{q}$ that satisfy (2.16). By searching through vectors $\boldsymbol{c}$ that can potentially appear on the right side of (2.15), we arrive at a set of "non-eliminated" candidate solutions $\boldsymbol{q}$ each of which satisfies (2.15). We evaluate the objective function at each element in that set. The one with the greatest objective value is an optimal solution.

Such a search may at first seem impractical because $\boldsymbol{c}$ is an n-dimensional vector. However, it turns out that the search through $\boldsymbol{c}$ can be reduced to a two-dimensional search. To motivate the approach, observe that if we let $K_1 = \sum_{j=1}^n q_j^\dagger$ and $K_2 = \sum_{j=1}^n q_j^\dagger/\gamma_j$, then $C_i(\boldsymbol{q}^\dagger)$ in (2.15) is given by $C_i(\boldsymbol{q}^\dagger) = \gamma_i K_2/(1-K_1) - \log(1-K_1) + 1 - y_i$. Hence, it suffices to search over values of K_1 and K_2 to find sales vectors that satisfy the necessary condition (2.15). To implement the algorithm one must specify discretization grids and stopping rules. See Rayfield et al. (2015) for such developments for a nested logit pricing problem without network effects. We do not pursue this here because it is not the main focus of this chapter. We summarize the above procedures in Algorithm 1. After running the algorithm, we can also compute the optimal prices using (2.5).

Algorithm 1. *1. Let γ_{min} and γ_{max} be the minimum and maximum values among $\gamma_1, \dots, \gamma_n$.*

2. Let $Q = \emptyset$. For K_1 from 0 to 1, for K_2 from K_1/γ_{max} to K_1/γ_{min}:

(a) Calculate $c_i = \gamma_i K_2/(1-K_1) - \log(1-K_1) + 1 - y_i$ for $i = 1, \dots, n$. If $c_i < 1 + \log(2\alpha_i)$ for any $i \in \mathcal{N}$, then skip steps (b)–(d) and continue to the next (K_1, K_2) pair.

(b) Solve for $\underline{q}_i^{c_i}$ and $\overline{q}_i^{c_i}$ that satisfy $2\alpha_i q - \log q = c_i$ for $i \in \mathcal{N}$.

(c) Let $\underline{\boldsymbol{q}} = (\underline{q}_1^{c_1}, \underline{q}_2^{c_2}, \dots, \underline{q}_n^{c_n})$ and for $i = 1, \dots, n$ let $\overline{\boldsymbol{q}}^i = (\underline{q}_1^{c_1}, \dots, \underline{q}_{i-1}^{c_{i-1}}, \overline{q}_i^{c_i}, \underline{q}_{i+1}^{c_{i+1}} \dots, \underline{q}_n^{c_n})$. Here, $\overline{\boldsymbol{q}}^i$ is the same as $\underline{\boldsymbol{q}}$ except that the i-th entry is replaced by $\overline{q}_i^{c_i}$. Let $R = \{\underline{\boldsymbol{q}}, \overline{\boldsymbol{q}}^1, \overline{\boldsymbol{q}}^2, \dots, \overline{\boldsymbol{q}}^n\}$.

(d) *For each* $\boldsymbol{q} \in R$*: if both* $K_1 = \sum_{j=1}^{n} q_j$ *and* $K_2 = \sum_{j=1}^{n} q_j/\gamma_j$ *hold then let* $Q = Q \cup \{\boldsymbol{q}\}$.

3. *End For-Loop over* (K_1, K_2).

4. *Evaluate (2.6) at each element in* Q *to find* $\boldsymbol{q}^\dagger$ *that maximizes (2.6).*

The above developments reveal an important structural property of the optimal solution and provide an effective approach to solve the general problem (P0). However, they do not tell us which product (if any) should have the high sales value. In the ensuing subsections, we describe two special cases for which we are able to more precisely identify the structure of the solution and more easily solve the problem.

2.4.1 Heterogeneous Network Sensitivities

In this section, we consider a version of (P0) in which only network sensitivities differ across products. To be more specific, we assume $\alpha_1 > \cdots > \alpha_n$, and $y_i = y$ and $\gamma_i = \gamma = 1$ for $i \in \mathcal{N}$. (There is no additional loss of generality in taking $\gamma = 1$.) In this case, the revenue function (2.6) becomes

$$\pi(\boldsymbol{q}) = \sum_{j=1}^{n} \alpha_j q_j^2 + \sum_{j=1}^{n} q_j \left(y + \log\left(1 - \sum_{j=1}^{n} q_j\right)\right) - \sum_{j=1}^{n} q_j \log q_j. \tag{2.17}$$

We refer to the revenue maximization problem (P0) with the objective function specialized to (2.17) as (P3). The following proposition summarizes properties of the optimal solution to (P3). The proof is in Appendix 4.2.

Proposition 2.4.4. *For any optimal solution* $\boldsymbol{q}^\dagger = (q_1^\dagger, \ldots, q_n^\dagger)$ *to (P3)*

1. $\boldsymbol{q}^\dagger$ *must satisfy* $q_1^\dagger > \cdots > q_n^\dagger$ *and*

$$2\alpha_i q_i^\dagger - \log q_i^\dagger = C\left(\sum\nolimits_{j\in\mathcal{N}} q_j^\dagger\right) \qquad \text{for all } i \in \mathcal{N}, \tag{2.18}$$

where $C(\sigma) = \frac{1}{1-\sigma} - \log(1-\sigma) - y$.

2. *For* $\boldsymbol{q}^\dagger$*, we have* $C \geq 1 + \log(2\alpha_1)$*, where* $C := C\left(\sum\nolimits_{j\in\mathcal{N}} q_j^\dagger\right)$*. In addition,* $q_1^\dagger \in \{\underline{q}_1^C, \overline{q}_1^C\}$ *and* $q_i^\dagger = \underline{q}_i^C$ *for* $i = 2, \ldots, n$.

3. *The optimal prices* $\boldsymbol{p}^\dagger = (p_1^\dagger, \ldots, p_n^\dagger) = (p_1(\boldsymbol{q}^\dagger), \ldots, p_n(\boldsymbol{q}^\dagger))$ *obtained from* $\boldsymbol{q}^\dagger$ *and (2.5) must satisfy* $p_1^\dagger < \cdots < p_n^\dagger$.

From the proposition, we see that an optimal solution to problem (P3) follows the similar structure as in Proposition 2.4.3, i.e., there is at most one entry in $\boldsymbol{q}^\dagger$ that takes the "high value". However, in this special case, we know for certain that if one of these products needs to be promoted, the optimal choice is always product 1 — the one associated with the largest network sensitivity parameter. The result is not surprising. If we are given a set of products that have the same intrinsic qualities and that possess equal price sensitivities and if the plan is to promote exactly one among them, then the seller benefits the most by choosing the product with the strongest network effect.

In this setting with intrinsic utilities and price sensitivities that are common across all products, Algorithm 1 can be simplified to a one-dimensional search. To see this, note that the right side of (2.18) depends just upon the sum of the entries of $\boldsymbol{q}$. Proposition 2.4.4 further simplifies the search for the optimal solution to (P3), because it tells us that given a value of c, merely two out of the 2^n vectors $\boldsymbol{q}$ that satisfy $2\alpha_i q_i - \log q_i = c$ are left to be evaluated. Specifically, we only need to consider the $\boldsymbol{q}$ where $q_i = \underline{q}_i^c$ for $i = 2, \ldots, n$ and either (a) $q_1 = \underline{q}_1^c$ or (b)

$q_1 = \overline{q}_1^c$. Algorithm 2 below outlines an approach to solve (P3) based upon these obervations.

Algorithm 2. *1. Let $Q = \emptyset$. For σ from 0 to 1:*

(a) Calculate $c = \frac{1}{1-\sigma} - \log(1-\sigma) - y$. If $c < 1 + \log(2\alpha_1)$, then skip steps (b)–(d) and continue to the next σ.

(b) Solve for $\overline{q}_1^c$ and $\underline{q}_i^c$ for $i \in \mathcal{N}$ that satisfy $2\alpha_i q - \log q = c$.

(c) Let $\underline{\boldsymbol{q}} = (\underline{q}_1^c, \underline{q}_2^c, \ldots, \underline{q}_n^c)$ and $\overline{\boldsymbol{q}}^1 = (\overline{q}_1^c, \underline{q}_2^c, \ldots, \underline{q}_n^c)$. Let $R = \{\underline{\boldsymbol{q}}, \overline{\boldsymbol{q}}^1\}$.

(d) For each $\boldsymbol{q} \in R$: if $C(\sum_{j\in\mathcal{N}} q_j) = c$ holds then let $Q = Q \cup \{\boldsymbol{q}\}$.

2. End For-Loop over σ.

3. Evaluate (2.17) at each element in Q to find $\boldsymbol{q}^\dagger$ that maximizes (2.6).

2.4.2 Heterogeneous Price Sensitivities

In this section, we consider a variation in which customers have heterogeneous price sensitivities for the n products, i.e., the γ_i are distinct while other parameters are the same across different products ($\alpha_i = \alpha$ and $y_i = y$). Throughout this section we assume $\gamma_1 < \cdots < \gamma_n$. In this setting, the revenue function (2.6) becomes

$$\pi(\boldsymbol{q}) = \alpha \sum_{j=1}^{n} \frac{q_j^2}{\gamma_j} + \sum_{j=1}^{n} \frac{q_j}{\gamma_j}\left(y + \log\left(1 - \sum_{i=1}^{n} q_i\right)\right) - \sum_{j=1}^{n} \frac{q_j}{\gamma_j} \log q_j. \tag{2.19}$$

We refer to (P0) with the objective function specialized to (2.19) as (P4). The next proposition describes the optimal solution to (P4). The proof is in Appendix 4.2.

Proposition 2.4.5. *For any optimal solution* $\boldsymbol{q}^\dagger = (q_1^\dagger, \ldots, q_n^\dagger)$ *to (P4)*

1. $\boldsymbol{q}^\dagger$ *must satisfy* $q_1^\dagger > \cdots > q_n^\dagger$ *and*

$$2\alpha q_i^\dagger - \log q_i^\dagger = C_i(\boldsymbol{q}^\dagger) \quad \text{for all } i \in \mathcal{N}, \tag{2.20}$$

   where $C_i(\boldsymbol{q}) = \gamma_i \frac{\sum_{j=1}^n q_j/\gamma_j}{1-\sum_{j=1}^n q_j} - \log(1 - \sum_{j=1}^n q_j) + 1 - y.$

2. *At* $\boldsymbol{q}^\dagger$, *we have* $C_i \geq 1 + \log(2\alpha)$, *where* $C_i := C_i(\boldsymbol{q}^\dagger)$ *for* $i = 1, \ldots, n$. *In addition,* $q_1^\dagger \in \{\underline{q}_1^{C_1}, \overline{q}_1^{C_1}\}$ *and* $q_i^\dagger = \underline{q}_i^{C_i}$ *for* $i = 2, \ldots, n$.

The preceding structure is similar to what is described in Proposition 2.4.4, and is also useful for simplifying the computation of $\boldsymbol{q}^\dagger$. In particular, the product with the lowest price sensitivity may take a "high" sales level while all the other products must take a "low" sales level. Proposition 2.4.5 can be easily modified to accommodate heterogeneous intrinsic utilities that are ordered in the direction opposite of the γ_i, i.e., $y_1 \geq y_2 \geq \cdots \geq y_n$.

The computations in the case of heterogeneous price sensitivities considered in this section are more difficult than those in subsection 2.4.1 (which considered heterogeneous α_i) because here $2\alpha q_i - \log q_i$ is no longer independent of i at optimality. Nevertheless, we obtain some simplifications to Algorithm 1. In particular, we can replace (b) and (c) by (b′) and (c′) as follows.

(b′) Solve for $\overline{q}^{c_1}$ and $\underline{q}^{c_i}$ that satisfy $2\alpha q - \log q = c_i$ for $i \in \mathcal{N}$.

(c′) Let $\underline{\boldsymbol{q}} = (\underline{q}^{c_1}, \underline{q}^{c_2}, \ldots, \underline{q}^{c_n})$ and let $\overline{\boldsymbol{q}}^1 = (\overline{q}^{c_1}, \underline{q}^{c_2}, \ldots, \underline{q}^{c_n})$. Here, $\overline{\boldsymbol{q}}^1$ is the same as $\underline{\boldsymbol{q}}$ except that the first entry is replaced by $\overline{q}^{c_1}$. Let $R = \{\underline{\boldsymbol{q}}, \overline{\boldsymbol{q}}^1\}$.

2.5 Three Extensions

In this section, we consider three extensions to the basic MNL model with network effects. In Section 2.5.1, we allow inter-product network effects. In Section 2.5.2, we consider the possibility that network effects may enter customers' expected utility in a non-linear fashion. In Section 2.5.3, we consider network effects for the no-purchase option.

2.5.1 Inter-product Network Effects

In this section, we incorporate inter-product network effects into consumers' utilities. More precisely, we replace the customer's utility function (2.1) with

$$v_i = y_i - \gamma_i p_i + \alpha_i x_i + \beta_i x_{-i},$$

where $\beta_i \geq 0$ is the inter-product network sensitivity parameter, and $x_{-i} = \sum_{j \in \mathcal{N} \setminus \{i\}} x_j$ is the total consumption of all products other than i. Hence, the new term $\beta_i x_{-i}$ represents the additional expected utility for product i from the market's consumption of products other than i. We assume $\alpha_i > \beta_i$, which means that the within-product sensitivity is stronger than the inter-product sensitivity.

In the following, we consider a homogeneous case in which $y_i = y, \gamma_i = 1, \alpha_i = \alpha$, and $\beta_i = \beta$ for all $i \in \mathcal{N}$. Now the expression (2.7) becomes

$$p_i(\boldsymbol{q}) = \alpha q_i + \beta q_{-i} - \log q_i + \log\left(1 - \sum_{j=1}^{n} q_j\right) + y,$$

and the revenue function (2.8) becomes

$$\pi(\boldsymbol{q}) = (\alpha - \beta)\sum_{j=1}^{n} q_j^2 + \sum_{j=1}^{n} q_j \left(y + \beta \sum_{j=1}^{n} q_j + \log\left(1 - \sum_{j=1}^{n} q_j\right)\right) - \sum_{j=1}^{n} q_j \log q_j.$$

As in (P1), we may assume $q_1 \geq \cdots \geq q_n$ without loss of optimality. With the same means of selecting a particular optimal solution as in (2.9), it turns out that most of the results in Section 2.3 for (P1) still hold in this new setting.

In particular, Theorem 2.3.1 holds with the minor modification that the threshold $\hat{\alpha}$ (which now depends upon y, n, and β) satisfies $\beta + 1/2 < \hat{\alpha} \leq \alpha^{\bar{R}}$ where $\alpha^{\bar{R}}$ is the smallest solution to $\bar{R}(\alpha) = 0$ where

$$\bar{R}(\alpha) = y + \log(2\alpha - 2\beta - n) - \frac{n}{2\alpha - 2\beta - n} + \frac{n\beta}{\alpha - \beta}.$$

Theorem 2.3.2 carries over without modification to this setting. The proofs of these results follow almost exactly as the proofs of Theorems 2.3.1 and 2.3.2 with only minor changes. A key step is using Proposition 2.3.3, but with α replaced by $\alpha - \beta$ and $C(\sigma)$ redefined as $\frac{1}{1-\sigma} - \log(1 - \sigma) - y - 2\beta\sigma$. (We note that we opted to consider the case without inter-product effects in the main body of the chapter to help keep arguments as transparent as possible.)

2.5.2 Non-linear Network Effects

In this section, we consider a more general form of network effects. More precisely, we assume the customer's utility function (2.1) is replaced by

$$v_i = y_i - \gamma_i p_i + f_i(x_i),$$

where $f_i(x_i)$ is the expected utility gained from the network effects. We assume that $f_i(0) = 0$ and that $f_i(\cdot)$ is increasing. For simplicity, we also assume $f_i(\cdot)$ is differentiable. With the same normalization as in (2.3), the expression for prices (2.5) becomes

$$p_i(\boldsymbol{q}) = \frac{1}{\gamma_i}\left(f_i(q_i) - \log q_i + \log\left(1 - \sum_{j=1}^{n} q_j\right) + y_i\right),$$

and the revenue function (2.6) becomes

$$\pi(\boldsymbol{q}) = \sum_{j=1}^{n} \frac{q_j f_j(q_j)}{\gamma_j} + \sum_{j=1}^{n} \frac{q_j}{\gamma_j} \log\left(1 - \sum_{j=1}^{n} q_j\right) + \sum_{j=1}^{n} \frac{q_j}{\gamma_j}\left(y_j - \log q_j\right). \tag{2.21}$$

Using an analysis as in Section 2.4, we see that the optimal $\boldsymbol{q}^\dagger$ must satisfy the first-order necessary conditions, and thus (2.16) becomes

$$q_i^\dagger f_i'(q_i^\dagger) + f_i(q_i^\dagger) - \log q_i^\dagger = C_i(\boldsymbol{q}^\dagger) \quad \text{for all } i \in \mathcal{N},$$

where $C_i(\boldsymbol{q})$ is as defined in Lemma 2.4.1. Note that if $f_i(q) = \alpha_i q$, then $q_i^\dagger f_i'(q_i^\dagger) + f_i(q_i^\dagger) = 2\alpha_i q_i^\dagger$ in which case the preceding condition simplifies to (2.15).

Define $H_i(q) = q f_i'(q) + f_i(q) - \log q$. Below, we assume $H_i''(q) > 0$ for $q > 0$. In that case, for $c \geq \min_q H_i(q)$, we re-define $\underline{q}_i^c \leq \overline{q}_i^c$ to be the solutions $H_i(q) = c$. (If there is only solution to $H_i(q) = c$, then we let both $\underline{q}_i^c$ and $\overline{q}_i^c$ be that solution.) The original definition (for problems with $f_i(q) = \alpha_i q$) appears in the paragraph that follows Lemma 2.4.2.

Proposition 2.5.1. *Suppose $H_i''(q) > 0$ for all $q > 0$ for each $i \in \mathcal{N}$. Then Proposition 2.4.3 holds in the more general setting of* (2.21)*, with $\underline{q}_i^c$ and $\overline{q}_i^c$ redefined as above.*

Examples of functions $f_i(\cdot)$ such that $H_i''(q) > 0$ for all $q > 0$ are $f_i(q) = \log(1 + q)$ and $f_i(q) = \alpha q^\theta$ where $\alpha > 0$ and $\theta \geq 1$. The proof of Proposition 2.5.1 follows almost exactly as the proofs of Lemma 2.4.2 and Proposition 2.4.3. As long as the network effect term in the utility function satisfies the conditions in Proposition 2.5.1, the optimal strategy remains to either maintain a semblance of balance among all products or else to boost the sales of just one product.

2.5.3 Network Effects for the No-Purchase Option

Here, we consider an extension of (P0) in which an individual customer's utility from not purchasing any product depends on the fraction of customers that do not purchase. This setting captures the idea that the customers who do not choose the products offered by our seller could buy a product from someone else. In this case, (2.1) remains unchanged but v_0 becomes

$$v_0 = \alpha_0 x_0,$$

where x_0 is the number of customers that do not buy any product offered by this firm. With the same normalization as in (2.3), the expression (2.5) becomes

$$p_i(\boldsymbol{q}) = \frac{1}{\gamma_i}\left(\alpha_i q_i - \log q_i + \log\left(1 - \sum_{j=1}^{n} q_j\right) + y_i - \alpha_0\left(1 - \sum_{j=1}^{n} q_j\right)\right),$$

and the revenue function (2.6) becomes

$$\begin{aligned}\pi(\boldsymbol{q}) = \sum_{j=1}^{n} q_j p_j(\boldsymbol{q}) = \sum_{j=1}^{n} \frac{\alpha_j}{\gamma_j} q_j^2 + \\ \sum_{j=1}^{n} \frac{q_j}{\gamma_j}\left(y_j + \log\left(1 - \sum_{j=1}^{n} q_j\right) - \log q_j - \alpha_0\left(1 - \sum_{j=1}^{n} q_j\right)\right).\end{aligned}$$

As in Section 2.4, the optimal $\boldsymbol{q}^\dagger$ must satisfy the first-order necessary conditions

$$2\alpha_i q_i^\dagger - \log q_i^\dagger = D_i(\boldsymbol{q}^\dagger) \quad \text{for all } i \in \mathcal{N}, \tag{2.22}$$

where $D_i(\boldsymbol{q}) = C_i(\boldsymbol{q}) + \alpha_0\left(1 - \sum_{j=1}^{n} q_j\right) - \alpha_0 \gamma_i \sum_{j=1}^{n} q_j/\gamma_j$ and $C_i(\boldsymbol{q})$ is defined in Lemma 2.4.1.

Since the left hand side of (2.22) remains the same as in (2.15), Lemma 2.4.2 still applies and the results of Proposition 2.4.3 still hold in this setting. The proofs follow exactly as before. Therefore even if the non-purchase option is sensitive to network effects, the optimal strategy follows the same structure as in our basic model.

2.6 Numerical Results

In this section, we describe some numerical experiments, which are divided into three parts. First, we investigate the issue of possible multiple equilibria at optimal prices and study how this may affect the implementation of the solution. Second, for both homogeneous and heterogeneous cases, we demonstrate the importance of taking into account network effects when making pricing decisions. In this portion of the numerical study, we also show that Algorithm 1 solves the general problem quickly. In the final portion of the numerical experiments, we describe tests showing that the decisions obtained from our model are robust with respect to estimation errors in the network strength parameters.

2.6.1 The Issue of Multiple Equilibria

As discussed earlier, one issue with the network MNL model is that given a price vector $\boldsymbol{p}$, there are possibly multiple $\boldsymbol{q}$ that satisfy the equilibrium condition (2.4). We are particularly interested in this issue at prices $\boldsymbol{p}^* = \boldsymbol{p}(\boldsymbol{q}^*)$, which are the optimal prices obtained from our model. Appendix 4.2 addresses the question of uniqueness and stability of solutions to (2.4). In this section, we study whether the sales vector will converge to $\boldsymbol{q}^*$ when prices are set at $\boldsymbol{p}^*$ if customers repeatedly adjust their purchases according to market conditions or their perceptions thereof.

Such issues are of interest even if $\boldsymbol{q}^*$ is the unique equilibrium at prices $\boldsymbol{p}^*$. In addition, if we find that there is convergence to $\boldsymbol{q}^*$ from all tested starting points of the adjustment process, then this provides some evidence that $\boldsymbol{q}^*$ is the unique equilibrium (or at least the unique stable equilibrium) associated with prices $\boldsymbol{p}^*$. For this numerical study, we do the following:

1. Given a set of parameters $\{(\alpha_i, \gamma_i, y_i) : i = \mathcal{N}\}$, we use Algorithm 1 to find the vector of optimal sales levels $\boldsymbol{q}^*$. Then we calculate $\boldsymbol{p}^* = \boldsymbol{p}(\boldsymbol{q}^*)$ from (2.5).

2. From a starting sales vector $\boldsymbol{q}^0$, we iteratively compute $\boldsymbol{q}^t$ using the following dynamics:

$$q_i^t = \frac{\exp(y_i - \gamma_i p_i^* + \alpha_i q_i^{t-1})}{1 + \sum_{j=1}^{n} \exp(y_j - \gamma_j p_j^* + \alpha_j q_j^{t-1})}\,. \tag{2.23}$$

3. We check whether the sequence $\{\boldsymbol{q}^t\}_{t=0}^{\infty}$ converges to $\boldsymbol{q}^*$.

The dynamics in (2.23) above could arise if there is a sequence of problem instances indexed by t, and customers make their purchase decisions in instance t based upon the sales levels in instance $t-1$. Alternatively, in a single problem instance, customers might hypothesize a vector of sales levels $\boldsymbol{q}^0$, which would suggest to the customers that sales levels would be $\boldsymbol{q}^1$. Realizing this, customers would then hypothesize sales levels $\boldsymbol{q}^1$, which would then suggest that sales levels would be $\boldsymbol{q}^2$. Customers would continue in this fashion until the hypothesized and suggested sales levels are (essentially) the same, at which point the customers would make their purchases, thereby producing an actual vector of sales $\boldsymbol{q}$. That vector will be an equilibrium (that is, satisfy (2.4)) for prices $\boldsymbol{p}^*$. If the seller initially announces that sales levels will be $\boldsymbol{q}^*$ and if the population believes the seller so that $\boldsymbol{q}^0 = \boldsymbol{q}^*$, then the iterative procedure will converge immediately to quantities $\boldsymbol{q}^*$. The above admittedly endows the population of customers with remarkably good mathematical faculties. (We should note that it is quite common that game-theoretic models of customer behavior make strong assumptions about what customers know and/or can compute.)

According to Proposition 4.2.2 in Appendix 4.2, the sales vector $\boldsymbol{q}^t$ must converge to $\boldsymbol{q}^*$ when $|\alpha_i| \leq 2$ for all $i \in \mathcal{N}$. We first show that this is true in our experiments with even more general α_i. In particular, we consider $n = 1, \ldots, 5$ and for each n, we test 1000 randomly generated problems. To generate a problem, we sample α_i, y_i and γ_i for $i = 1, \ldots, n$ from uniform distributions over $[0, 4]$, $[0, 2]$, and $[1, 2]$ respectively. In this first set of experiments, we focus on the starting point $\boldsymbol{q}^0 = \boldsymbol{0} := (0, \ldots, 0)$, which is a natural choice if initially no customers buy any of the products. For each n, we compute the average of $\|\boldsymbol{q}^t - \boldsymbol{q}^*\| = \left(\sum_{i=1}^n (q_i^t - q_i^*)^2\right)^{1/2}$ over the 1000 experiments for different values of t. The results are plotted in Figure 2.4. The figure shows that for each n, the average of $\|\boldsymbol{q}^t - \boldsymbol{q}^*\|$ converges to 0 as t increases. In fact, in all of these experiments $\boldsymbol{q}^t$ converges to $\boldsymbol{q}^*$ starting from $\boldsymbol{q}^0 = \boldsymbol{0}$, and in most cases, $\boldsymbol{q}^t$ becomes very close to $\boldsymbol{q}^*$ in fewer than 20 steps.

In the above experiments, the network strength parameters α_i are not "very large". In the following, we consider some settings with larger values of α_i in which case we will see that $\boldsymbol{q}^t$ may no longer converge to $\boldsymbol{q}^*$. For simplicity, we consider only the homogeneous case in the following. We apply steps 1-3 above for values of α ranging from 0 to 7 with $n = 5$ and $y_i = y = 2, \gamma_i = \gamma = 1$ for

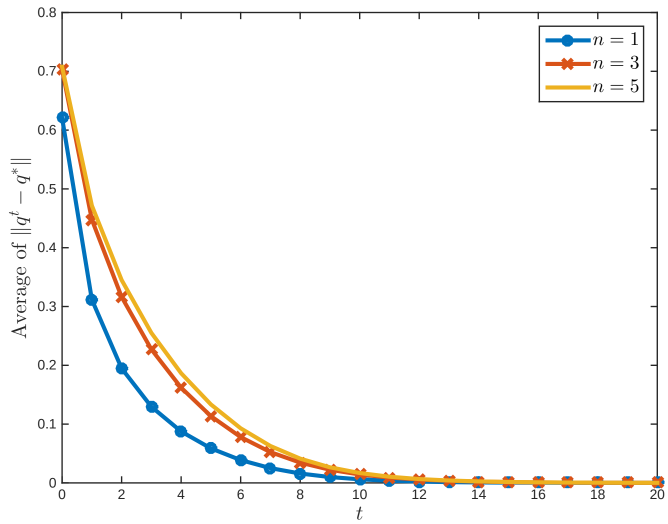

Figure 2.4: Average of $\|\boldsymbol{q}^t - \boldsymbol{q}^*\|$ versus t

$i = 1, \ldots, 5$. For each value of α, we first study whether $\boldsymbol{q}^t$ converges to $\boldsymbol{q}^*$ when starting from $\boldsymbol{q}^0 = \mathbf{0}$. Then, we consider 1000 different values for the initial vector $\boldsymbol{q}^0$ by sampling uniformly at random from the n-simplex, and we determine the fraction of those 1000 cases in which the vector $\boldsymbol{q}^t$ converges to $\boldsymbol{q}^*$. The results are summarized in Table 2.1. Within a cell in the row labeled "$n = 5$", a 'Yes' at the top means that $\boldsymbol{q}^t$ converges to $\boldsymbol{q}^*$ when starting from $\boldsymbol{q}^0 = \mathbf{0}$ and a 'No' means it does not. The value below the 'Yes'/'No' indicates the percentage of the 1000 initial points for which $\boldsymbol{q}^t$ converges to $\boldsymbol{q}^*$. The bottom value in each cell in the $n = 5$ row is the optimal objective function value for $n = 5$; i.e., $\pi_5^* = \sum_{i=1}^5 p_i(\boldsymbol{q}^*)q_i^* = \sum_{i=1}^5 p_i^* q_i^*$. We emphasize that this value is obtained from sales quantities $\boldsymbol{q}^*$.

From Table 2.1, we can see that when starting from $\boldsymbol{q}^0 = \mathbf{0}$, the vector of sales levels $\boldsymbol{q}^t$ converges to $\boldsymbol{q}^*$ for $\alpha \leq 4$, but not for $\alpha \geq 5$. In fact, for $\alpha \leq 4$, there is convergence to $\boldsymbol{q}^*$ from all sampled starting points. However, for $\alpha \geq 5$, the vector

α		0	1	2	3	4	5	6
$n=5$	$\boldsymbol{q}^0=\mathbf{0}$	Yes	Yes	Yes	Yes	Yes	No	No
	Rd $\boldsymbol{q}^0$	100%	100%	100%	100%	100%	11.8%	1.8%
	π_5^*	1.9445	2.0343	2.1295	2.4413	3.1210	3.8807	4.6843
$n=1$	$q^0=0$	Yes	Yes	Yes	Yes	Yes	No	No
	Rd q^0	100%	100%	100%	100%	100%	64.0%	41.7%
	π_1^*	1.0000	1.3216	1.8013	2.4142	3.1170	3.8801	4.6843
$(\pi_5^*-\pi_1^*)/\pi_5^*$		48.57%	35.03%	15.41%	1.11%	0.13%	0.02%	<0.01%

Table 2.1: Convergence results for different values of α

$\boldsymbol{q}^t$ converges to $\boldsymbol{q}^*$ only when $\boldsymbol{q}^0$ is "close enough" to $\boldsymbol{q}^*$. The region that is close enough shrinks as α increases, which is reflected by the decreasing percentage of the randomly drawn initial conditions for which there is convergence to $\boldsymbol{q}^*$. For $\alpha = 7$, just two of the 1000 samples give us convergence to $\boldsymbol{q}^*$.

When $\alpha \geq 5$, the optimal $\boldsymbol{q}^*$ follows the "one high, others low" pattern described in part (b) of Theorem 2.3.1. In the cases for which there is not convergence to $\boldsymbol{q}^*$ from most initial conditions (the columns on the right in the table), the low value in $\boldsymbol{q}^*$ is very close to 0 and the high value is close to 1. For instance, in the $\alpha = 7$ case, we have $q_H^* = 0.9196$ and $q_L^* = 8 \times 10^{-6}$, with corresponding $\pi^* = 5.5178$. In that case, if we start from $\boldsymbol{q}^0 = \mathbf{0}$, then $\boldsymbol{q}^t$ will converge to a different equilibrium point with $q'_H = 0.0207$ and $q'_L = 1 \times 10^{-4}$ and the corresponding revenue is $\pi' = 0.1287$ (we will further comment on this phenomenon at the end of this section). It is important to note that it is not the "one high, others low" pattern itself that yields the lack of convergence to $\boldsymbol{q}^*$. Rather, a lack of convergence to $\boldsymbol{q}^*$ seems to arise in settings for which the low value is very close to zero. If the low value is not "very low", then there is still convergence to $\boldsymbol{q}^*$. For instance, if $\alpha = 4$, then the optimal solution is "one high, others low", with $q_H^* = 0.8546$ and $q_L^* = 9 \times 10^{-4}$, but as can be seen in the table, there is convergence to $\boldsymbol{q}^*$ from $\boldsymbol{q}^0 = \mathbf{0}$ as well as from all 1000 random initial conditions.

Based upon the observation in the preceding paragraph, if α is large enough that the low value in $\boldsymbol{q}^*$ is very small, then the seller may simply stop offering an assortment of n products and instead offer only the one product with the high value in $\boldsymbol{q}^*$. Since the low value in $\boldsymbol{q}^*$ is very small, the revenue loss from dropping those products entirely would be negligible. In this case, the seller would simply solve a single-product pricing problem ($n = 1$) in hope that the single-product problem is such that the optimal (single) price $p(q^*)$ yields a unique equilibrium quantity or at least there is convergence to q^* starting from a larger set of initial conditions.

In the rows of Table 2.1 labeled $n = 1$, we examine this idea. We again consider the convergence of q^t to q^* (for the $n = 1$ case) starting from either 0 or a random initial value selected uniformly on $[0, 1]$. Then we compute the optimal revenue $\pi_1^* = p(q^*)q^*$. From Table 2.1, we can see that the optimal revenue from selling only one product is very close to that from selling all five products when $\alpha \geq 4$. Moreover, the convergence behavior is better when $n = 1$ than when $n = 5$. Even though the issue of multiple equilibria still exists, the chance of convergence starting from a random point is markedly greater when $n = 1$ than when $n = 5$. For example, when $\alpha = 5$, we have the desired convergence from 64% of initial points when $n = 1$ compared to 11.8% when $n = 5$.

We close this section by noting that the issue of multiple equilibria is common in problems involving network effects (see Appendix 4.2 for references) and more generally in problems involving strategic agents. As we note in the appendix, there is no single "right" answer for what to do when multiple equilibria are present. Our analysis covers many settings in which the equilibrium is unique. If there are multiple equilibria, it is possible that a suboptimal equilibrium has very poor performance compared to the optimal one, as we showed earlier in the $\alpha = 7$ case ($\pi' = 0.1287$ versus $\pi^* = 5.5178$). To avoid such an unsatisfactory result, the seller can implement prices $\boldsymbol{p}(\boldsymbol{q}^*)$ and announce the anticipated sales vector $\boldsymbol{q}^*$. If the customers believe the seller's announcement, then the sales will actually be $\boldsymbol{q}^*$. If the customers believe that the sales will be in a neighborhood of $\boldsymbol{q}^*$ then adjustments as described above will lead sales to $\boldsymbol{q}^*$. (As shown above, the neighborhood can be large and even include all possible initial conditions, or it can be small.) If the seller is concerned that some equilibrium other than $\boldsymbol{q}^*$ may prevail when prices $\boldsymbol{p}(\boldsymbol{q}^*)$ are implemented, then that seller may wish to consider a different formulation than (P0). For example, if a seller is particularly pessimistic or sensitive to potential negative outcomes, it may wish to instead maximize $\sum_{i=1}^{n} p_i q_i^{\mathrm{LR}}(\boldsymbol{p})$ over $\boldsymbol{p}$ where $\boldsymbol{q}^{\mathrm{LR}}(\boldsymbol{p}) = (q_1^{\mathrm{LR}}(\boldsymbol{p}), \ldots, q_n^{\mathrm{LR}}(\boldsymbol{p}))$ is the vector of quantities that yields the lowest revenue among those $\boldsymbol{q}$ that satisfy (2.4) for prices $\boldsymbol{p}$. This idea is addressed in Chapter 3.

2.6.2 Importance of Considering Network Effects

In this section, we demonstrate that if products do indeed exhibit network effects, then it is important to take that into account when making pricing decisions. Otherwise, a significant portion of the optimal revenue may be lost. We also show that it is important to consider the possibility of setting different prices for different products even in the case of homogeneous products (recall that optimal prices for homogeneous products will be identical in the absence of network effects).

We start our tests by considering the homogeneous case. In our experiments, we fix $n = 5$, $y = 2$, and $\gamma = 1$. We assume that the network effect parameter α ranges from 0 to 5. For each value of the true α, we consider three pricing strategies:

1. We solve problem (P1) and obtain the optimal sales levels $\boldsymbol{q}^*$. Then we use price vector $\boldsymbol{p}^* = \boldsymbol{p}(\boldsymbol{q}^*)$, which gives revenue $\pi(\boldsymbol{q}^*) = \sum_{i=1}^n p_i^* q_i^*$.

2. We consider the optimal symmetric solution. That is, we impose an additional constraint that $q_1 = \cdots = q_n$ in (P1). Let the optimal solution to this problem be $\boldsymbol{q}^u$ and the corresponding price be $\boldsymbol{p}^u = \boldsymbol{p}(\boldsymbol{q}^u)$. The revenue is $\pi^u = \sum_{i=1}^n p_i^u q_i^u$.

3. We ignore the network effect by solving (P1) with $\alpha = 0$. Let $\boldsymbol{q}^\circ$ denote the optimal solution, and let $\boldsymbol{p}^\circ(\boldsymbol{q})$ denote (2.7) with $\alpha = 0$. We compute prices $\boldsymbol{p}^\circ = \boldsymbol{p}^\circ(\boldsymbol{q}^\circ)$. Note that $\boldsymbol{q}^\circ$ will not be an equilibrium for these prices (except when α is actually 0) because the equilibrium condition (2.4) uses the actual value of α. If there are multiple sales vectors $\boldsymbol{q}$ that satisfy the equilibrium condition for prices $\boldsymbol{p}^\circ$, then in our comparison we select the one (which we will call $\bar{\boldsymbol{q}}^\circ$) that achieves the highest revenue in order to give the method of ignoring network effects the "benefit of doubt". We denote the corresponding revenue of this price by $\pi^\circ = \sum_{i=1}^n p_i^\circ \bar{q}_i^\circ$.

Note that the price vector in Case 3 is the optimal price vector for the classical

MNL model without network effects (which is a uniform price across all products in this homogeneous case). We summarize the results in Table 2.2. In the table, ℓ^u represents the percentage of the revenue $\pi(\boldsymbol{q}^*)$ that is lost if we use prices $\boldsymbol{p}^u$. Similarly, ℓ° is the percentage lost if we use prices $\boldsymbol{p}^\circ$.

α	0	1	2	3	4	5
$\pi(\boldsymbol{q}^*)$	1.9445	2.0343	2.1295	2.4413	3.1209	3.8807
π^u	1.9445	2.0343	2.1295	2.2299	2.3353	2.4456
ℓ^u	0%	0%	0%	8.66%	25.17%	36.98%
π°	1.9445	2.0336	2.1255	2.2185	2.3103	2.3984
ℓ°	0%	0.04%	0.19%	9.12%	25.97%	38.20%

Table 2.2: Comparison of revenues for different pricing decisions.

From Table 2.2, we see that when $\alpha \leq 2$, the optimal pricing strategy is to set uniform prices across all products; therefore, the revenue in Case 2 is identical to that in Case 1. When $\alpha \geq 3$, a uniform pricing strategy is no longer optimal and Case 1 generates a higher revenue than does Case 2. Moreover, for strictly positive α, the revenue in Case 3 is smaller than the optimal revenue, and the difference becomes larger when α increases.

Next we do the same test with heterogeneous parameters. In the following experiments, we consider problems with n ranging from 2 to 20. For each n, we generate 1000 different sets of parameters to obtain 1000 randomly generated problem instances. We apply Algorithm 1 to solve each problem instance. To generate an instance, we sample α_i, y_i and γ_i from uniform distributions over $[0, 4]$, $[1, 2]$ and $[0, 2]$, respectively. We consider only Cases 1 and 3 because a uniform pricing strategy is not optimal — even when there are no network effects — in this heterogeneous setting. The results are summarized in Table 2.3, where $\bar{\ell}^\circ$ denotes the percentage loss from ignoring network effects, averaged over the

1000 instances.

n	2	3	5	10	15	20
$\bar{T}$ (secs.)	4.00	4.23	4.68	5.24	5.89	6.34
ℓ°	10.31%	12.30%	13.03%	13.82%	12.45%	11.47%

Table 2.3: Comparison of revenues in heterogeneous settings

In Table 2.3, we can observe that the loss from ignoring network effects is considerable for each n. In the row labeled $\bar{T}$, we show the average run-time (in seconds) of Algorithm 1 on a Mac desktop computer with a 2.7 GHz Intel Quad-Core i5 processor and 8 GB of memory. We see that Algorithm 1 can be carried out very quickly, even with large values of n. As n increases from 2 to 20, the average running time experiences a modest increase from about four seconds to about six seconds.

2.6.3 Robustness of the Solution

In this section, we test the robustness of the solution to (P1) with respect to estimation errors of the network strength parameters. For simplicity, we confine our attention to problems with homogeneous products. Suppose that the seller believes the network strength parameter is $\tilde{\alpha}$ (this may be viewed as the seller's estimate) and solves (P1) with $\tilde{\alpha}$ in place of α. This yields solution $\tilde{\boldsymbol{q}}$. Then, using (2.7) with $\tilde{\alpha}$ in place of α, the seller obtains prices $\tilde{\boldsymbol{p}}(\tilde{\boldsymbol{q}})$. What happens if the seller implements those prices when the network strength parameter upon which customers base their purchase decisions is actually α rather than $\tilde{\alpha}$? (Note that by the discussion after (2.3), our test can also be viewed as a test for robustness of the solution with respect to the estimation errors of the market size M.)

To test this, we fix $n = 5$, $y = 2$, and $\gamma = 1$. In this setting, the threshold from Theorem 2.3.1 at which the actual optimal policy switches from uniform pricing to two different prices is $\hat{\alpha} = 2.513$. We consider two different true values for the network strength parameter, $\alpha = 2.4$ and $\alpha = 2.6$. Note that one of these is below the threshold and one is above. In these cases, there is no issue with multiple equilibria at prices $\boldsymbol{p}^*$ or $\tilde{\boldsymbol{p}}(\tilde{\boldsymbol{q}})$. Let $\boldsymbol{q}'$ denote the resulting equilibrium sales vector when the seller implements prices $\tilde{\boldsymbol{p}}(\tilde{\boldsymbol{q}})$ in the problem that has actual network parameter α. Note that $\boldsymbol{q}' \neq \tilde{\boldsymbol{q}}$ when $\tilde{\alpha} \neq \alpha$. On the other hand, if $\tilde{\alpha} = \alpha$, then the seller does not make an estimation error and therefore $\boldsymbol{q}' = \tilde{\boldsymbol{q}}$. Let $\tilde{\pi} = \sum_{i=1}^{5} \tilde{p}_i(\tilde{\boldsymbol{q}})q_i'$ be seller's revenue from implementing prices $\tilde{\boldsymbol{p}}(\tilde{\boldsymbol{q}})$ and let $\tilde{\ell}$ be the percentage revenue loss in comparison to the optimal solution (without estimation errors); i.e., $\tilde{\ell} = 100 \times (\pi(\boldsymbol{q}^*) - \tilde{\pi})/\pi(\boldsymbol{q}^*)$. Results are summarized in Table 2.4.

$\tilde{\alpha}$		2	2.2	2.4	2.6	2.8	3
$\tilde{\boldsymbol{p}}(\tilde{\boldsymbol{q}})$	p^H	3.0496	3.0614	3.0645	4.4127	4.8514	5.2434
	p^L	3.0496	3.0614	3.0645	2.6694	2.7966	2.9269
$\alpha = 2.4$	$\tilde{\pi}$	2.1689	2.1690	2.1690	2.1226	2.0779	1.9943
	$\tilde{\ell}$	0.01%	$<$ 0.01%	0%	2.14%	4.20%	8.05%
$\alpha = 2.6$	$\tilde{\pi}$	2.1886	2.1889	2.1889	2.2199	2.2046	2.1505
	$\tilde{\ell}$	1.41%	1.40%	1.39%	0%	0.69%	3.13%

Table 2.4: Sensitivity to errors in estimates of α

As the table shows, the optimal strategy is to price equally if the true network parameter is $\alpha = 2.4$. If α is underestimated (columns to the left of the column labeled 2.4 in the row labeled 2.4), then the strategy is still to price equally, and the revenue loss from underestimating α is negligible. This occurs because the optimal price is quite insensitive to α in this case, as can be seen (for different examples) in Figure 2.3, which shows prices as "nearly" constant in the region where a single price is optimal. When α is overestimated (columns to the right of the column labeled 2.4 in the row labeled 2.4), then the resulting strategy is to price differently and the revenue loss is comparatively more sensitive to the estimation error. We should not be surprised to see a relatively large (roughly 8%) loss in the column for $\tilde{\alpha} = 3$ because in that case the estimate of the network strength parameter differs from the true value of $\alpha = 2.4$ by 25%. If the true network parameter is $\alpha = 2.6$, then it is above the threshold $\hat{\alpha}$, and therefore the optimal strategy has two distinct prices. When the estimate $\tilde{\alpha}$ is below $\hat{\alpha}$, the revenue loss is quite insensitive to $\tilde{\alpha}$ due to the insensitivity of the optimal price in that range. When α is overestimated, the revenue loss is also low as long as the estimation error is reasonable.

2.7 Concluding Remarks

In this chapter, we considered pricing problems faced by a seller whose products exhibit network effects. For a setting with a homogeneous assortment of products, we established that the optimal solution has the following form: if the network effects are weak, then the seller should set the same price for all products; if the network effects are strong, then the seller should boost the sales of a single product by setting its price low and setting the prices of all other products at a single higher value. We also provided comparative statics and extended our results to nonhomogeneous settings. In view of the particularly clean structure that arises from our model, we are optimistic that the results in this chapter can serve as a building block for much future research.

There are, of course, limitations to our model. One simple example is our underlying assumption that customers each buy at most one product. As a practical matter, it is natural to expect that some customers may wish to buy more than one product. A possibly more complex issue is that in our current formulation, purchase probabilities satisfy the equilibrium condition (2.4), which means that customers have rational expectations about market outcomes. However, it may be that customers behave in a more myopic fashion. For example, for a single instance of the pricing problem, sales may build dynamically over time and customers may make purchase decisions based on current cumulative sales levels of the products. Such sequential myopic decision making may lead to an outcome that does not satisfy (2.4). This also in turn raises the possibility of the seller responding to such behavior by adjusting prices over time to gain higher revenue. These issues may be the topic of future work.

Chapter 3

Multiplicity of Equilibria and Worst-Case Pricing

3.1 Introduction

In last chapter, we showed that multiple equilibria could arise in the multinomial logit model with network effects. More precisely, for a fixed pricing decision, there could be multiple sales quantities that satisfy the equilibrium condition. In the previous discussions, when such a situation happens, we optimistically assumed that the sales quantity that leads to the highest (best-case) revenue would prevail. Yet in practice, this may not necessarily be the case.

To address this issue of multiple equilibria, in this chapter, we study the pricing problem from a robust perspective. Specifically, if there are multiple equilibria of sales quantities to a pricing decision, we take a conservative attitude and assume that the resulting equilibrium is the one that leads to the lowest (worst-case) revenue. We call this the “worst-case problem” and the optimal solution the pessimistic solution.

We first consider the single-product case. We show that if the network effect is weak, there is no issue of such multiple equilibria and the optimistic solution and the pessimistic solution are equivalent; if the network effect is strong, the multiple equilibria issue arises but we can still solve the worst-case problem efficiently. Particularly, we show that the optimal price in the pessimistic case would be lower than that in the optimistic case, which means that the firm should be more conservative with its pricing policy. We also show that as the network effect increases, both product sales and revenue will increase accordingly. In addition, the potential loss from assuming optimism on the realized sales could be significant if the actual sales quantity is the pessimistic one, especially when the network effect is strong.

We next study the two-product case. When the firm has two products to sell, there are many possibilities for the total number of equilibria for a fixed pair of prices. The equilibrium condition becomes much more complicated and it is difficult to derive analytical solutions to the worst-case problem. To numerically solve this problem, we propose two methods to compute the pessimistic optimal solution: divide and search, and linear relaxation. Numerical experiments suggest that at the pessimistic optimal prices, there is always only one equilibrium of sales quantities. In addition, we find that when the network effect is strong, the revenue from offering just one product is almost as good as that from selling a portfolio of multiple products. This suggests that a company selling multiple products with strong network effects may be wise to simply offer a single product.

3.2 The General Model

The model settings are the same as in Chapter 2. A single seller offers n products, labeled as $i \in \mathcal{N} = \{1, \cdots, n\}$, to the customers. Each individual customer in the market chooses at most one product among them, or could also choose not to buy, in which case we label this "no purchase" option as product 0. The utility a customer obtains from purchasing product i is

$$u_i = v_i + \epsilon_i, \tag{3.1}$$

where v_i is the expected utility from purchasing product i and ϵ_i is a random variable that represents customer-specific idiosyncracies. For the no purchase option, without loss of generality, we let $v_0 = 0$. As in the standard MNL model, we assume $\epsilon_0, \epsilon_1, \cdots, \epsilon_n$ are i.i.d. Gumbel random variables. By incorporating the network effects, each v_i includes a network effect term that reflects an individual's additional utility from others' usage of product i. More precisely, for $i \in \mathcal{N}$, we have

$$v_i = y_i - p_i + \alpha_i q_i \tag{3.2}$$

where $y_i \geq 0$ is the intrinsic utility of product i, p_i is the price, α_i is the network effect sensitivity parameter, and q_i is the sales quantity. The parameter α_i represents the strength of network effects for product i, and a higher value of α_i indicates that the consumers of product i are more sensitive to others' usage of that product. As in Chapter 2, we consider a fluid model of "infinitesimal" customers with the market size normalized to 1. Consequently, the sales quantity q_i can also be interpreted as the probability that an individual customer will purchase product i.

By standard results for the MNL model, if we are given sales quantities $\boldsymbol{q} = (q_1, \ldots, q_n)$ and prices $\boldsymbol{p} = (p_1, \ldots, p_n)$, then the probability that a customer purchases product $i \in \mathcal{N}$ is

$$q_i = P(u_i = \max_{j\in\{0,1,\ldots,n\}} u_j) = \frac{\exp(y_i - p_i + \alpha_i q_i)}{1 + \sum_{j=1}^{n} \exp(y_j - p_j + \alpha_j q_j)}. \tag{3.3}$$

Upon defining

$$F_i(\boldsymbol{p}, \boldsymbol{q}) = \frac{\exp(y_i - p_i + \alpha_i q_i)}{1 + \sum_{j=1}^{n} \exp(y_j - p_j + \alpha_j q_j)}, \tag{3.4}$$

we may re-write (3.3) as

$$q_i = F_i(\boldsymbol{p}, \boldsymbol{q}) \text{ for all } i \in \mathcal{N}. \tag{3.5}$$

The goal of the seller is to select prices that maximize its total revenue $\pi(\boldsymbol{p}, \boldsymbol{q}) = \boldsymbol{p}^T \boldsymbol{q}$ subject to the equilibrium constraints (3.3). The seller implements prices $\boldsymbol{p}$, and the market responds with quantities $\boldsymbol{q}$ that satisfy (3.3). For a given vector of prices $\boldsymbol{p}$, there may exist multiple sales vectors that satisfy (3.3) and these sales vectors are associated with different revenues. In Chapter 2, we implicitly took an optimistic attitude that for any prices $\boldsymbol{p}$, the sales vector that arises is the one with the highest revenue among those satisfying (3.3). Here, we seek to better understand the effects of non-unique sales equilibria by studying the "worst-case" problem assuming that for any prices $\boldsymbol{p}$, the sales vector that arises is the one with the lowest revenue among those satisfying (3.3).

Define $F(\boldsymbol{p}, \boldsymbol{q}) = (F_1(\boldsymbol{p}, \boldsymbol{q}), \ldots, F_n(\boldsymbol{p}, \boldsymbol{q}))$. For any vector of prices $\boldsymbol{p}$, define $Q(\boldsymbol{p})$ to be the set containing all $\boldsymbol{q}$ that satisfy (3.5), that is, $Q(\boldsymbol{p}) = \{\boldsymbol{q} \in [0,1]^n : \boldsymbol{q} = F(\boldsymbol{p}, \boldsymbol{q})\}$. We can now formulate the best-case and worst-case pricing problems.

The best-case pricing problem is

$$\begin{aligned}\overline{\pi} &= \sup_{\boldsymbol{p}} \overline{\pi}(\boldsymbol{p}) \\ \overline{\pi}(\boldsymbol{p}) &= \max_{\boldsymbol{q}} \left\{\pi(\boldsymbol{p},\boldsymbol{q}) : \boldsymbol{q} \in Q(\boldsymbol{p})\right\} .\end{aligned} \tag{BC}$$

Likewise, the worst-case pricing problem is

$$\begin{aligned}\underline{\pi} &= \sup_{\boldsymbol{p}} \underline{\pi}(\boldsymbol{p}) \\ \underline{\pi}(\boldsymbol{p}) &= \min_{\boldsymbol{q}} \left\{\pi(\boldsymbol{p},\boldsymbol{q}) : \boldsymbol{q} \in Q(\boldsymbol{p})\right\} .\end{aligned} \tag{WC}$$

Lemma **??** in the appendix establishes that $Q(\boldsymbol{p})$ is a compact set. Moreover, for given $\boldsymbol{p}$, the function $\pi(\boldsymbol{p},\boldsymbol{q})$ is continuous in $\boldsymbol{q}$ and therefore $\overline{\pi}(\boldsymbol{p})$ and $\underline{\pi}(\boldsymbol{p})$ both exist. As we will see later, $\underline{\pi}(\boldsymbol{p})$ may be discontinuous in $\boldsymbol{p}$, and therefore there may be no optimal solution to $\sup_{\boldsymbol{p}} \underline{\pi}(\boldsymbol{p})$ in (WC). In such cases, we will settle on an ϵ-optimal solution $\boldsymbol{p}^{\epsilon}$ wherein $\underline{\pi}(\boldsymbol{p}^{\epsilon}) > \sup_{\boldsymbol{p}} \underline{\pi}(\boldsymbol{p}) - \epsilon$. Such an ϵ-optimal solution exists for any $\epsilon > 0$.

3.3 Single-Product Case

In this section, we consider the scenario that there is only one product. The problem is to set a price to maximize the total revenue from selling this product. Throughout this section, we will drop the subscripts from the notations because such subscripts are unnecessary when there is only one product. As in (3.2), the expected utility a customer obtains from purchasing the product for a given price p is $u = y - p + \alpha$, and the sales quantity q in equilibrium satisfies

$$q = F(p,q) := \frac{\exp(y - p + \alpha q)}{1 + \exp(y - p + \alpha q)}. \tag{3.6}$$

We begin by describing an approach from Du et al. (2016) to solve the best-case version of this problem. This is the approach from Chapter 2, specialized to a setting with one product ($n = 1$). By transforming (3.6), price p can be written as

$$p(q) = y + \alpha q - \log q + \log(1 - q). \tag{3.7}$$

For any given $q \in (0, 1)$, simple algebra shows that $p = p(q)$ is the unique price for which (3.6) holds. [Note, however, this does not preclude the possibility that there is some other value of sales (say q') such that $p(q)$ and q' also together satisfy (3.6).] For $n = 1$, if $q \in (0, 1)$ then we have that $q \in Q(p)$ if and only if $p(q) = p$.

Define

$$\widetilde{\pi}(q) = p(q)q = yq + \alpha q^2 - q\log(q) + q\log(1 - q), \tag{3.8}$$

and consider the maximization problem

$$\widetilde{\pi}^* \quad = \quad \max_q \{\widetilde{\pi}(q) : 0 < q < 1\}\ . \tag{P0}$$

We now have that

$$\overline{\pi} = \sup_p \max_q \{pq : q \in Q(p)\} = \sup_p \max_{q\in(0,1)} \{pq : p(q) = p\} = \max_{q\in(0,1)} \{p(q)q\} = \widetilde{\pi}^*$$

and we can summarize our developments so far with the following result.

Lemma 3.3.1. *For $n = 1$, the best-case problem* (BC) *is equivalent to* (P0)*; i.e., $\overline{\pi} = \widetilde{\pi}^*$ and a pair (p, q) is an optimal solution to* (BC) *if and only if q is an optimal solution to* (P0) *and $p(q) = p$.*

Given an optimal solution $(\overline{p}, \overline{q})$ to the best-case problem (BC), what might happen if the seller implements the price $\overline{p}$? In particular, what value of sales quantity other than $\overline{q}$ might arise? In Figure 3.1, we present an illustrative example. Figure 3.1 plots the function $p(q)$ defined in (3.7). Note that the (p, q)-pairs that satisfy (3.6) are simply the points in two-dimensional space that fall on the plot of $p(q)$. Therefore, we can determine the number of equilibria for a given price p by counting the number of times a horizontal line at p intersects the curve of $p(q)$. Figure 3.1 shows that if the price p is such that $p \in (p^L, p^H)$, then there are three different q satisfying (3.7). Therefore, if $\overline{p}$ falls into this range, then the sales could have three possible values. By imposing an optimistic attitude, we assume that the actual sales level q is the largest one among these three values. If, contrary to the optimistic assumption, the actual sales level q turns out to be the smallest value (which is consistent with the pessimistic assumption), then the sales — as well as the revenue — would be much lower than that in the optimistic assumption. For example, in Figure 3.1(b), when the implemented price is $p' = 4$, then the largest sales quantity in equilibrium would be around 0.929, while the smallest sales would be slightly above 0.07. Thus there is a drastic difference between these two realizations of sales quantities. The following lemma describes the structure of the function $p(q)$ and is essential for understanding when multiple sales equilibria do and do not arise.

Lemma 3.3.2. *For $n = 1$, the equilibrium condition has the following properties.*

1) Suppose $\alpha \leq 4$. The function $p(q)$ defined in (3.7) *is a strictly decreasing function. For each $p > 0$, there is a unique q such that (p, q) satisfies* (3.6).

2) Suppose $\alpha > 4$. The function $p(q)$ has a unique local minimum and a unique local maximum at respectively, $q^L = 1/2 - \sqrt{1/4 - 1/\alpha}$ and $q^H = 1/2 + \sqrt{1/4 - 1/\alpha}$. In addition, for $p^L := p(q^L)$ and $p^H := p(q^H)$ we have the following:

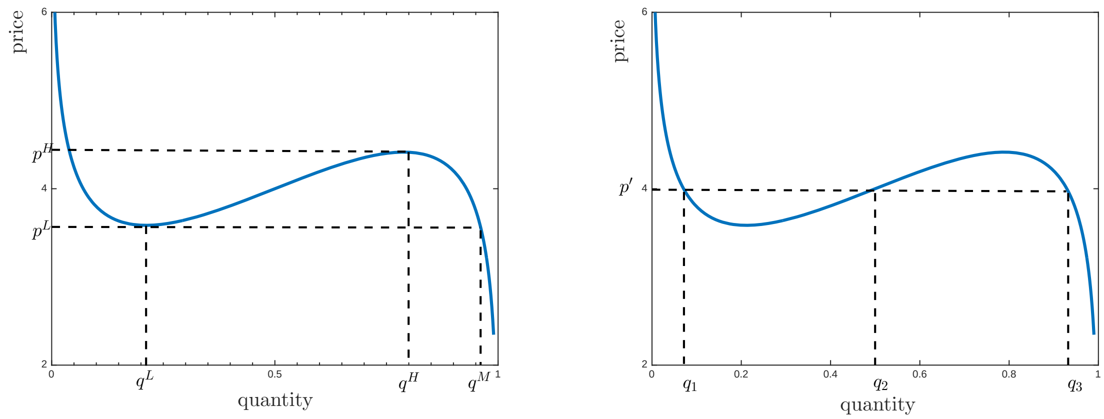

(a) Plot of function $p(q)$ for parameter values $\alpha = 6$, $y = 1$.

(b) An example of three equilibria

Figure 3.1: Plot of function $p(q)$

(a) For each $p \in (p^L, p^H)$, *there are three distinct* q *such that* (p, q) *satisfies* (3.6)*;*

(b) for each $p \in \{p^L, p^H\}$, *there are two distinct* q *such that* (p, q) *satisfies* (3.6)*;*

(c) for each $p \notin [p^L, p^H]$, *there is a unique* q *such that* (p, q) *satisfies* (3.6).

Below, we use q^M to denote the larger one of the two sales quantities for which (p^L, q) satisfies (3.6). We have $q^M > q^L$ and $p(q^M) = p^L$(see Figure 3.1).

Lemma 3.3.2 indicates that, in single product scenario the issue of multiple equilibria arises only when the network effect is strong enough ($\alpha > 4$). When the product exhibits weak network effects, the problems (BC) and (WC) are equivalent, because there is a unique sales equilibrium to every price. Because of this, in the remainder of this section, we focus on the setting with strong network effects ($\alpha > 4$). Here, it is obvious that if the optimal price (resp., sales quantity) in (BC) is less than p^L (resp., greater than q^M), then (BC) and (WC) have the same optimal solution. To further understand the relationship between (BC) and (WC), we present the following result.

Lemma 3.3.3. *Suppose $\alpha > 4$. Then $\widetilde{\pi}(q)$ is a unimodal function and has a unique local and global maximum point on $\widetilde{q} \in (0,1)$. The point $\widetilde{q}$ satisfies the first order condition*

$$\widetilde{\pi}'(q) = y + 2\alpha q - \log(q) + \log(1-q) - \frac{1}{1-q} = 0. \tag{3.9}$$

In addition, $\widetilde{q} > q^H$ and $\widetilde{q}$ increases in α.

By Lemma 3.3.3, there is a unique solution $\widetilde{q}$ to (P0) and therefore (BC) can be computed efficiently using a standard root-finding algorithm, for example, the Newton's method. Consider now the following optimization problem

$$\widetilde{\pi}^\circ \quad = \quad \max_q \{\widetilde{\pi}(q) : q^M \leq q < 1\}, \tag{P1}$$

which can be viewed as a tighter version of (P0).

Define the optimal solution to (P1) as q°, and the corresponding price as $p^\circ = p(q^\circ)$. Similarly, there is a unique optimal solution to (P1), because $\widetilde{\pi}(q)$ is unimodal. In addition, we have the following result.

Lemma 3.3.4. *Suppose $\alpha > 4$. For $n = 1$, the worst-case problem* (WC) *can be solved using* (P1)*; i.e., $\underline{\pi} = \widetilde{\pi}^\circ$. In addition:*

1) *If $q^\circ > q^M$, then there exists an optimal solution $(\underline{p}, \underline{q})$ to* (WC) *and $\underline{p} = p^\circ$ and $\underline{q} = q^\circ$;*

2) If $q^\circ = q^M$*, then there is no optimal solution to* (WC)*, and for any* $\epsilon > 0$*,* $(\underline{p}^\epsilon, \underline{q}^\epsilon) := (p^\circ - \epsilon, q^\circ + \delta(\epsilon)) = (p^L - \epsilon, q^M + \delta(\epsilon))$ *is an* ϵ*-optimal solution to* (WC)*, where* $\delta(\epsilon)$ *is the unique solution to* $p(q^M + \delta) = p^L - \epsilon$*.*

The preceding results reveal, via (P0) and (P1), a simple relationship between (BC) and (WC): solving for the pessimistic case is almost equivalent to solving for the optimistic case on a restricted domain. Lemma 3.3.4 also implies that the optimal sales quantity $\underline{q}$ must be higher than q^M, therefore the optimal price p° must be lower than p^L. Therefore, an optimal worst-case quantity can lie only in the region $q \in [q^M, 1)$. Lemma 3.3.4 also states that even if the worst case problem (WC) has no optimal solution, the seller should set the price at a level just below p°. To simplify our notations, we let $\underline{p} = p^\circ$ from now on.

Based on Lemmas 3.3.3 and 3.3.4, the following theorem compares the solutions to the two settings based on different grounds.

Theorem 3.3.5. *The optimal solution* $\overline{p}$ *to Problem* (BC) *and the optimal solution* $\underline{p}$ *to Problem* (WC) *satisfy:*

1) If $\widetilde{q} > q^M$*, then* $\overline{q} = \underline{q}$*,* $\overline{p} = \underline{p}$*, and* $\overline{\pi} = \underline{\pi}$*;*

2) If $\widetilde{q} \leq q^M$*, (WC) has no optimal solution but for any* $\epsilon > 0$*,* $(\underline{p}^\epsilon, \underline{q}^\epsilon)$ *is an* ϵ*-optimal solution to* (WC)*, and it satisfies* $\overline{q} < \underline{q}^\epsilon$*,* $\overline{p} > \underline{p}^\epsilon$*, and* $\overline{\pi} > \underline{p}^\epsilon \underline{q}^\epsilon$*.*

Here $\widetilde{q}$ *is the unique maximum point of* $\widetilde{\pi}(q)$*, and* q^M *is the larger solution to* $p(q^M) = p^L$*.*

Theorem 3.3.5 gives an explicit relationship between the optimal price decisions based on two beliefs: the optimistic price should always be higher than the pessimistic price. When $\widetilde{q} > q^M$, they have the same price decision. Therefore, we do not need to be concerned about the issue of multiple equilibria. However, when $\widetilde{q} \le q^M$, the firm should implement a price lower than the optimistic optimal price to eliminate the risk that a sales equilibrium arises at a value lower than expected. The following algorithm solves Problem (WC).

Algorithm 3. *1. Solve for $\widetilde{q}$ by solving the equation $\widetilde{\pi}'(q) = y + 2\alpha q - \log(q) + \log(1-q) - \frac{1}{1-q} = 0$ over $[1/2 + \sqrt{1/4 - 1/\alpha}, 1]$ through bisection method.*

2. Solve for q^M by solving the equation $p(q^L) = y + \alpha q - \log(q) + \log(1-q)$ over $[1/2 + \sqrt{1/4 - 1/\alpha}, 1]$ through bisection method.

3. Compare $\widetilde{q}$ with q^M to obtain $\underline{q} = \max(\widetilde{q}, q^M)$.

Theorem 3.3.6. *$\underline{q}$ and $\underline{\pi}$ increase in α.*

Theorem 3.3.6 says that as network effects become stronger, the seller should sell more products to achieve a higher revenue. This property is similar to the optimistic case, therefore the seller should always boost sales when it observes a stronger network effect, no matter which belief he holds.

From the view of the seller, it is useful to know the loss they would suffer from implementing the "optimal" optimistic price when the realized sales quantity is actually the pessimistic one. More specifically, the seller would like to know that in the pessimistic condition, what is the difference in revenue from using the optimistic optimal price $\overline{p}$ to using the pessimistic optimal price $\underline{p}$. By Theorem 3.3.5, these two revenues differ only when $\widetilde{q} < q^M$. Even though $\underline{p}$ does not exist, a pessimistic seller can adopt the price $p^L - \epsilon$. This is an ϵ-optimal solution that always leads to a unique revenue. In contrast, an optimistic seller implementing price $\overline{p} = p(\overline{q})$ could lead to three different sales quantities according to our model, and the realized one may not be the highest one. Suppose the smallest sales quantity corresponding to price $\overline{p}$ is denoted as q^S. We would like to compare the revenue $\pi^S = \overline{p}q^S$ with $\underline{\pi} = p^L q^M$. $\underline{\pi}$ is the total revenue a pessimistic seller would earn when the realized sales is indeed in the worst case, while π^S is the total revenue an optimistic seller would earn when the realized sales counters its wish. We have the following results about π^S and $\underline{\pi}$.

Theorem 3.3.7. *1) If $\widetilde{q} > q^M$, then $\pi^S = \underline{\pi}$.*

2) If $\widetilde{q} \leq q^M$, then $\frac{\pi^S}{\underline{\pi}} \leq \frac{q^L}{q^H} = \frac{1/2-\sqrt{1/4-1/\alpha}}{1/2+\sqrt{1/4-1/\alpha}}$. Furthermore, $\underline{\pi} - \pi^S \geq p^L(q^M - q^L)$, and $p^L(q^M - q^L)$ increases in α.

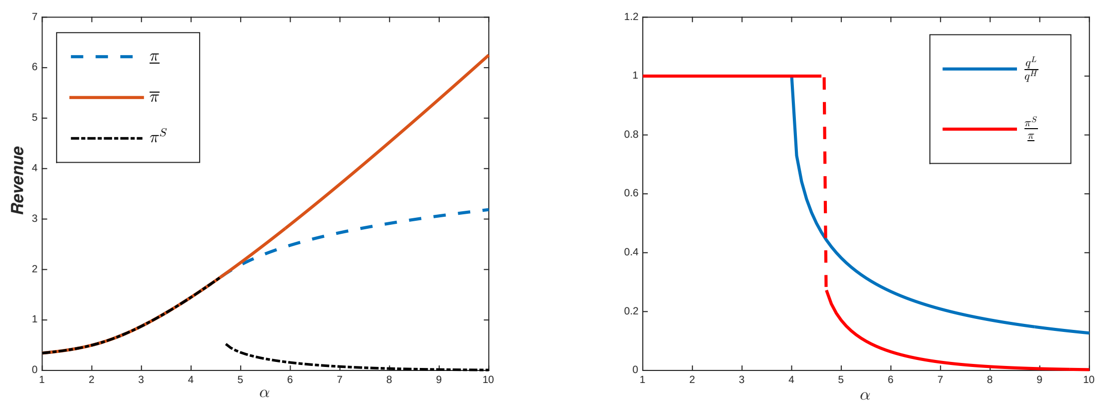

(a) Comparison of $\underline{\pi}$, $\overline{\pi}$, and π^S

(b) Revenue ratio comparison

Figure 3.2: Illustration of Theorem 3.3.7

From Theorem 3.3.7, the potential loss from assuming optimism on the realized sales could be significant, especially in the presence of a strong network effect. In Figure 3.2, it can be observed that there exists a discontinuity in π^S as well as in $\frac{\pi^S}{\underline{\pi}}$, when α exceeds a threshold value. This is due to the multiple equilibria from implementing the optimal price $\overline{p}$. We also observe that the issue of multiple equilibria aggravates the loss of revenue for higher values of α. Therefore, a company selling a product with stronger network effect has more incentives to take a pessimistic attitude and implement a conservative pricing policy.

3.4 Two-Product Case

In this section, we study the scenario that a seller offers two products with network effects to the market. The two products have intrinsic utilities y_1, y_2 with network effect sensitivities α_1, α_2. For simplicity, we assume $y_1 = y_2 = 0$. Suppose the prices of those two products are set to $\boldsymbol{p} = (p_1, p_2)$, the sales quantities must satisfy

$$q_1 = \frac{e^{-p_1+\alpha_1 q_1}}{1 + e^{-p_1+\alpha_1 q_1} + e^{-p_2+\alpha_2 q_2}}, \qquad q_2 = \frac{e^{-p_2+\alpha_2 q_2}}{1 + e^{-p_1+\alpha_1 q_1} + e^{-p_2+\alpha_2 q_2}}.$$

The above equations follow directly from (3.3). For the convenience of analysis, we rewrite the two equations into a system of two closed-form functions in terms of q_1, q_2

$$q_1 = 1 - q_2 - q_2 e^{p_2 - \alpha_2 q_2}, \qquad q_2 = 1 - q_1 - q_1 e^{p_1 - \alpha_1 q_1}. \tag{3.10}$$

Define $f_i(q_i|p_i, \alpha_i) := 1 - q_i - q_i e^{p_i - \alpha_i q_i}$, the two equations in (3.10) can be represented as $q_2 = f_1(q_1|p_1, \alpha_1)$ and $q_1 = f_2(q_2|p_2, \alpha_2)$. The structure of f_1, f_2 shows that the sales quantities q_1, q_2 are determined not only by pricing decisions (p_1, p_2), but also by themselves.

With the above knowledge, the problem is to find the pessimistic optimal pricing decision $\underline{\boldsymbol{p}}$ that maximizes the total revenues $\underline{\pi}(\boldsymbol{p})$ in the worst-case

$$\underline{\pi}(\boldsymbol{p}) \quad = \quad \min_{\boldsymbol{q}} \left\{ \pi(\boldsymbol{p}, \boldsymbol{q}) = p_1 q_1 + p_2 q_2 | (q_1, q_2) = \boldsymbol{q} \in Q(\boldsymbol{p}) \right\},$$

where $Q(\boldsymbol{p})$ is the set of all $\boldsymbol{q}$ satisfying (3.10).

The two-product (WC) problem is not as easily solvable as the one-product (WC) problem. In the one-product scenario, the equilibrium condition (3.7) is a one-variable function $p(q)$ with a simpler analytic structure, which allows us to easily identify the worst-case sales q for any price p. In addition, the unimodularity property of function $\widetilde{\pi}(q)$ enables us to apply a standard root-finding algorithm on q to find the optimistic optimal solution $\widetilde{q}$, and then we can solve the pessimistic problem by comparing $\tilde{q}$ with q^M. However, in the two-product scenario, the equilibrium condition consists of two functions $f_1(\cdot), f_2(\cdot)$ with two variables q_1, q_2. The extra equation and variable impose more complexity to the system, and the structure of the worst-case equilibrium is not straightforward to characterize. In fact, even if prices $\boldsymbol{p} = (p_1, p_2)$ are known, the total number of equilibria may not be easy to analyze. For example, Figure 3.3 plots the equation system $q_2 = f_1(q_1|p_1, \alpha_1), q_1 = f_2(q_2|p_2, \alpha_2)$ where $p_1 = 4, p_2 = 2.8, \alpha_1 = 14, \alpha_2 = 8$. The equilibria are the points in two-dimensional space where the two curves (the solid lines) intersect. (The dashed lines in the figure will be explained later.) In this particular example, there exist a total number of 5 equilibria. In general, there does not appear to be a straightforward way to compute all equilibria. Thus in order to find the pessimistic optimal prices, we have to search over all the price combinations.

In the following discussions, we propose two different directions for finding

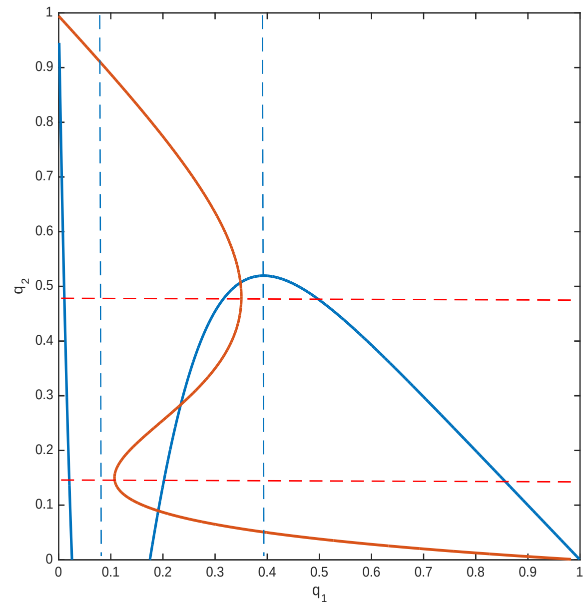

Figure 3.3: Example with $p_1 = 4, p_2 = 2.8, \alpha_1 = 14, \alpha_2 = 8$

the pessimistic solution $\underline{\boldsymbol{q}}$ for a given pair of prices $\boldsymbol{p} = (p_1, p_2)$. The first idea is to divide the feasible two-dimensional domain into smaller regions where the equilibrium functions (3.10) in each region remain monotone. The monotonicity structure enables us to apply bisection algorithm in each region separately and then find all the equilibria. Another method we propose is to formulate a linear relaxation problem for the original one and find a lower bound of the pessimistic solution $\underline{\boldsymbol{q}}$. We elaborate the two ideas in the next two subsections.

3.4.1 Divide and Search

In this subsection, we discuss the idea of dividing and searching for all the equilibria to (3.10) for a given pair of prices $\boldsymbol{p} = (p_1, p_2)$. We take advantage of the structure of the equilibrium functions, and divide the feasible domain into smaller regions. In each region, we exploit the properties of functions therein and apply appropriate algorithms to compute the equilibria contained in that region.

First, we explain the basic reasoning of dividing the two-dimension domain. Note that all the feasible equilibria in $Q(\boldsymbol{p})$ are contained in $[0,1]^2$. For a sub-region $\mathcal{R} \subseteq [0,1]^2$, define $\underline{\pi}^{\mathcal{R}}(\boldsymbol{p})$ as the worst-case revenue from implementing $\boldsymbol{p}$ with sales quantities restricted in $\mathcal{R}$, i.e.

$$\underline{\pi}^{\mathcal{R}}(\boldsymbol{p}) \quad = \quad \min_{\boldsymbol{q}} \left\{ \pi(\boldsymbol{p}, \boldsymbol{q}) : \boldsymbol{q} \in \{Q(\boldsymbol{p}) \cap \mathcal{R}\} \right\} . \tag{3.11}$$

Then we have

Proposition 3.4.1. *If regions $\mathcal{R}_1, \cdots, \mathcal{R}_N \subseteq [0,1]^2$ are such that $\bigcup_{i=1,\cdots,N} \mathcal{R}_i = [0,1]^2$, then*

$$\underline{\pi}(\boldsymbol{p}) = \min_{i=1,\cdots,N} \underline{\pi}^{\mathcal{R}_i}(\boldsymbol{p}).$$

Proposition 3.4.1 says that the worst equilibrium in $[0,1]^2$ can be found by comparing the worst ones within each sub-region. Then a question remains on how to divide the feasible region in a useful way. A natural idea is to divide it such that in functions (3.10) the equilibrium conditions are monotone within each region. The following lemma provides a guidance to do that.

Lemma 3.4.2. *The function $f(q) = 1 - q - qe^{p-\alpha q}$ satisfies the following properties:*

1) $f(q)$ is convex on $q \in (0, 2/\alpha)$, and concave on $q \in (2/\alpha, 1)$;

2) $f'(q) = 0$ has at most two solutions;

3) $f(0) = 1, f(1) < 0$ and $f(q) = 0$ has at most three solutions.

By Lemma 3.4.2, $f_i'(q) = 0$ has at most two solutions, therefore q_i can be separated into three monotone regions by the two solutions of $f_i'(q) = 0$. There are two variables q_1, q_2 in the two-product case and $f_1(q_1), f_2(q_2)$ are cut along the q_1, q_2 axes respectively. Therefore, there are at most 9 regions to study. The dashed lines in Figure 3.3 illustrates an example of dividing the domain $[0,1]^2$ into 9 regions. (Note that $f'(q) = 0$ may have only one or even no solution by Lemma 3.4.2, and in that case we separate it into two regions at the point $f''(q) = 0$.) In each region, $f_1(\cdot), f_2(\cdot)$ are monotone functions and this structure enables us to apply efficient algorithms on them to quickly find all the equilibria. Define $g(q) = f_1(q) - f_2^{-1}(q)$, then any equilibrium must satisfy $g(q) = 0$. The following proposition elaborates the results.

Proposition 3.4.3. *For $i = 1, 2$, suppose $f_i'(q_i) = 0$ has two different solutions q_i^a, q_i^b and $q_i^a < q_i^b$. The feasible domain $q_1 \in [0,1]$ is divided into three intervals $\mathcal{I}_1 = [0, q_1^a], \mathcal{I}_2 = (q_1^a, q_1^b]$ and $\mathcal{I}_3 = (q_1^b, 1]$, and $q_2 \in [0,1]$ is divided into three intervals $\mathcal{J}_1 = [0, q_2^a], \mathcal{J}_2 = (q_2^a, q_2^b]$ and $\mathcal{J}_3 = (q_2^b, 1]$. Therefore the domain of $\boldsymbol{q} \in [0,1]^2$ can be divided into 9 regions. Table 3.1 shows the maximum number of equilibria in each region. Within each region, Algorithm 4 is applied to compute all the equilibria efficiently.*

$\mathcal{J}_3$	2	1	1
$\mathcal{J}_2$	1	①	1
$\mathcal{J}_1$	①	1	2
	$\mathcal{I}_1$	$\mathcal{I}_2$	$\mathcal{I}_3$

Table 3.1: Maximum number of equilibria in each region

Algorithm 4. *1. For cells in which there is at most 1 equilibrium:*

(a) *Find the feasible lower bound and upper bound* $[x_L, x_U]$*;*

(b) *Evaluate* $g(x_L)$ *and* $g(x_U)$*. If they have opposite signs, then apply bisection method to* $g(x)$ *over* $[x_L, x_U]$*.*

2. *For cells in which there are at most* 2 *equilibria:*

(a) *Find the feasible lower bound and upper bound* $[x_L, x_U]$*;*

(b) *Evaluate* $g(x_L)$ *and* $g(x_U)$*, if they have opposite signs, apply bisection method to* $g(x)$ *over* $[x_L, x_U]$*;*

(c) *If* $g(x_L)$ *and* $g(x_U)$ *have the same sign, search for the unique local maxima or minima* x_M*. If* x_M *exists and* $g(x_M)$ *has opposite signs with* $g(x_L), g(x_U)$*, then apply bisection method to* $g(x)$ *over* $[x_L, x_M]$ *and* $[x_M, x_U]$ *separately;*

On the two regions $[\mathcal{I}_1 \times \mathcal{J}_1]$ and $[\mathcal{I}_2 \times \mathcal{J}_2]$, we circle the number in Table 3.1 to indicate that we are unable to prove the maximum number of equilibria. However, in our extensive numerical experiments, we find that there is at most 1 equilibrium in each region. Therefore Algorithm 4 works well for finding all the equilibria in practice.

3.4.2 Linear Relaxation

In this subsection, instead of computing all the equilibria exactly, we aim to compute a lower bound on $\underline{\pi}^{\mathcal{R}}(\boldsymbol{p})$. The approach is based on an idea proposed by Yamamura and Fujioka (2003) for finding all solutions of systems of nonlinear equations. One of their main ideas is to identify regions without any solution by checking (in)feasibility of a linear program.

Suppose $\boldsymbol{q} = (q_1, q_2)$ is restricted on one region $\mathcal{R} = [L_1, U_1] \times [L_2, U_2] \in [0, 1]^2$, then the pessimistic optimization problem on the bounded region can be formulated

$$\begin{aligned} \underline{\pi}^{\mathcal{R}}(\boldsymbol{p}) &= \min_{\boldsymbol{q}} \pi(\boldsymbol{p}, \boldsymbol{q}) \\ \text{s.t.} \quad & q_i e^{p_i - \alpha q_i} + q_1 + q_2 - 1 = 0, \\ & L_i \le q_i \le U_i. \end{aligned} \tag{PES-bd}$$

For our purposes below, we will be interested only in two choices $[L_i, U_i] \subseteq [0, 2/\alpha]$ or $[L_i, U_i] \subseteq [2/\alpha, r_i]$, where $r_i = q_i^b$ is defined in Proposition 3.4.3. Define $h_i(q_i) := q_i e^{p_i - \alpha q_i}$ and formulate the following linear problem as a relaxation of Problem (PES-bd).

$$\begin{aligned} \rho^{\mathcal{R}}(\boldsymbol{p}) &= \min_{\boldsymbol{q}} \pi(\boldsymbol{p}, \boldsymbol{q}) \\ \text{s.t.} \quad & y_i + q_1 + q_2 - 1 = 0, \quad i = 1, 2 \\ & L_i \le q_i \le U_i, \quad i = 1, 2 \\ & y_i \ge A_i^k, \quad i = 1, 2 \text{ and all } k \\ & y_i \le B_i^k, \quad i = 1, 2 \text{ and all } k \end{aligned} \tag{LN}$$

where

1) if $[L_i, U_i] \subseteq [0, 2/\alpha]$, then $A_i^1 = h_i(L_i) + \frac{h_i(U_i) - h_i(L_i)}{U_i - L_i}(q_i - L_i)$, $B_i^1 = h_i'(L_i) q_i + h_i(L_i) - h_i'(L_i) L_i$, $B_i^2 = h_i'(U_i) q_i + h_i(U_i) - h_i'(U_i) U_i$ and $B_i^3 = h_i(1/\alpha)$.

2) if $[L_i, U_i] \subseteq [2/\alpha, r_i]$, then $A_i^1 = h_i'(U_i) q_i + h_i(U_i) - h_i'(U_i) U_i$, $A_i^2 = h_i'(L_i) q_i + h_i(L_i) - h_i'(L_i) L_i$ and $B_i^1 = h_i(L_i) + \frac{h_i(U_i) - h_i(L_i)}{U_i - L_i}(q_i - L_i)$.

The idea using Problem (LN) as a relaxation of Problem (PES-bd) is illustrated in Figure 3.4. In this example, $\alpha = 5$. In Problem (LN), a set of new variables

y_i are added to substitute the nonlinear terms $h_i(q_i)$ in Problem (PES-bd). On the interval $q_i \in [0, 2/\alpha]$, function $h_i(q_i)$ is concave as shown in Figure 3.4(a), thus bounded within the polygon formed by A_i^1, B_i^1, B_i^2, and B_i^3. On the interval $q_i \in [2/\alpha, 1]$, function $h_i(q_i)$ is convex as shown in Figure 3.4(b), and thus bounded within the polygon formed by A_i^1, A_i^2, and B_i^1. Therefore, all the points satisfying the constraints in Problem (PES-bd) must also satisfy the constraints in Problem (LN). The following lemma summarizes the relationship between Problem (PES-bd) and (LN).

Proposition 3.4.4. *Given fixed prices* $\boldsymbol{p} = (p_1, p_2)$ *and region* $\mathcal{R} = [L_1, U_1] \times [L_2, U_2]$, *Problem (LN) is a linear relaxation of Problem (PES-bd) and*

(I) If Problem (LN) is infeasible, then Problem (PES-bd) is also infeasible, and therefore there is no equilibrium on $\mathcal{R}$.

(II) If Problem (LN) has an optimal value $\rho^{\mathcal{R}}(\boldsymbol{p})$, *then either (i.) Problem (PES-bd) is infeasible, or (ii.) Problem (PES-bd) is feasible and* $\underline{\pi}^{\mathcal{R}}(\boldsymbol{p}) \geq \rho^{\mathcal{R}}(\boldsymbol{p})$.

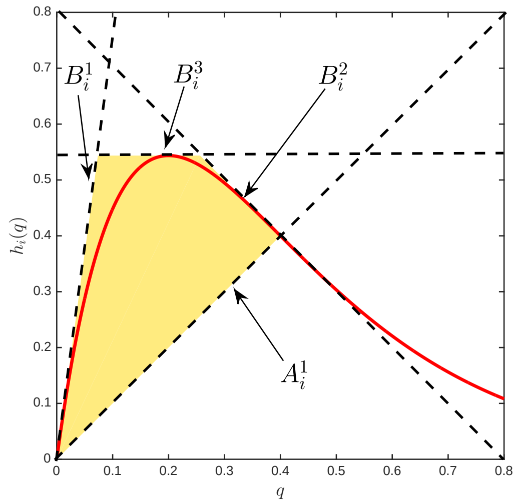

(a) On the interval $[0, 2/\alpha]$

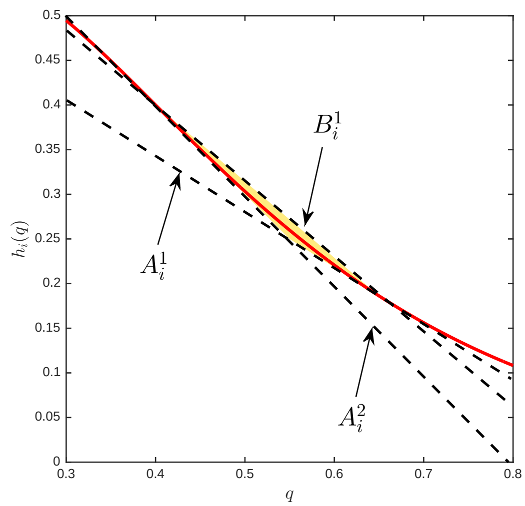

(b) On the interval $[2/\alpha, r_i]$

Figure 3.4: Illustration of relaxation

By Proposition 3.4.4, for a given region $\mathcal{R}$, solving problem (LN) gives us a lower bound of the optimal objective value $\underline{\pi}^{\mathcal{R}}(\boldsymbol{p})$ to problem (PES-bd). To derive a tighter value of $\rho^{\mathcal{R}}(\boldsymbol{p})$ to $\underline{\pi}^{\mathcal{R}}(\boldsymbol{p})$, we can further restrict the initial domain $\mathcal{R}$ to a smaller region which contains all equilibria. The idea is to change the objective function in (LN), from minimizing the total revenue to a new objective function of maximizing/minimize q_i, while keeping all constraints the same as in (LN). Namely, given the initial bounds $\mathcal{R} = [L_i, U_i]^2$, we formulate a set of new problems

$$\begin{aligned} &\min / \max q_i \\ \text{s.t.} \quad & y_i + q_1 + q_2 - 1 = 0, \quad i = 1, 2 \\ & L_i \leq q_i \leq U_i, \quad i = 1, 2 \\ & y_i \geq A_i^k, \quad i = 1, 2 \text{ and all } k \\ & y_i \leq B_i^k, \quad i = 1, 2 \text{ and all } k \end{aligned} \tag{LN-bd}$$

where A_i^ks and B_i^ks are defined in Problem (LN).

By Proposition 3.4.4, all the equilibria in (PES-bd) represented by $Q(\boldsymbol{p}) \cap \mathcal{R}$ are contained in the feasible domain of (LN). Since (LN) and (LN-bd) share the same constraints, $Q(\boldsymbol{p}) \cap \mathcal{R}$ is also contained in the feasible domain of (LN-bd). Therefore, maximizing/minimizing the decision variable q_i results in a stricter feasible region for q_i that contains $Q(\boldsymbol{p}) \cap \mathcal{R}$, as well as a smaller region $\mathcal{R}'$. By using $\mathcal{R}'$ as the new region in (LN), we can again easily solve it and derive $\rho^{\mathcal{R}'}(\boldsymbol{p})$. If $\rho^{\mathcal{R}'}(\boldsymbol{p})$ is strictly higher than $\rho^{\mathcal{R}}(\boldsymbol{p})$, then it is a tighter lower bound for $\underline{\pi}^{\mathcal{R}}(\boldsymbol{p})$. By following this procedure recursively, we could derive an increasingly tighter bound for $\underline{\pi}^{\mathcal{R}}(\boldsymbol{p})$. The detailed algorithm is described in the following.

Algorithm 5. *Find a lower bound $\rho^{\mathcal{R}}(\boldsymbol{p})$ for $\underline{\pi}^{\mathcal{R}}(\boldsymbol{p})$. Input price $\boldsymbol{p} = (p_1, p_2)$ and the initial region $\mathcal{R}^0 = [L_i, U_i]^2$. Set $\mathcal{R} = \mathcal{R}^0$ and starting from $\rho^{\mathcal{R}}(\boldsymbol{p}) = 0$*

1. *Test the feasibility of Problem (LN). If (LN) is feasible, go to step* 2*; if (LN) is infeasible, end the algorithm.*

2. *Solve Problem (LN). If the objective value at optimum is strictly higher than* $\underline{\pi}$*, save it as the new* $\underline{\pi}$*.*

3. *Solve Problem (LN-bd) for the new bounds of* q_i*, and save it as* $\mathcal{R}'$*. If* $\mathcal{R}'$ *and* $\mathcal{R}$ *are equivalent, end the algorithm; otherwise, update* $\mathcal{R}$ *with* $\mathcal{R}'$ *and go back to Step* 1*.*

Output $\rho^{\mathcal{R}}(\boldsymbol{p})$*.*

3.5 Comparison Between One and Two-product Cases

In this section, we study how the number of products available on the market impacts the revenue performance. In particular, we compare the one-product case with the two-product case. First, we consider the (BC) problems and compare their revenues in the two scenarios. We have the following proposition.

Proposition 3.5.1. *Let* $\overline{\pi}^m$ *denote the optimistic optimal revenue for the* m*-product problem, then* $\lim_{\alpha\to\infty} \overline{\pi}^{m+1} - \overline{\pi}^1 = 0$*.*

Proposition 3.5.1 implies that when the network effect is strong, the revenue from offering just one product is almost as high as that from selling an assortment of multiple products. This is because the sales quantity of the second product goes to zero when α is infinitely large.

Next we consider the pessimistic case. Define the pessimistic optimal pricing decisions as $\underline{\boldsymbol{p}}$, the worst-case sales in equilibrium as $\underline{\boldsymbol{q}}$, and the corresponding total revenue as $\underline{\pi}_2 = \underline{\boldsymbol{p}}^T \cdot \underline{\boldsymbol{q}}$. If the supremum does not exist in Problem (WC), then we must interpret $\boldsymbol{p}$ and $\boldsymbol{q}$ in a limiting sense similar to part 2 of Lemma 3.3.4. In comparison, the equilibrium condition (3.6) in the one-product case can also be transformed into $f(q|p) = 0$. We define the optimal price, the resulting sales, and the total revenue in the pessimistic scenario as $\underline{p}, \underline{q}$ and $\underline{\pi}_1$, then they have the following property.

Lemma 3.5.2. *1) For any $\epsilon > 0$, if $p = \alpha - \log \epsilon$, then $q \in [0, \epsilon)$ is a necessary condition for $f(q|p) > 0$.*

2) For any $p' < p^L$, if the price in one-product case is set to $p_1 = p'$ and the price in the two-product case is set to $\boldsymbol{p}^\epsilon = (p', \alpha - \log \epsilon)$, and the pessimistic revenues from these two cases are $\underline{\pi}_1(p')$ and $\underline{\pi}_2(\boldsymbol{p}^\epsilon)$ respectively, then we have

$$\lim_{\epsilon \to 0} \|\underline{\pi}_1(p') - \underline{\pi}_2(\boldsymbol{p}^\epsilon)\| = 0. \tag{3.12}$$

3) For the optimal pessimistic revenue $\underline{\pi}_1$ in one-product case and the optimal pessimistic revenue $\underline{\pi}_2$ in two-product case, they follow

$$\underline{\pi}_1 \leq \underline{\pi}_2 \tag{3.13}$$

The first part of Lemma 3.5.2 says that if the price is high enough, the function $f(q|p)$ will be almost overlapping the vertical axis. Due to the fact that $\underline{\boldsymbol{q}}$ has to satisfy both equations in (3.10), it implies that when p_2 is set at a high value, the condition $q_1 = f(q_2|p_2)$ is almost equivalent to $q_2 = 0$. Therefore, the condition $q_2 = f(q_1|p_1)$ is almost equivalent to $f(q_1|p_1) = 0$, which is exactly the equilibrium condition in the one-product case. The second part of Lemma 3.5.2 further exploits this result: for the two-product problem, we can always find a pair of prices such that the resulting pessimistic revenue is close to the pessimistic revenue $\underline{\pi}_1$ in the one-product case. In other words, it shows that the optimal price $\underline{p}$ and the resulting sales quantity $\underline{q}$ in the one-product case can always be approximately achieved by setting $p_1 = \underline{p}$ and $p_2 = L$ in the two-product case, where L is a very large number. It also gives us the result that $\underline{\pi}_1 \leq \underline{\pi}_2$.

In Table 3.2, we compare the optimal revenues between $n = 1$ and $n = 2$ for various α. Naturally, a two-product portfolio always brings higher revenue than a single-product case. In addition, we also compare the pessimistic price $\underline{p}_1$ in the $n = 1$ case, and the lower price $\underline{p}_2^L$ of the pessimistic optimal prices in $n = 2$ case. It is shown that ulp_1 is always higher than $\underline{p}_2^L$. However, the difference in revenue as well as between optimal prices shrinks and when $\alpha \geq 5$, there is almost no difference between them. Therefore, Table 3.2 shows us that when the network effect is strong, the revenue from offering exactly one product is almost equivalent to that from selling a portfolio of multiple products. This result suggests that a company selling products with high network effects may be wise to offer just one product.

3.6 Numerical Study

In this section, we conduct some numerical experiments to understand more about the worst-case problem. In particular, we demonstrate that under the two different perspectives — optimistic or pessimistic — the optimal pricing decisions as well as the resulting revenues could be drastically different. In addition, their difference

α	3	3.5	4	4.5	5	5.5	6
$\underline{\pi}_1$	0.8769	1.1456	1.4510	1.7840	2.0930	2.3128	2.4811
$\underline{\pi}_2$	0.8881	1.1494	1.4523	1.7844	2.0931	2.3128	2.4811
$\frac{\underline{\pi}_2-\underline{\pi}_1}{\underline{\pi}_1}$	1.276%	0.334%	0.092%	0.023%	$<0.001\%$	$<0.001\%$	$<0.001\%$
$\underline{p}_1$	1.2493	1.5072	1.8124	2.1492	2.3444	2.4727	2.5849
$\underline{p}_2^L$	1.240	1.501	1.810	2.148	2.343	2.472	2.584
$\frac{\underline{p}_2^L-\underline{p}_1}{\underline{p}_1}$	0.74%	0.42%	0.13%	0.06%	0.06%	0.03%	0.03%

Table 3.2: Revenue comparision between $n = 1$ and $n = 2$ for different α.

becomes more obvious when the network effect becomes stronger. In the following, we study a two-product example. We fix the intrinsic values y to be 0. The network effect parameter α is taken in the range of $[3, 5.5]$, at an increasing rate of 0.1 in each step. Problem (BC) is solved to derive $\overline{\boldsymbol{p}}(\alpha)$ and it can be computed by using the algorithm in Du et al. (2016), but Problem (WC) is more complicated and the detailed procedure is as follows:

1. Given any α, set the bounds for prices $p_1 \in [0, 30]$ and $p_2 \in [0, 30]$, then construct a grid collection $\{p_i^t : t = 1, \ldots, T_i\}$ at the precision 0.001.

2. For each grid point $\boldsymbol{p} = (p_1, p_2)$, apply Algorithm 4 on equations (3.10) to derive all the equilibria and store them in $Q(\boldsymbol{p})$. If there is more than one equilibrium, evaluate their revenues and choose the equilibrium $\underline{q}(\boldsymbol{p})$ with the lowest revenue $\underline{\pi}(\boldsymbol{p})$.

3. Compare $\underline{\pi}(\boldsymbol{p})$ for all the grid points and find the price grid point with the highest value.

The results of this numerical experiment are summarized in Figure 3.5. Figure 3.5(a) illustrates how $\underline{\boldsymbol{p}}(\alpha)$ and $\overline{\boldsymbol{p}}(\alpha)$ change with α. We notice that even though the two products have same parameters, their prices at optimality are asymmetric. When the network effect is weak $\alpha < 4.5$, the optimal pricing decisions under the two different scenarios are exactly the same and so are the optimal revenues. However, it no longer holds when $\alpha > 4.5$. The lower price $\underline{p}_1$ at optimality in the pessimistic case, is smaller than the lower optimal price $\overline{p}_1$ in the optimistic case. Meanwhile, the higher optimal price $\underline{p}_2$ in the pessimistic case is larger than $\overline{p}_2$. Furthermore, as α increase, $\underline{p}_2$ grows rapidly close to 30, which is the upper-bound limit imposed in the test. We believe the "jagged" appearance in $\underline{p}_2$ in Figure 3.5(b) arises from discretization.

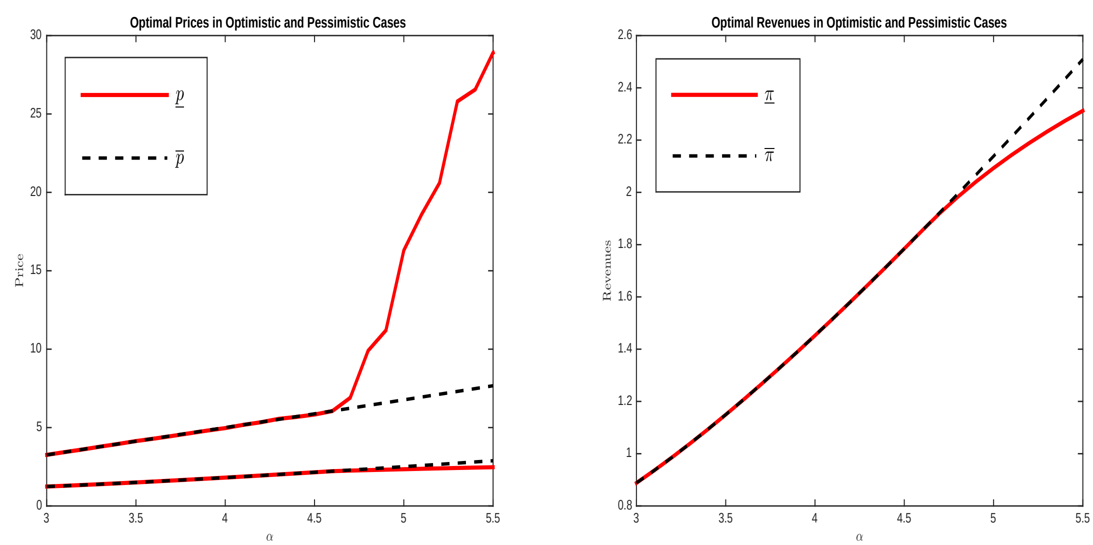

(a) Optimal prices (b) Optimal revenues

Figure 3.5: Comparison of pessimistic and optimistic cases

The behavior of $\overline{\boldsymbol{p}}$ and $\underline{\boldsymbol{p}}$ can be explained by investigating the equilibrium equations (3.10) for different α. In Figure 3.6, we plot the equilibrium equations at the optimal prices in these four different scenarios: (a) Pessimistic case with $\alpha = 3.5$; (b) Optimistic case with $\alpha = 3.5$; (c) Pessimistic case with $\alpha = 5.5$; (d) Optimistic case with $\alpha = 5.5$. By comparing Figure 3.6(a) and 3.6(b), we find that when α is relatively small, equilibrium equations (3.10) at the optimistic pricing decisions has only one solution. In this situation, it is easy to see that the optimal solutions in the pessimistic case are equal to that in the optimistic case. In contrast, when α is large, equilibrium equations (3.10) at the optimistic pricing decisions have multiple solutions. In Figure 3.6(c) and 3.6(d), the dashed line for the plot of $q_1 = f(q_2)$ is very close (also indistinguishable) to the horizontal axis because p_2 is very high. In Figure 3.6(c), at $\overline{\boldsymbol{p}} = (2.885, 7.890)$, there are three different equilibria $\boldsymbol{q}_a = (7.97, 0.034)$, $\boldsymbol{q}_b = (59.32, 0.015)$ and $\boldsymbol{q}_c = (86.96, 0.005)$. The highest revenue is $\pi_c = 2.5092$, in comparison with the lowest revenue $\pi_a = 0.2324$. These two numbers show that implementing the optimistic optimal prices would put the seller at the risk of losing 90.73% of his revenue, if the actual sales is the pessimistic one. In contrast, Figure 3.6(d) shows that when implementing the pessimistic optimal prices $\underline{\boldsymbol{p}} = (2.472, 29.912)$, there is only one equilibrium at $\boldsymbol{q} = (93.56, 0.00)$. The dashed curve "squeezes" just between the solid curve and the horizontal axis near $q_1 = 2.3$. So there is no equilibrium there and implementing the only pessimistic prices could secure the seller a revenue of 2.313. Therefore, there is a high incentive for the seller to implement the pessimistic optimal price.

As shown in Figure 3.6(c), lower-left equilibrium $\boldsymbol{q}_a$ is always unappealing because it results in a low revenue. To avoid this situation, p_1 tends to have a lower value such that the local minimum point of $q_2 = f(q_1)$ is above the x-axis. At the same time, p_2 tends to have a higher value to ensure that $q_1 = f(q_2)$ has no equilibrium at the lower-left region. These two effects lead to the observation that in Figure 3.6(d) there is only one equilibrium. In fact, throughout all our experiments, we find that implementing the optimal prices of Problem (WC) always leads to a unique equilibrium. Unfortunately, we are not able to prove this property.

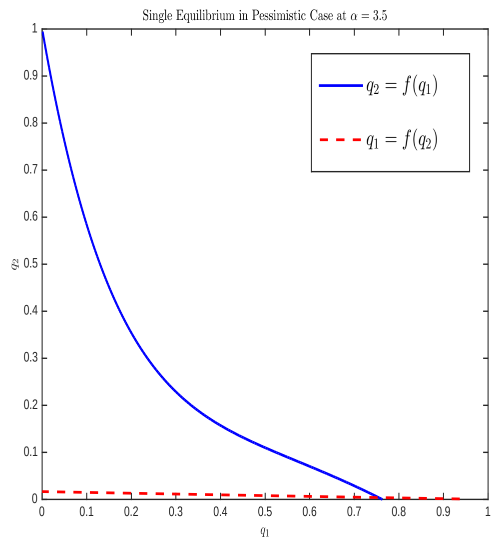

(a) Pessimistic case $\alpha = 3.5$

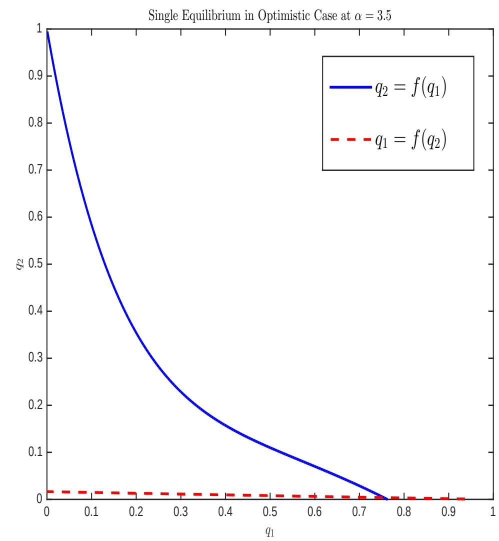

(b) Optimistic case $\alpha = 3.5$

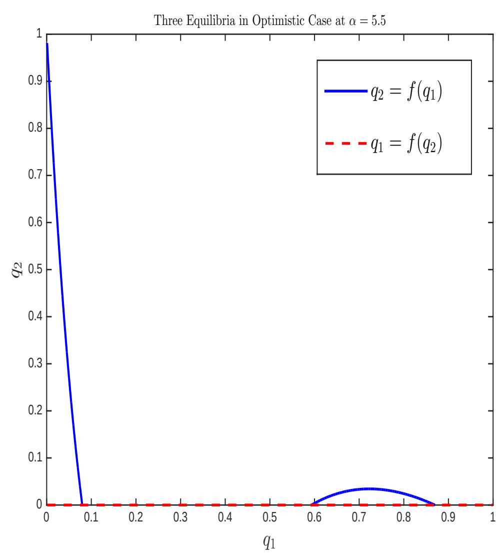

(c) Optimistic case $\alpha = 5.5$

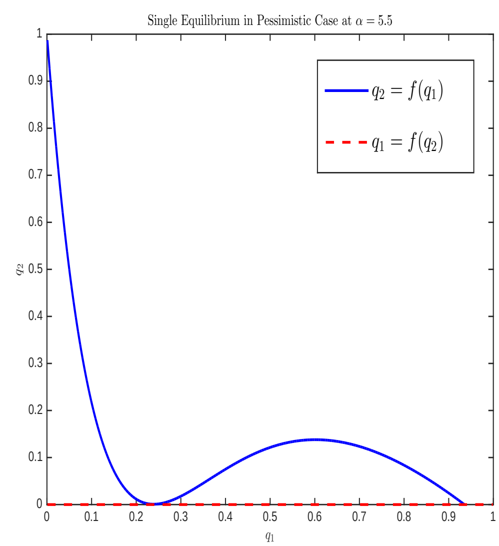

(d) Pessimistic case $\alpha = 5.5$

Figure 3.6: Equilibrium equations at optimality

In addition to compare the optimal decisions and their revenues directly in the two scenarios, we are also interested in how an optimal decision in a preassumed perspective would behave in the other one. Namely, we would like to understand $\underline{\pi}(\overline{\boldsymbol{p}})$ and $\overline{\pi}(\underline{\boldsymbol{p}})$ and the results are shown in Figure 3.7. As we have mentioned, in our numerical experiments, there is always only one equilibrium at optimality in the pessimistic case. Therefore, it is not surprising to see that the two revenue outcomes perfectly match in Figure 3.7(a) when implementing the pessimistic optimal prices $\underline{\boldsymbol{p}}$. In contrast, Figure 3.7(b) shows us that implementing the optimistic optimal prices $\overline{\boldsymbol{p}}$ could lead to a huge loss of revenue if the pessimistic outcome happens. To make things worse, the resulting revenue drops rapidly when α increases. Therefore, for a seller with a strong risk-averse attitude, the pessimistic optimal pricing decision is a safer choice that brings in a more secured amount of revenue.

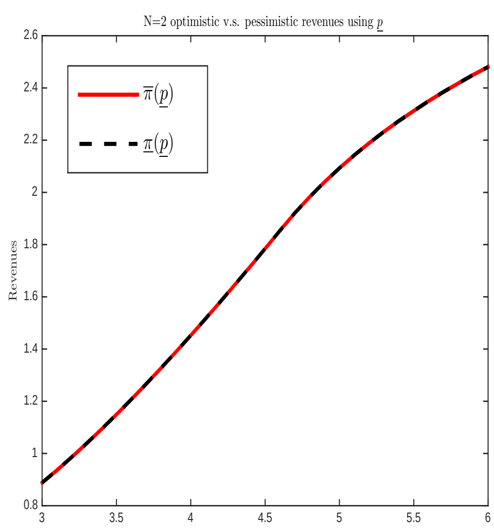

(a) Seller implements $\underline{\boldsymbol{p}}$

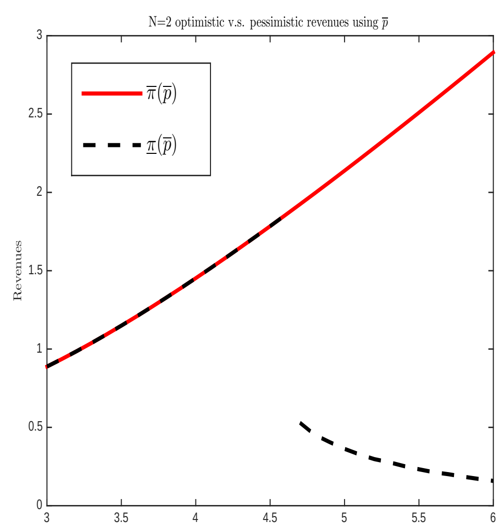

(b) Seller implements $\overline{\boldsymbol{p}}$

Figure 3.7: Comparison of two revenue outcomes

3.7 Conclusion

In this chapter, we employ a robust optimization perspective to study the pricing optimization problem with network effects. The focus is to resolve the issue of multiple equilibria at optimal prices and study the optimal pricing problem from a pessimistic perspective. When there is only one product to sell, we find a close relationship between the best-case problem and worst-case problem. Moreover, the worst-case problem can be solved by adding an extra constraint to the best-case problem. When the firm has two products to sell, we propose two different methods for computing the pessimistic solution: divide and search, and linear relaxation. Numerical studies imply that at the pessimistic optimal prices there is always only one equilibrium for sales quantities. To put the preceding statement on more rigorous ground, we need to determine whether the supremum in (WC) is a maximum for the case when $n = 2$. Other results in the numerical studies show that the potential loss of revenue from using the optimistic optimal prices could be huge, compared with using the pessimistic optimal prices.

Chapter 4

Future Research Directions

In this chapter, we briefly discuss a number of future research directions based on the topic of the multinomial logit choice model with network effects.

4.1 Joint Assortment Planning and Pricing Optimization

One important topic in revenue management is assortment planning. Many companies face the problem of choosing the "right" products for the market among many potential products. There is often a limit on the number of products that can be contained in the assortment. The constraint on the assortment capacity could be due to reasons like budget, marketing strategies, production cost, etc. In Wang (2012), the author discusses the joint assortment planning and pricing problem under the standard MNL model. Here we are interested in how network effects would influence assortment planning and pricing decisions. In this scenario, the seller will make two decisions with the goal of maximizing the revenue: 1) Choose an assortment from the potential products; 2) Set prices for the products in the chosen assortment.

To model this scenario, we use the same setting as in Section 2.2 with one extra restriction: the seller can only choose an assortment $\mathcal{S}$ with at most C products from the product set $\mathcal{N} = \{1, \ldots, n\}$. Given an assortment $\mathcal{S} \subseteq \mathcal{N}$, the relationships of utilities, purchase probabilities and prices still follow (2.1) – (2.4), with product set $\mathcal{N}$ replaced by $\mathcal{S}$. In the joint assortment planning and pricing problem, we have $\mathcal{S}$ and $\boldsymbol{q}^{\mathcal{S}} = (q_1^{\mathcal{S}}, \ldots, q_n^{\mathcal{S}})$ as the decision variables and the total revenue is

$$\pi(\mathcal{S}, \boldsymbol{q}^{\mathcal{S}}) = \sum_{j \in \mathcal{S}} \frac{\alpha_j}{\gamma_j} q_j^2 + \sum_{j \in \mathcal{S}} \frac{q_j}{\gamma_j} \log \left(1 - \sum_{j \in \mathcal{S}} q_j \right) + \sum_{j \in \mathcal{S}} \frac{q_j}{\gamma_j} \left(y_j - \log q_j \right). \tag{4.1}$$

The seller's joint assortment planning and pricing problem can now be formulated as the following optimization problem:

$$\begin{aligned} \max \quad & \pi(\mathcal{S}, \boldsymbol{q}^{\mathcal{S}}) \\ \text{s.t.} \quad & |\mathcal{S}| \leq C, \\ & \sum_{j \in \mathcal{S}} q_j \leq 1, \quad q_i \geq 0, \ i \in \mathcal{S}, \end{aligned}$$

where $|\mathcal{S}|$ denotes the cardinality of set $\mathcal{S}$.

There are many papers discussing assortment optimization problems under variants of MNL models. In the pioneering work of Talluri and van Ryzin (2004), the authors study assortment problems under MNL without any constraints and show that the optimal assortment is always a revenue-ordered assortment, which consists of products in decreasing price order. Wang and Wang (2016) consider the same problem under the network MNL model: They find that the problem is generally NP-Hard, but a new class of assortments called quasi-revenue-ordered assortments is optimal under some mild conditions. Adding constraints in the assortment problems (e.g., on the size of the assortment) often makes the problems more challenging. Davis et al. (2013) show that a class of problems with certain types of constraints on the assortment can be reformulated and solved as linear programming problems. Gallego and Topaloglu (2014) study assortment optimization problems under the nested logit model, with constraints on the set of products offered in each nest. They also solve joint assortment optimization and pricing problems under the nested logit model.

Back to our problem, the first goal is to find an efficient algorithm to identify the optimal assortment $\mathcal{S}^*$, and the optimal price for each product. Another interesting question arises from comparing the optimal assortment of the above problem with the original problem. From the discussion in Section 2.4, we know that in the case of no capacity constraint the optimal solution could be to boost the sales of a single product. One natural question is, when we limit the size of assortment, does the optimal assortment always include this product? In the extreme case, if we only allow 1 product in the assortment, is it optimal to choose that product which is boosted in the unconstrained setting? This remains an interesting question to be answered.

4.2 Nested Logit Model with Network Effects

In order to address the IIA property of the standard MNL model, McFadden (1978, 1980) proposed a model with a nested choice structure, namely, the nested logit (NL) model. The nested logit model is derived from the idea that consumers often make choices in a hierarchical way: a potential buyer first chooses a group of products among several possible groups, and then limits her subsequent selection within the chosen group. In recent years, pricing optimization problem under the NL model has been an area of active research. Li and Huh (2011) study the pricing problem for a profit-maximizing monopoly and show that the total profit function is concave in the market shares. In addition, they consider an oligopoly price-competition setting where each firm owns one nest of products. By solving for the price equilibrium, they find that competition drives up the total market share and drives down the prices, but the revenue for a particular product may go in either way, depending on the product-specific parameters. Gallego and Wang (2014) study the multi-product pricing problem under the NL model. They define the "adjusted markup" for each product as its price minus the sum of its cost and the reciprocal of its price sensitivity, and show that at optimality adjusted markup is constant for all products within a nest. They further show that each nest has an adjusted nest-level markup that is nest invariant. This property of optimal solution helps reduce the problem to a single variable optimization over a bounded interval. In Li et al. (2015), the authors provide an efficient algorithm to find the optimal assortment under a d-level nested logit model. They also study the price optimization problem with the assortment fixed, and develop a price-converging algorithm that performs much faster than gradient-based methods in their numerical experiments.

To our knowledge, no one has considered the NL model with network effects, and we want to investigate how network effects would influence the sellers' decisions in this framework. To describe the model more clearly, we start with a two-stage nested logit model but the idea can be extended to multiple stages. Suppose all the products can be categorized into n nests, and in nest $i \in \mathcal{N} = \{1, \ldots, n\}$ there are m_i products. A customer will first choose a nest, then choose within that nest. The utility from choosing product j in nest i is

$$u_{ij} = v_{ij} + \epsilon_{ij} = y_{ij} - p_{ij} + \alpha_{ij} x_{ij} + \beta_i \sum_{k=1}^{m_i} x_{ik} + \epsilon_{ij}. \tag{4.2}$$

Here y_{ij} is the intrinsic utility of that product, p_{ij} is the product price (for simplicity, we assume all the customers are equally sensitive to prices and normalize the price sensitivity to 1). We use parameter α_{ij} to represent the strength of networks for product ij and x_{ij} to denote the overall consumption of product ij. A customer gains more utility from purchasing product ij if more other customers purchase product ij. All the above interpretations are the same as in the standard MNL with networks effects. A main difference of using the nested logit model comes from the term $\beta_i \sum_{k=1}^{m_i} x_{ik}$, where β_i can be interpreted as the strength of network effects for nest i and $\sum_{k=1}^{m_i} x_{ik}$ is the aggregate consumption of all the products in nest i. Therefore $\beta_i \sum_{k=1}^{m_i} x_{ik}$ is the extra utility from the influence of aggregate consumption in nest i.

In the nested logit model, we use $q_{j|i}$ to denote the probability that a customer purchases product j given he has chosen category i, and Q_i to denote the probability for a customer to choose category i. Then we have

$$q_{j|i} = \frac{\exp(v_{ij})}{\sum_{k=1}^{m_i} \exp(v_{ik})}, \quad Q_i = \frac{\eta_i^{\gamma_i}}{1 + \sum_{l=1}^{n} \eta_l^{\gamma_l}},$$

where $\eta_i = \sum_{j=1}^{m_i} \exp(v_{ij})$. The γ_l's are called nest dissimilarity indices of product categories and represent the degree of inter-nest heterogeneity. Here we assume $0 < \gamma_l < 1$, which means products are more similar within nest i than across nests.

Then the probability of choosing product j in nest i is $q_{ij} = Q_i q_{j|i}$ and the profit generated from that product is

$$\pi_{ij} = (p_{ij} - c_{ij})q_{ij}. \tag{4.3}$$

To explain the reasoning behind the model, we can take the mobile games industry as an example. Mobile games can be grouped into different categories such as strategy, adventures, actions, etc. Customers have their own preferences over the categories, and when they are choosing among the wide variety of games, they first decide one category and then select one particular game within that category. To validate the term $\beta_i \sum_{k=1}^{m_i} x_{ik}$ in the utility function, think about this scenario. Suppose one game gains popularity in a short time. On one hand, it benefits the original game buyers because more players means more interaction in the gameplay; on the other hand, this game's rising fame would attract attention of customers who have no prior interest in this game category, and there is a chance that they choose other products other than the popular one. Therefore, extra consumption of one product benefits the players of all games in the same category.

Unlike in the standard model where all products are sold by one seller, here we do not specify whether one monopoly firm sells all the products or multiple firms sell their own exclusive products. One possible situation is that there are n firms, and each firm focuses on one category and sells products within that category. Therefore the firms compete over nests but monopolize within their own nests. The video game industry falls into this setting, because most game developers have their expertise and strength in one category of games and put all of their efforts into that category. Another possible situation is that there are multiple firms selling exclusive products across different nests, and they not only compete over different nests but also compete within the nests. The cloud storage service is a suitable example for this setting. Dropbox and Box.com both provide cloud storage service for differentiated space. Light users may find the free space suffices for their demands, while enterprise users need to pay much more for storing their massive data. Then cloud storage service can be grouped by their tiered storage capacity and the two companies compete with each other in each market segment.

The decision makers can take advantage of the nested logit model with network effects for their flexible purpose. In the product planning stage, they can use this model to evaluate which nest is best for a new product. In the product developing stage, they can use the model to design the product with appropriate features. In the selling stage, prices are the decision variables to help them maximize the revenue.

References

S. P. Anderson and A. de Palma. Multiproduct firms: A nested logit approach. *The Journal of Industrial Economics*, 40(3):261–276, 1992.

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

G. Aydin and E. L. Porteus. Joint inventory and pricing decisions for an assortment. *Operations Research*, 56(5):1247–1255, 2008.

P. Belobaba and C. Hopperstad. Boeing/mit simulation study: Pods results update. AGIFORS Reservations and Yield Management Study Group Symposium, 1999.

O. Besbes and A. Zeevi. On the (surprising) sufficiency of linear models for dynamic pricing with demand learning. *Management Science*, 61(4):723–739, 2015.

F. Bloch and N. Quérou. Pricing in social networks. *Games and Economic Behavior*, 80:243–261, 2012.

B. A. Brock and S. N. Durlauf. A multinomial-choice model of neighborhood effects. *The American Economic Review*, 92(2):298–303, 2002.

S. Brumelle and J. I. McGill. Airline seat allocation with multiple nested fare classes. *Operations Research*, 41:127–137, 1993.

S. Brumelle, J. I. McGill, T. H. Oum, K. Sawaki, and T. M. W. Allocation of airline seat between stochastically dependent demands. *Transportation Science*, 24:183–192, 1990.

O. Candogan, K. Bimpikis, and A. Ozdaglar. Optimal pricing in networks with externalities. *Operations Research*, 60(4):883–905, 2012.

W. L. Cooper and L. Li. On the use of buy up as a model of customer choice in revenue management. *Production and Operations Management*, 21:833–850, 2012.

W. L. Cooper, T. Homem-de Mello, and A. J. Kleywegt. Models of the spiral-down effect in revenue management. *Operations Research*, 54:968–987, 2006.

J. Davis, G. Gallego, and H. Topaloglu. Assortment planning under the multinomial logit model with totally unimodular constraint structures. Working paper, Cornell University, 2013.

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

C. Du, W. L. Cooper, and Z. Wang. Optimal pricing for a multinomial logit choice model with network effects. *Operations Research*, 64(2):441–455, 2016.

P. H. Dybvig and C. S. Spatt. Adoption externalities as public goods. *Journal of Public Economics*, 20(2):231–247, 1983.

N. Economides. The economics of networks. *International Journal of Industrial Organization*, 14(6):673–699, 1996a.

N. Economides. Network externalities, complementarities, and invitations to enter. *European Journal of Political Economy*, 12(2):211–233, 1996b.

J. Farrell and G. Saloner. Standardization, compatibility, and innovation. *The RAND Journal of Economics*, 16(1):70–83, 1985.

A. Galeotti, S. Goyal, M. O. Jackson, F. Vega-Redondo, and L. Yariv. Network games. *The Review of Economic Studies*, 77(1):218–244, 2010.

G. Gallego and H. Topaloglu. Constrained assortment optimization for the nested logit model. *Management Science*, 60(10):2583 – 2601, 2014.

G. Gallego and R. Wang. Multiproduct price optimization and competition under the nested logit model with product differentiated price sensitivities. *Operations Research*, 62(2):450–461, 2014.

W. Hanson and K. Martin. Optimizing multinomial logit profit functions. *Management Science*, 42(7):992–1003, 1996.

M. O. Jackson and L. Yariv. Diffusion of behavior and equilibrium properties in network games. *The American Economic Review*, 97(2):92–98, 2007.

M. Katz and C. Shapiro. Network externalities, competition, and compatibility. *The American Economic Review*, 75(3):424–440, 1985.

M. Katz and C. Shapiro. Technology adoption in the presence of network externalities. *Journal of Political Economy*, 94(4):822–841, 1986.

D. M. Kreps. *Microeconomic Foundations I: Choice and Competitive Markets*. Princeton University Press, 2012.

G. Li, P. Rusmevichientong, and H. Topaloglu. The d-level nested logit model: Assortment and price optimization problems. *Operations Research*, 63(2):325 – 342, 2015.

H. Li and W. T. Huh. Pricing multiple products with the multinomial logit and nested logit models: Concavity and implications. *Manufacturing & Service Operations Management*, 13(4):549–564, 2011.

K. Littlewood. Forecasting and control of passenger bookings. *Proceedings of 12th AGIFORS Symposium*, 4, 1972.

D. McFadden. Conditional logit analysis of qualitative choice behavior. In P. Zarembka, editor, *Frontiers in Econometrics*, pages 105–142. Academic Press, New York, 1974.

D. Merkin. *Introduction to the Theory of Stability*. Springer, 1997.

T. Miyao and P. Shapiro. Discrete choice and variable returns to scale. *International Economic Review*, 22(2):257–273, 1981.

W. Z. Rayfield, P. Rusmevichientong, and H. Topaloglu. Approximation methods for pricing problems under the nested logit model with price bounds. *INFORMS Journal on Computing*, 2015. Forthcoming.

L. W. Robinson. Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes. *Operations Research*, 43:252–263, 1995.

J. Rohlfs. A theory of interdependent demand for a communication service. *The Bell Journal of Economics and Management Science*, 5(1):16–37, 1974.

C. Starkweather. Network effects and externalities with logit demand. Discussion Papers in Economics, University of Colorado, 2003.

M. Suh and G. Aydin. Dynamic pricing of substitutable products with limited inventories under logit demand. *IIE Transactions*, 43(5):323–331, 2011.

A. Sundararajan. Local network effects and complex network structure. *The B.E. Journal of Theoretical Economics*, 7(1), 2007.

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

D. M. Topkis. *Supermodularity and Complementarity.* Princeton University Press, 1998.

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

G. J. van Ryzin and J. I. McGill. Revenue management without forecasting or optimization: An adaptive algorithm for determining seat protection levels. *Management Science*, 46:760–775, 2000.

R. Wang. Capacitated assortment and price optimization under the multinomial logit model. *Operations Research Letters*, 40(6):492–497, 2012.

R. Wang and Z. Wang. Consumer choice models with endogenous network effects. *Management Science*, 2016. Articles in Advance.

R. D. Wollmer. An airline seat management model for a single leg route when lower fare classes book first. *Operations Research*, 40:26–37, 1992.

Z. Xue and J. Song. Demand management and inventory control for substitutable products. Working paper, Duke University, 2007.

K. Yamamura and T. Fujioka. Finding all solutions of nonlinear equations using the dual simplex method. *Journal of Computational and Applied Mathematics*, 152:587–595, 2003.

Appendix

Uniqueness and Stability of Optimal Sales Levels

In our developments, we have used $\boldsymbol{q}$ as the decision variable in the seller's revenue optimization problem. However, as mentioned in Section 2.2, given a price vector $\boldsymbol{p}$, there could be multiple $\boldsymbol{q}$ that satisfy the equilibrium condition (2.4), which we write here as $\boldsymbol{q} = F(\boldsymbol{p}, \boldsymbol{q})$ to show the dependence upon $\boldsymbol{p}$. We effectively assumed that the seller can select sales quantities $\boldsymbol{q}$ by using prices $\boldsymbol{p}(\boldsymbol{q}) = (p_1(\boldsymbol{q}), \ldots, p_n(\boldsymbol{q}))$ given by (2.5). In this section, we provide some justification for this assumption.

Observe first that for any $(\boldsymbol{p}, \boldsymbol{q})$ pair that satisfies $\boldsymbol{q} = F(\boldsymbol{p}, \boldsymbol{q})$, it must be that $\boldsymbol{p} = \boldsymbol{p}(\boldsymbol{q})$ because as noted in Section 2.2, given a $\boldsymbol{q}$, there is a unique $\boldsymbol{p}$ for which $\boldsymbol{q} = F(\boldsymbol{p}, \boldsymbol{q})$. Consequently, if a $(\boldsymbol{p}, \boldsymbol{q})$ pair satisfies $\boldsymbol{q} = F(\boldsymbol{p}, \boldsymbol{q})$, then $\sum_j q_j p_j = \sum_j q_j p_j(\boldsymbol{q}) \le \sum_j q_j^* p_j(\boldsymbol{q}^*) = \sum_j q_j^* p_j^*$. That is, if a $(\boldsymbol{p}, \boldsymbol{q})$ pair is an equilibrium, then the revenue accrued at that equilibrium is no greater than that accrued at the equilibrium sales quantities and prices identified in Theorems 2.3.1 and 2.3.2.

The seller will implement prices $\boldsymbol{p}^* = \boldsymbol{p}(\boldsymbol{q}^*)$, so it is of particular interest to look at the possibility that for the optimal $\boldsymbol{q}^*$, there could be some other $\boldsymbol{q}'$ that also satisfies the equilibrium condition for $\boldsymbol{p}^* = \boldsymbol{p}(\boldsymbol{q}^*)$, i.e., $\boldsymbol{q}' = F(\boldsymbol{p}^*, \boldsymbol{q}')$. By the preceding argument, we see that the revenue associated with $(\boldsymbol{q}', \boldsymbol{p}^*)$ cannot exceed that of $(\boldsymbol{q}^*, \boldsymbol{p}^*)$. Hence, if there are multiple equilibria associated with prices $\boldsymbol{p}(\boldsymbol{q}^*)$, then we may view the seller as picking the best one that can arise from those prices.

Proposition 4.2.1 below provides a sufficient condition that ensures that for each $\boldsymbol{p}$, there exists a unique $\boldsymbol{q}$ that satisfies (2.4), i.e., that satisfies $\boldsymbol{q} = F(\boldsymbol{p}, \boldsymbol{q})$. Consequently, for problems that satisfy the sufficient condition, if the seller implements prices $\boldsymbol{p}(\boldsymbol{q}^*)$, then the only sales levels that satisfy (2.4) are $\boldsymbol{q}^*$ and hence the issue of multiple equilibria is not present. The proposition is proved in Miyao and Shapiro (1981) for a model more general than ours and is specialized to the network choice model (2.4) in Wang and Wang (2016).

Proposition 4.2.1. *For any values of $\{(\alpha_i, \gamma_i, y_i) : i \in \mathcal{N}\}$ and any $\boldsymbol{p}$, there exists at least one solution to (2.4). Moreover, if $\alpha_i \leq 2$ for all $i \in \mathcal{N}$, then for any $\{(\gamma_i, y_i) : i \in \mathcal{N}\}$ and any $\boldsymbol{p}$, the solution to (2.4) is unique.*

One important question about the choice model we consider is whether sales levels $\boldsymbol{q}$ will converge to an equilibrium if they are initially out of equilibrium and customers repeatedly adjust their purchase decisions according to market conditions. To address this, we consider the following dynamics of the sales quantities:

$$q_i^t = \frac{\exp(y_i - \gamma_i p_i + \alpha_i q_i^{t-1})}{1 + \sum_{j=1}^n \exp(y_j - \gamma_j p_j + \alpha_j q_j^{t-1})}. \tag{4.4}$$

The following result is proved in Miyao and Shapiro (1981) and Wang and Wang (2016):

Proposition 4.2.2. *Suppose $|\alpha_i| \leq 2$ for all $i \in \mathcal{N}$. Fix any $\boldsymbol{p}$ and any $\boldsymbol{q}^0 = (q_1^0, \ldots, q_n^0)$ and consider $\{\boldsymbol{q}^t = (q_1^t, \ldots, q_n^t)\}$ in (4.4). Then $\{\boldsymbol{q}^t\}$ converges to the unique solution to (2.4).*

In Wang and Wang (2016), the authors use a set of DVD purchase data from a major online retailer to calibrate the network MNL choice model, and find the optimal fit of the coefficient α (they assume homogeneous α_i across products) is 0.998 (statistically significant). Therefore, in that case, the conditions in Propositions 4.2.1 and 4.2.2 hold, the equilibrium is unique for any prices, and dynamic customer adjustments will give us convergence to that equilibrium.

In general, if the α_i do not satisfy the conditions in the above propositions, then it is possible that there are multiple equilibria for (2.4) and the above dynamic adjustments (4.4) may converge to different equilibria depending on the starting point (see Section 2.6.1 for further discussion). In fact, such phenomenon of multiple equilibria is quite common in models that incorporate network effects; see, for example, Galeotti et al. (2010), Jackson and Yariv (2007), Sundararajan (2007), Economides (1996b), Katz and Shapiro (1985), Dybvig and Spatt (1983), and Rohlfs (1974), among many others. There is evidently not a single "right" answer to the question of what will happen in the presence of multiple equilibria. However, one notion that has been used to explain why a particular equilibrium might arise while another might not is that of (local) stability.

We will argue next that $\boldsymbol{q}^*$ is a stable equilibrium under prices $\boldsymbol{p}(\boldsymbol{q}^*)$. For the ensuing discussion we say that $\boldsymbol{q}^*$ is a stable equilibrium if all the eigenvalues of the Jacobian matrix of $F(\boldsymbol{q})$ at $\boldsymbol{q}^*$ have real part less than 1 (see Chapter 4 in Merkin 1997 for a reference). Such stability ensures that there exists a neighborhood of $\boldsymbol{q}^*$

such that the differential equation system $\frac{\partial \boldsymbol{q}(t)}{\partial t} = F(\boldsymbol{q}(t)) - \boldsymbol{q}(t)$ with starting point $\boldsymbol{q}(0)$ in a neighborhood of $\boldsymbol{q}^*$ will converge to $\boldsymbol{q}^*$ as t goes to infinity. In practice, the seller may first guide the customers to a neighborhood of $\boldsymbol{q}^*$ (e.g., by posting expected sales) and then the dynamics of the system will make the sales levels converge to $\boldsymbol{q}^*$, thereby justifying the choice of $\boldsymbol{q}^*$. Related notions of equilibrium stability are discussed in, e.g., Jackson and Yariv (2007) and Economides (1996b).

Proposition 4.2.3. *Consider the homogeneous case with $\alpha_i = \alpha \leq \hat{\alpha}$ ($\hat{\alpha}$ is defined in Theorem 2.3.1), $\gamma_i = 1$, and $y_i = y$ for all $i \in \mathcal{N}$. The optimal $\boldsymbol{q}^*$ is a stable equilibrium.*

Proof. Consider the Jacobian matrix $J = \frac{\partial F}{\partial \boldsymbol{q}}$. We have

$$\begin{aligned} \frac{\partial F_i}{\partial q_i} &= \frac{\alpha_i B_i (1 + \sum_{k=1}^n B_k) - \alpha_i B_i^2}{(1 + \sum_{k=1}^n B_k)^2} = \alpha_i q_i (1 - q_i), \\ \frac{\partial F_i}{\partial q_j} &= \frac{-\alpha_j B_i B_j}{(1 + \sum_{k=1}^n B_k)^2} = -\alpha_j q_i q_j, \quad i \neq j. \end{aligned}$$

where $B_i = \exp(y_i - \gamma_i p_i + \alpha_i q_i)$. When $\alpha_i = \alpha \leq \hat{\alpha}$, $\gamma_i = 1$, and $y_i = y$, the Jacobian matrix J can be written as:

$$J = \left[\frac{\partial F_i}{\partial q_j}\right]_{(i,j)} = \alpha \begin{bmatrix} q_1(1-q_1) & -q_1 q_2 & \cdots & -q_1 q_n \\ -q_1 q_2 & q_2(1-q_2) & & -q_2 q_n \\ \vdots & & \ddots & \vdots \\ -q_1 q_n & -q_2 q_n & \cdots & q_n(1-q_n) \end{bmatrix}.$$

By Theorem 2.3.1(a), $q_1^* = \cdots = q_n^* = q^*$ when $\alpha \leq \hat{\alpha}$. In this case, $J = \alpha q^* [I - q^* \mathbf{e}\mathbf{e}^T]$, where I is an identity matrix. The largest eigenvalue of J is αq^*. To see this, note that the eigenvalues of $\mathbf{e}\mathbf{e}^T$ are $0, \ldots, 0, n$ and therefore the eigenvalues of $[I - q^* \mathbf{e}\mathbf{e}^T]$ are $1, \ldots, 1, 1 - nq^*$ and the eigenvalues of J are

$\alpha q^*, \ldots, \alpha q^*, \alpha q^*(1 - nq^*)$. By Lemma 4.2.4, we have $\frac{\partial^2 \tilde{\pi}}{\partial s_i^2}(\boldsymbol{s}^*) = 2\alpha - \frac{1}{q^*} \leq 0$ in this case. Thus we have proved that $\boldsymbol{q}^*$ must be a stable fixed point when $\alpha \leq \hat{\alpha}$.

□

Unfortunately, we have not been able to prove the stability of the optimal $\boldsymbol{q}^*$ in the general case. However, we have conducted a large number of numerical experiments (with both homogeneous and heterogeneous parameters) and in all experiments, the optimal $\boldsymbol{q}^*$ is stable.

Proofs for Section 2.3

To prove the theorems and propositions in Section 2.3, we need the following lemma.

Lemma 4.2.4. *For any optimal solution $\hat{\boldsymbol{s}}$ to (P2), we must have $\frac{\partial \tilde{\pi}}{\partial s_i}(\hat{\boldsymbol{s}}) = 0$ and $\frac{\partial^2 \tilde{\pi}}{\partial s_i^2}(\hat{\boldsymbol{s}}) \leq 0$ for all $i \in \mathcal{N}$. Similarly, for any optimal solution $\hat{\boldsymbol{q}}$ to (P1), we must have $\frac{\partial \pi}{\partial q_i}(\hat{\boldsymbol{q}}) = 0$ and $\frac{\partial^2 \pi}{\partial q_i^2}(\hat{\boldsymbol{q}}) \leq 0$ for all $i \in \mathcal{N}$.*

Proof. Consider (P0) with homogeneous parameters. If we substitute for $\boldsymbol{q}$ in terms of $\boldsymbol{s}$, then we get a problem (P0′) that is the same as (P2) except the feasible region is $\{\boldsymbol{s} | 1/\sqrt{n} \geq s_1 \geq 0, q_i \geq 0\}$ where the q_i are given by (2.13). Next we show that any optimal solution $\tilde{\boldsymbol{s}}$ to (P0′) must be an interior point in the feasible region. For this, it suffices to show that $0 < \tilde{s}_1 < 1/\sqrt{n}$ and $\tilde{q}_i > 0$. The KKT condition for s_1 at optimality is

$$2\alpha s_1 + \sqrt{n}\left(y + \log(1 - \sqrt{n} s_1)\right) - \frac{\sqrt{n}}{1 - s_1\sqrt{n}} - \frac{1}{\sqrt{n}} \sum_{j=1}^{n} \log q_j + \nu_1 - \sum_{j=1}^{n} \lambda_j \frac{1}{\sqrt{n}} = 0,$$

where ν_1 and λ_j are Lagrange multipliers satisfying $\nu_1 \geq 0, \lambda_j \geq 0$. From the KKT condition, it can be seen that at optimality, we must have $s_1 < 1/\sqrt{n}$ and

$q_i > 0$ for all $i \in \mathcal{N}$. It then follows that $s_1 > 0$ by (2.12). Thus, any optimal $\tilde{\boldsymbol{s}}$ must be in the interior, and therefore $\frac{\partial \tilde{\pi}}{\partial s_i}(\tilde{\boldsymbol{s}}) = 0$ and $\frac{\partial^2 \tilde{\pi}}{\partial s_i^2}(\tilde{\boldsymbol{s}}) \leq 0$ must hold for any such $\tilde{\boldsymbol{s}}$. Recall that $\hat{\boldsymbol{s}}$ is an optimal solution to (P2), and $A^{-1}\hat{\boldsymbol{s}}$ is an optimal solution to (P1). Thus $A^{-1}\hat{\boldsymbol{s}}$ is also optimal for (P0). It follows that $\hat{\boldsymbol{s}}$ is also optimal to (P0$'$). Therefore, we must have $\frac{\partial \tilde{\pi}}{\partial s_i}(\hat{\boldsymbol{s}}) = 0$ and $\frac{\partial^2 \tilde{\pi}}{\partial s_i^2}(\hat{\boldsymbol{s}}) \leq 0$. The second half of the lemma follows immediately, because $\frac{\partial \pi}{\partial \boldsymbol{q}} = A\frac{\partial \tilde{\pi}}{\partial \boldsymbol{s}}$ and $\frac{\partial^2 \pi}{\partial \boldsymbol{q}^2} = A^T \frac{\partial^2 \tilde{\pi}}{\partial \boldsymbol{s}^2} A$. □

We are now ready to prove Propositions 2.3.3 and 2.3.4.

Proof of Proposition 2.3.3. By Lemma 4.2.4, we must have $\frac{\partial \pi}{\partial q_i}(\boldsymbol{q}^*) = 0$, thus (2.11) follows. Note that $2\alpha x - \log x$ is strictly convex in x, and hence for any σ, $2\alpha q - \log q = C(\sigma)$ has at most two different solutions. Therefore $\boldsymbol{q}^*$ has at most two distinct entries. □

Proof of Proposition 2.3.4. To prove that $\tilde{\pi}(s_1, \ldots, s_n)$ is supermodular in $(\boldsymbol{s}, \alpha)$, it suffices to show that all the cross partial derivatives are non-negative. We have

$$\frac{\partial^2 \tilde{\pi}}{\partial \alpha \partial s_i} = 2s_i \quad \text{for all } i$$

$$\frac{\partial^2 \tilde{\pi}}{\partial s_i \partial s_j} = \begin{cases} \frac{1}{\sqrt{nj(j-1)}} \sum_{k=1}^{j-1} (\frac{1}{q_j} - \frac{1}{q_k}), & j > i = 1 \\ \frac{1}{\sqrt{ij(i-1)(j-1)}} \sum_{k=1}^{i-1} (\frac{1}{q_i} - \frac{1}{q_k}), & j > i \geq 2. \end{cases}$$

By our assumption, $q_1 \geq \cdots \geq q_n$. Therefore, all the cross partials are non-negative and $\tilde{\pi}(s_1, \ldots, s_n)$ is supermodular in $(\boldsymbol{s}, \alpha)$.

Next we prove that the feasible set is a sublattice on $\mathbb{R}^n$. This is true because all the constraints $1/\sqrt{n} \geq s_1 \geq \sqrt{n-1}s_n, s_2 \geq 0, s_{i+1} \geq \sqrt{(i-1)/(i+1)}s_i$ are bimonotone linear inequalities, i.e., for each inequality there are at most two non-zero coefficients that are of opposite signs. The rest of the lemma follows directly from Theorem 2.8.2 of Topkis (1998). $\square$

Lemma 4.2.5. *If $\boldsymbol{q}^*$ has two distinct entries q_H and q_L with $d = q_H - q_L > 0$, then we must have $q_H = q_H(d)$ and $q_L = q_L(d)$ where*

$$q_H(d) = \frac{de^{2\alpha d}}{e^{2\alpha d} - 1}, \quad q_L(d) = \frac{d}{e^{2\alpha d} - 1}. \tag{4.5}$$

Moreover, $q_H(d)$ is increasing and $q_L(d)$ is decreasing in d with $q_H(d) > \frac{1}{2\alpha} > q_L(d)$ for all $d > 0$.

Proof. If $\boldsymbol{q}^*$ has two distinct entries q_H, q_L with $d = q_H - q_L > 0$, then (2.11) implies that $\log q_H - \log q_L = 2\alpha(q_H - q_L)$. From this we can solve for q_H and q_L to obtain (4.5).

Define $q_H(0) = \lim_{d\to 0} q_H(d) = 1/2\alpha$ and $q_L(0) = \lim_{d\to 0} q_L(d) = 1/2\alpha$. First consider $q_L(d)$. We have

$$q_L'(d) = \frac{e^{2\alpha d} - 2\alpha de^{2\alpha d} - 1}{(e^{2\alpha d} - 1)^2}.$$

Define $x = 2\alpha d$ and $g_1(x) = e^x - xe^x - 1$. To prove $q_L(d)$ is decreasing, it suffices to prove $g_1(x) < 0$ on $x > 0$. The latter condition is indeed true, because $g_1'(x) = -xe^x < 0$ and $g_1(0) = 0$. Thus $q_L(d)$ is decreasing in $d > 0$ and $q_L(d) < q_L(0) = 1/2\alpha$.

Similarly, we have

$$q_H'(d) = \frac{e^{4\alpha d} - e^{2\alpha d} - 2\alpha de^{2\alpha d}}{\left(e^{2\alpha d} - 1\right)^2}.$$

To show $q_H(d)$ is increasing, we again take $x = 2\alpha d$. It suffices to prove $g_2(x) = e^{2x} - e^x - xe^x > 0$ on $x > 0$. Because $e^x > 1 + x$ on $x > 0$, we have $g_2(x) = e^x(e^x - 1 - x) > 0$. Therefore $q_H(d)$ is increasing in $d > 0$ with $q_H(d) > q_H(0) = 1/2\alpha$. $\square$

Proof of Theorem 2.3.1. First we prove that $\boldsymbol{q}^*$ must be of the form $q_1^* \geq q_2^* = \cdots = q_n^*$. By Proposition 2.3.3, entries of $\boldsymbol{q}^*$ can take at most two distinct values. Because of the constraint $q_1 \geq \cdots \geq q_n$, we can assume

$$q_i^* = \begin{cases} q_H & i = 1, \ldots, k \\ q_L & i = k+1, \ldots, n, \end{cases}$$

where $q_H > q_L$ for some k (the case $q_1 = \cdots = q_n$ corresponds to $k = 0$). Next we will prove $k \leq 1$ by contradiction.

If $k \geq 2$, we have by (2.12) and (2.14),

$$\frac{\partial^2 \tilde{\pi}}{\partial s_k^2}(\boldsymbol{s}^*) = 2\alpha - \sum_{j=1}^{n} \frac{1}{q_j} \left(\frac{\partial q_j}{\partial s_k}\right)^2 = 2\alpha - \sum_{j=1}^{k-1} \frac{1}{k(k-1)q_j^*} - \frac{k-1}{k}\frac{1}{q_k^*} = 2\alpha - \frac{1}{q_H}.$$

By Lemma 4.2.5, $q_H > \frac{1}{2\alpha}$ and thus $\frac{\partial^2 \tilde{\pi}}{\partial s_i^2}(\boldsymbol{s}^*) > 0$, which contradicts Lemma 4.2.4. Therefore $\boldsymbol{q}^*$ must be of the form $q_1^* \geq q_2^* = \cdots = q_n^*$.

Now by Proposition 2.3.4, we know that $s_2^* = (q_1^* - q_2^*)/\sqrt{2}$ increases in α. Therefore, there is a threshold $\hat{\alpha}$ for α below which $q_1^* = \cdots = q_n^*$, and above which $q_1^* > q_2^* = \cdots = q_n^*$. Furthermore, for the case when $q_1^* = \cdots = q_n^*$, by Proposition 2.3.4 we have that $s_1^* = \sqrt{n} q_1^*$ increases in α. Therefore q_i^* increases in α for all $i \in \mathcal{N}$ in this case. Part (a) is thus proved.

When $q_1^* > q_2^* = \cdots = q_n^*$, by Lemma 4.2.5, $q_L(d)$ decreases in d and $d = q_1^* - q_2^* = \sqrt{2} s_2^*$ increases in α by Proposition 2.3.4. Furthermore, with d fixed,

$q_L = \frac{d}{e^{2\alpha d}-1}$ decreases in α. Thus q_2^* decreases in α. Moreover, by Proposition 2.3.4, $s_1^* = \sum_{i=1}^n q_i^*/\sqrt{n} = (q_H + (n-1)q_L)/\sqrt{n}$ increases in α, and hence $q_1^* = q_H$ must increase in α.

As α goes to infinity, we have $q_L < 1/2\alpha$ by Lemma 4.2.5. Therefore $q_i = q_L$ goes to 0 in the limit for all $i \geq 2$. Furthermore, by Lemma 4.2.4, we have

$$\frac{\partial \pi}{\partial q_1}(\boldsymbol{q}^*) = 2\alpha q_1^* + \log(1 - \sqrt{n}s_1^*) + y - \frac{1}{1 - \sqrt{n}s_1^*} - \log q_1^* = 0,$$

and thus $\log(1 - \sqrt{n}s_1^*) - \frac{1}{1-\sqrt{n}s_1^*} = -2\alpha q_1^* - y + \log q_1^*$. Now we have proved that q_1^* increases in α. Thus, $-2\alpha q_1^* - y + \log q_1^*$ goes to $-\infty$ as α goes to infinity. Therefore s_1^* goes to $1/\sqrt{n}$ in the limit, which also means q_1^* goes to 1. Hence, part (b) is also proved.

Proposition 4.2.6 below completes the proof of the theorem. □

Proposition 4.2.6. *Let* $R(\alpha) = y + \log(2\alpha - n) - \frac{n}{2\alpha - n}$. *There is a unique solution* α^R *to* $R(\alpha) = 0$. *Moreover,* $1/2 < \hat{\alpha} \leq \alpha^R$. *If* $n = 2$, *then* $\hat{\alpha} = \alpha^R$.

Proof. We first prove the uniqueness of α^R. First, we note that $R(\alpha) \to -\infty$ as $\alpha \downarrow n/2$ and $R(\alpha) \to \infty$ as $\alpha \to \infty$. Moreover, $R(\alpha)$ is continuous and increasing in α on $(n/2, \infty)$, so there is a unique solution α^R to $R(\alpha) = 0$. And we have

$$R(\alpha) > 0 \quad \text{for } \alpha > \alpha^R. \tag{4.6}$$

To show that $\hat{\alpha} > 1/2$, observe that when $\alpha \leq 1/2$, we must have $q_1^* = \cdots = q_n^*$. Otherwise, Lemma 4.2.5 implies that $q_H > \frac{1}{2\alpha} \geq 1$, which is a contradiction. Hence, $\hat{\alpha} > 1/2$.

Next, we will prove that $q_1^* > q_2^* = \cdots = q_n^*$ when $\alpha > \alpha^R$, from which it follows that $\hat{\alpha} \leq \alpha^R$. By the results we have proved, it suffices to rule out the possibility $q_1^* = \cdots = q_n^*$. In the ensuing argument we shall use the expressions (4.7)–(4.9) for derivatives of $\tilde{\pi}(\cdot)$, which follow from (2.12)–(2.14):

$$\frac{\partial \tilde{\pi}}{\partial s_1} = 2\alpha s_1 + \sqrt{n}(y + \log(1-\sqrt{n}s_1)) - \frac{ns_1}{1-\sqrt{n}s_1} - \sum_{j=1}^{n} \frac{1}{\sqrt{n}}(\log q_j + 1) \quad (4.7)$$

$$\frac{\partial^2 \tilde{\pi}}{\partial s_1^2} = 2\alpha - \frac{n}{1-\sqrt{n}s_1} - \frac{n}{(1-\sqrt{n}s_1)^2} - \sum_{j=1}^{n} \frac{1}{nq_j} \quad (4.8)$$

$$\frac{\partial^2 \tilde{\pi}}{\partial s_i^2} = 2\alpha - \sum_{j=1}^{i-1} \frac{1}{(i-1)iq_j} - \frac{i-1}{iq_i} \qquad \text{for } i = 2, \ldots, n. \quad (4.9)$$

Consider $\alpha > \alpha^R$ and suppose for a contradiction that $q_1^* = \cdots = q_n^* = q^*$. By (2.12) we have that $\boldsymbol{s}^*$ is given by $s_1^* = \sqrt{n}q^*$ and $s_2^* = \cdots = s_n^* = 0$. By Lemma 4.2.4 and (4.9), we have $0 \geq \frac{\partial^2 \tilde{\pi}}{\partial s_i^2}(\boldsymbol{s}^*) = 2\alpha - \frac{1}{q^*}$ for $i = 2, \ldots, n$. Therefore, $q^* \leq \frac{1}{2\alpha}$.

Next we show $q^* > 1/3n$. From (4.7) we obtain

$$\frac{\partial \tilde{\pi}}{\partial s_1}(\boldsymbol{s}^*) = \sqrt{n}\left(2\alpha q^* + y + \log(1-nq^*) - \frac{1}{1-nq^*} - \log q^*\right) = \sqrt{n}\varphi(q^*), \quad (4.10)$$

where we define $\varphi(q) = 2\alpha q + y + \log(1-nq) - \frac{1}{1-nq} - \log q$. Note that α^R must satisfy $2\alpha^R - n > 0$, and therefore $2\alpha > n$ because we are considering an $\alpha > \alpha^R$. Thus

$$\frac{\partial \tilde{\pi}}{\partial s_1}(\boldsymbol{s}^*) > \sqrt{n}\left(nq^* + \log(1-nq^*) - \frac{1}{1-nq^*} - \log q^*\right) = \sqrt{n}\widehat{\varphi}(q^*),$$

where we define $\widehat{\varphi}(q) = nq + \log(1-nq) - \frac{1}{1-nq} - \log q$. We claim $\widehat{\varphi}(q) > 0$ when $q \leq 1/3n$. This is true because $\widehat{\varphi}(1/3n) = \log 2n - 7/6 > 0$ and $\widehat{\varphi}'(q) = n - 1/q(1-nq)^2$ is negative for $0 < q < 1/3n$. Therefore, by Lemma 4.2.4, we must have $q^* > 1/3n$.

Consider now the second derivative of $\tilde{\pi}$ with respect to s_1 evaluated at $\boldsymbol{s}^*$. By (4.8) and Lemma 4.2.4 we have

$$0 \;\geq\; \frac{\partial^2 \tilde{\pi}}{\partial s_1^2}(\boldsymbol{s}^*) \;=\; 2\alpha - \frac{1}{q^*(1-nq^*)^2} \;=\; \psi(q^*),$$

where we define $\psi(q) = 2\alpha - \frac{1}{q(1-nq)^2}$. It is easy to verify that $\psi(q)$ decreases in q for $q \geq 1/3n$, thus $\psi(q) \leq 0$ for all $q \in [q^*, \frac{1}{2\alpha}]$. Note also that $\psi(q) = \varphi'(q)$, and hence $\varphi'(q) \leq 0$ for all $q \in [q^*, \frac{1}{2\alpha}]$. It follows immediately that $\varphi(q^*) \geq \varphi(1/2\alpha)$. Therefore,

$$\sqrt{n}\varphi(q^*) \geq \sqrt{n}\varphi(1/2\alpha) = \sqrt{n}\left(y + \log(2\alpha - n) - \frac{n}{2\alpha - n}\right) = \sqrt{n}R(\alpha) > 0,$$

where the final inequality follows from (4.6). Combining the preceding with (4.10), we see that $\frac{\partial \tilde{\pi}}{\partial s_1}(\boldsymbol{s}^*) > 0$, which is a contradiction with Lemma 4.2.4. Thus we have proved that $q_1^* > q_2^* = \cdots = q_n^*$ for $\alpha > \alpha^R$. Hence, $\hat{\alpha} \leq \alpha^R$.

To complete the proof it remains only to establish that $\hat{\alpha} = \alpha^R$ when $n = 2$. We just proved that $\hat{\alpha} \leq \alpha^R$ for any n. Hence, it suffices to show that for $n = 2$, $q_1^* = q_2^*$ at $\alpha = \alpha^R$.

Suppose $n = 2$ and $\alpha = \alpha^R$. By (2.12) and (4.7), the first-order condition for s_1 is

$$\begin{aligned}\frac{\partial \tilde{\pi}}{\partial s_1} &= \frac{1}{\sqrt{2}}\left(2\alpha s + 2y + 2\log 2 + 2\log(1-s) - \frac{2}{1-s} - \log(s+d) - \log(s-d)\right) \\ &= 0 \end{aligned} \tag{4.11}$$

where $s = q_1 + q_2$ and $d = q_1 - q_2$. The feasible region for (P2) is $0 \leq d \leq s \leq 1$. From the definition of α^R, we have $2y + 2\log 2 = -2\log(\alpha - 1) + \frac{2}{\alpha - 1}$ at $\alpha = \alpha^R$. By Proposition 2.3.3, if $q_1^* - q_2^* = d > 0$, then $s = s(d)$ where $s(d)$ is defined as $s(d) = q_H(d) + q_L(d) = \frac{d\exp(2\alpha d) + d}{\exp(2\alpha d) - 1}$. We also let $s(0) = \lim_{d \to 0} s(d) = 1/\alpha$. With this we can re-write the first order condition (4.11) without the leading $1/\sqrt{2}$ as $f(d) = 0$ where

$$f(d) = 2\alpha s(d) - 2\log(\alpha - 1) + \frac{2}{\alpha - 1} + \\ 2\log(1 - s(d)) - \frac{2}{1 - s(d)} - \log(s(d) + d) - \log(s(d) - d).$$

Note that $f(0) = 0$. To establish that $q_1^* = q_2^*$, it is sufficient to verify that $f(d) \neq 0$ on $(0, \bar{d}]$ where $\bar{d}$ is such that $s(\bar{d}) = 1$.

With Lemma 4.2.7 below, we know that $f(d) \neq 0$ on $(0, \bar{d}]$ and thus the proposition holds. □

Lemma 4.2.7. *$f(d)$ defined in the proof of Proposition 4.2.6 is strictly decreasing on $[0, \bar{d})$.*

Proof. We first remove the terms that do not depend upon d in $f(d)$ and define

$$g(d) = 2\alpha s(d) + 2\log(1 - s(d)) - \frac{2}{1 - s(d)} - \log(s(d)^2 - d^2).$$

Now it suffices to prove $g(d)$ decreases on $[0, \bar{d})$. To do so, we write $g(d) = 2g_1(d) + g_2(d)$ where

$$g_1(d) = \alpha s(d) - \frac{1}{1 - s(d)}, \quad g_2(d) = 2\log(1 - s(d)) - \log(s(d)^2 - d^2).$$

It suffices to prove that both $g_1(d)$ and $g_2(d)$ are decreasing on $[0, \bar{d})$.

We first consider $g_1(d)$. We have $g_1'(d) = (\alpha - 1/(1 - s(d))^2)s'(d)$. We claim that $s(d)$ is increasing on $d > 0$. To prove this, we define $x = \exp(2\alpha d)$. Then $s(d) = \bar{s}(x) = \frac{(x+1)\log x}{2\alpha(x-1)}$. Differentiating, we get $\bar{s}'(x) = \frac{x - 1/x - 2\log x}{2\alpha(x-1)^2}$. The numerator $x - 1/x - 2\log x$ is 0 at $x = 1$; taking the derivative of this expression yields $1 + 1/x^2 - 2/x = (1 - 1/x)^2$. Thus the numerator of $\bar{s}'(x)$ is zero at $x = 1$ and strictly positive for $x > 1$, and thus the claim is proved. Therefore, $s(d) > s(0) = 1/\alpha$ when $d > 0$. For $d \in (0, \bar{d})$ we have $0 < 1 - s(d) < 1 - 1/\alpha$ and therefore

$$\begin{aligned} g_1'(d) < (\alpha - 1/(1-1/\alpha)^2)s'(d) &= (\alpha - \alpha^2/(\alpha-1)^2)s'(d) \\ &= \frac{\alpha s'(d)}{(1-\alpha)^2}(\alpha^2 - 3\alpha + 1). \end{aligned}$$

From the above, $g_1'(d) < 0$ if $\alpha^R \in (1, \frac{3+\sqrt{5}}{2})$. (Recall we have taken $\alpha = \alpha^R$.) It is evident that $\alpha^R > 1$ from the definition of $R(\cdot)$. Moreover $R(\cdot)$ increases in y, so $\alpha^R < \alpha_0 = 2.18 < \frac{3+\sqrt{5}}{2}$ where α_0 satisfies $0 + \log(2\alpha_0 - 2) - \frac{1}{\alpha_0 - 1} = 0$ (i.e., $R(\alpha_0) = 0$ when $y = 0$). Hence, $g_1'(d) < 0$.

Next we consider $g_2(d)$. We have

$$\begin{aligned} g_2(d) &= 2\log(1 - s(d)) - \log(s(d)^2 - d^2) \\ &= 2\log\left(\frac{e^{\alpha d} - e^{-\alpha d}}{2d} - \frac{e^{\alpha d} - e^{-\alpha d}}{2}\right) \doteq 2\log h(d). \end{aligned}$$

Therefore, in order to prove that $g_2(d)$ is decreasing, it suffices to show that $h(d)$ is decreasing in d on $[0, \bar{d})$. We take the derivative, and we have

$$h'(d) = \frac{1}{2d^2 e^{\alpha d}}\left\{e^{2\alpha d}(-\alpha d^2 + \alpha d - 1) + \alpha d^2 + \alpha d + 1\right\}.$$

Now we want to show that $h'(d) \le 0$. For this, it suffices to show that the numerator is less than 0. Denote the numerator by $h_1(d)$. We have $h_1(0) = 0$, and $h_1'(d) = \alpha e^{2\alpha d}(-2d - 2\alpha d^2 + 2\alpha d - 1) + 2\alpha d + \alpha$. Thus $h_1'(0) = 0$. Now it suffices to show that $h_1''(d) \le 0$ for all $d \in [0, \bar{d})$. Taking another derivative, we get $h_1''(d) = \alpha e^{2\alpha d}(-2 - 8\alpha d - 4\alpha^2 d^2 + 4\alpha^2 d) + 2\alpha$. We have $h_1''(0) = 0$ so it suffices to show that $h_1'''(d) \le 0$ for all $d \in [0, \bar{d})$. Taking yet another derivative gives us $h_1'''(d) = \alpha e^{2\alpha d}(-4\alpha(3-\alpha) - 8\alpha^2 d(3-\alpha) - 8\alpha^3 d^2)$. Since $0 < \alpha^R \le 2.18$ when $n = 2$, $h_1'''(d) < 0$. Thus we have proved that $g_2(d)$ is decreasing, which completes the proof. $\square$

Proof of Theorem 2.3.2. First we prove part (a). When $\alpha \leq \hat{\alpha}$, we have $q_1^* = \cdots = q_n^* = q^*$ by Theorem 2.3.1. By (2.7), the entries of the optimal price vector $\boldsymbol{p}^*$ must also be identical and given by

$$p(q^*) = \alpha q^* - \log q^* + \log(1 - nq^*) + y. \tag{4.12}$$

Furthermore, with the condition $q_1 = \cdots = q_n = q^*$ and (2.11), we can see that q^* must satisfy

$$2\alpha q^* + y + \log(1 - nq^*) - \frac{1}{1 - nq^*} - \log q^* = 0.$$

Using the above equation to substitute for αq^* in (4.12), we have

$$p(q^*) = \frac{1}{2}\left(\frac{1}{1 - nq^*} + \log(1 - nq^*) - \log q^* + y\right) \text{ and } p'(q^*) = \frac{2nq^* - 1}{2q^*(1 - nq^*)^2}.$$

From the preceding expression, it can be seen that the behavior of $p(q^*)$ depends on the sign of $2nq^* - 1$. When $q^* \leq 1/2n$, $p(q^*)$ decreases in α. When $q^* \geq 1/2n$, $p(q^*)$ increases in α.

By Theorem 2.3.1, q^* monotonically increases in $\alpha \in [0, \hat{\alpha}]$. Thus if $q^* \geq 1/2n$ at $\alpha = 0$, then $p(q^*)$ increases in α; if $q^* \leq 1/2n$ at $\alpha = \hat{\alpha}$, then $p(q^*)$ decreases in α; if $q^* = 1/2n$ at $\alpha \in (0, \hat{\alpha})$, then $p(q^*)$ first decreases and then increases in α. This completes the proof of part (a).

Next we prove part (b). When $\alpha > \hat{\alpha}$, we have $q_1^* > q_2^* = \cdots = q_n^*$ by Theorem 2.3.1. Also, by (2.7), we have

$$p_1^* = \alpha q_1^* - \log q_1^* + \log(1 - s^*) + y, \quad p_2^* = \alpha q_2^* - \log q_2^* + \log(1 - s^*) + y, \tag{4.13}$$

where $s^* = q_1^* + (n-1)q_2^*$. Thus,

$$p_1^* - p_2^* = \alpha(q_1^* - q_2^*) - (\log q_1^* - \log q_2^*) = -\alpha(q_1^* - q_2^*) < 0,$$

so, $p_1^* < p_2^* = \cdots = p_n^*$. Similar to part (a), by (2.11), we have

$$\begin{aligned} 2\alpha q_1^* + \log(1-s^*) + y - \frac{1}{1-s^*} - \log q_1^* = 0, \\ 2\alpha q_2^* + \log(1-s^*) + y - \frac{1}{1-s^*} - \log q_2^* = 0. \end{aligned}$$

Substituting for αq_1^* and αq_2^* in (4.13), we obtain

$$\begin{aligned} p_2^* &= \frac{1}{2}\left(\frac{1}{1-s^*} + \log(1-s^*) + y - \log q_2^*\right), \\ p_1^* &= \frac{1}{2}\left(\frac{1}{1-s^*} + \log(1-s^*) + y - \log q_1^*\right). \end{aligned}$$

It follows that p_2^* increases in α because $\log q_2^*$ decreases in α by Theorem 2.3.1, s^* increases in α by Proposition 2.3.4, and $f(s) = \frac{1}{1-s} + \log(1-s)$ increases in s. By Theorem 2.3.1, $\lim_{\alpha\to\infty} s^* = \lim_{\alpha\to\infty} q_1^* = 1$. Consequently, $\lim_{\alpha\to\infty} p_1^* = \infty$. Therefore, part (b) is proved. $\square$

Proof of Continuity of q^* in α for $n \leq 2$. We prove the $n = 1$ scenario first. When $n = 1$, the problem becomes a one-variable optimization problem. The objective function is

$$\pi(q) = \alpha q^2 + q\big(y + \log(1-q)\big) - q\log(q).$$

The function π is jointly continuous in q and α. By Corollary A4.8 of Kreps (2012), to prove the continuity of q^* in α it is sufficient to establish that for each α there is a unique q that maximizes $\pi(q)$.

For fixed α, observe that $\pi(q)$ is continuous in q and $\pi'(0) > 0$ and $\pi(1) = -\infty$, so $\pi(q)$ has at least one maximizer on $(0, 1)$. Hence, the argument will be complete if we can show that for fixed α, there is at most one maximizer of $\pi(q)$. We do this next by showing that for fixed α, there is at most one local maximizer of $\pi(q)$.

The first order and the second order derivatives of π are:

$$\pi'(q) = 2\alpha q + y + \log(1-q) - \frac{1}{1-q} - \log(q), \quad \pi''(q) = 2\alpha - \frac{1}{q(1-q)^2}.$$

It is easy to prove that $\frac{1}{q(1-q)^2}$ achieves its minimum at $q = 1/3$ and the minimal value is $\frac{27}{4}$. Therefore if $\alpha < 27/8$, then $\pi''(q) < 0$ always holds and $\pi(q)$ is strictly concave. Thus, there is at most one local maximum when $\alpha < 27/8$.

We next consider $\alpha \geq 27/8$. In that case we have (recall that $y \geq 0$)

$$\pi'(q) = 2\alpha q + y + \log(1-q) - \frac{1}{1-q} - \log(q) \geq \frac{27}{4}q + \log(1-q) - \frac{1}{1-q} - \log(q).$$

Let $f(q) = \frac{27}{4}q + \log(1-q) - \frac{1}{1-q} - \log(q)$ denote the expression on the right side of the above inequality. Observe that $f(1/3) = 3/4 + \log 2 > 0$ and $f'(q) = \frac{27}{4} - \frac{1}{q(1-q)^2} \leq 0$ on $(0, 1/3]$. Hence, $f(q) > 0$ for $q \in (0, 1/3]$. Therefore, $\pi'(q) > 0$ for $q \in (0, 1/3]$.

For $q \in [1/3, 1)$, recall from above that $\pi''(q) = 2\alpha - \frac{1}{q(1-q)^2}$. It is easy to prove that $\pi''(q) > 0$ on $[1/3, q')$ and $\pi''(q) < 0$ on $(q', 1)$, where q' is the unique solution to $2\alpha = \frac{1}{q(1-q)^2}$ on $[1/3, 1)$. Therefore $\pi'(q)$ will either first increase and then decrease, or strictly decrease on $[1/3, 1)$. Combining this with the discussion for $q \in (0, 1/3)$, we see that there is at most one point where $\pi'(q) = 0$ on $(0, 1)$. Therefore there is at most one local maximum.

We have established that for fixed α, there is at most one local maximum of $\pi(q)$. Hence, we are done for $n = 1$.

Next we consider the $n = 2$ scenario. By Theorem 2.3.1, we have $\hat{\alpha} = \alpha^R$. For $\alpha = \hat{\alpha}$, (2.11) implies that $\boldsymbol{q}^* = (q^*, q^*)$ must satisfy

$$2\hat{\alpha}q^* + \log(1 - 2q^*) + y - \frac{1}{1 - 2q^*} - \log q^* = 0. \tag{4.14}$$

From the definition of α^R, we have $y = 1/(\hat{\alpha} - 1) - \log(2\hat{\alpha} - 2)$. Now we claim that the unique solution to (4.14) is $q^* = 1/(2\hat{\alpha})$.

First, it is easy to see that $1/(2\hat{\alpha})$ is indeed a solution to (4.14). Next we show that the left hand side of (4.14) is strictly decreasing in q^*, thus the solution must be unique. Let $l(q) = 2\hat{\alpha}q + \log(1 - 2q) - \frac{1}{1-2q} - \log q$. We have

$$l'(q) = 2\hat{\alpha} - \frac{1}{q(1 - 2q)} - \frac{2}{(1 - 2q)^2} < 2\hat{\alpha} - \frac{1}{q(1 - 2q)} \leq 2(\hat{\alpha} - 4)$$

where the last inequality is because $q(1 - 2q) \leq 1/8$.

Now it remains to show that $\hat{\alpha} \leq 4$. We note that given y, the function $R(\alpha)$ is increasing in α, therefore, $\hat{\alpha}$ is decreasing in y. Furthermore, when $y = 0$, $R(4) = \log 6 - 1/3 > 0$, therefore, it must hold for all y that $\hat{\alpha} < 4$.

Finally, by Lemma 4.2.5, $q_H(d) = \frac{de^{2\alpha d}}{e^{2\alpha d}-1}$ and $q_L(d) = \frac{d}{e^{2\alpha d}-1}$ when $d > 0$. Note that $\lim_{d\to 0} q_H(d) = \lim_{d\to 0} q_L(d) = 1/(2\alpha)$, which is the same as q^* at $\alpha = \hat{\alpha}$. Therefore the continuity is proved. □

Proofs for Section 2.4

Proof of Lemma 2.4.1. Consider problem (P0), it is easy to see that at optimal $\boldsymbol{q}^\dagger$ we must have $1 - \sum_{j=1}^n q_j^\dagger > 0$ and $q_i^\dagger > 0$ for all $i \in \mathcal{N}$. Therefore $\boldsymbol{q}^\dagger$ must satisfy the first-order necessary condition, i.e.,

$$\frac{\partial \pi}{\partial q_i} = \frac{2\alpha q_i}{\gamma_i} + \frac{1}{\gamma_i}\log\left(1 - \sum_{j=1}^n q_j\right) - \frac{\sum_{j=1}^n q_j/\gamma_j}{1 - \sum_{j=1}^n q_j} + \frac{y_i}{\gamma_i} - \frac{1}{\gamma_i} - \frac{\log q_i}{\gamma_i} = 0.$$

Thus (2.15) follows. $\square$

Proof of Lemma 2.4.2. We have $h''_\alpha(q) = 1/q^2 > 0$. Thus $h_\alpha(q)$ is a convex function and achieves its minimal value $h_\alpha(q) = 1 + \log(2\alpha)$ at $q = 1/2\alpha$. Furthermore, we know that $h'_\alpha(q) = 2\alpha - 1/q$. Therefore $h_\alpha(q)$ decreases on $(0, 1/2\alpha)$ and increases on $(1/2\alpha, \infty)$. The lemma is thus proved. $\square$

Proof of Proposition 2.4.3. We prove the result by contradiction. Suppose there exists an optimal solution $\boldsymbol{q}'$ in which $q'_i = \overline{q}_i^{C_i}$, $q'_j = \overline{q}_j^{C_j}$ for some $i, j \in \mathcal{N}$. Consider $\boldsymbol{q}^\epsilon$ where $q_i^\epsilon = q'_i + \epsilon$ and $q_j^\epsilon = q'_j - \epsilon$, while all the other entries remain the same as $\boldsymbol{q}'$. Define $\Delta(\epsilon) = \pi(\boldsymbol{q}^\epsilon) - \pi(\boldsymbol{q}')$. When ϵ is sufficiently small, $\boldsymbol{q}^\epsilon$ is still feasible. Because $\boldsymbol{q}'$ is optimal, $\epsilon = 0$ should be a local maximizer of $\Delta(\epsilon)$. Thus $\epsilon = 0$ should satisfy the first- and second-order necessary conditions. Taking the second-order derivative of $\Delta(\epsilon)$, we obtain $\Delta''(0) = \frac{1}{\gamma_i}(2\alpha_i - 1/q'_i) + \frac{1}{\gamma_j}(2\alpha_j - 1/q'_j)$. Since $q'_i = \overline{q}_i^{C_i}$ and $q'_j = \overline{q}_j^{C_j}$, we have $q'_i > 1/2\alpha_i$ and $q'_j > 1/2\alpha_j$ by Lemma 2.4.2. Hence $\Delta''(0) > 0$, indicating $\boldsymbol{q}'$ is not optimal. Thus we reach a contradiction and the proposition holds. $\square$

Proof of Proposition 2.4.4. The objective function (2.17) is symmetric in $(q_1, \ldots, q_n)$ except for the first term $\sum_{j=1}^n \alpha_j q_j^2$. Therefore, $q_1^\dagger > \cdots > q_n^\dagger$ because $\alpha_1 > \cdots > \alpha_n$.

Any optimal solution for (P3) must be an interior point, and hence the first-order optimality conditions are necessary. The first-order conditions are

$$\frac{\partial \pi}{\partial q_i} = 2\alpha_i q_i - \log q_i + y + \log\left(1 - \sum_{j=1}^n q_j\right) - \frac{1}{1 - \sum_{j=1}^n q_j} = 0 \text{ for all } i \in \mathcal{N}$$

and hence (2.18) follows.

For part 2, Proposition 2.4.4 part 1 and Lemma 2.4.2 with $\alpha = \alpha_1$ imply $C \geq 1 + \log(2\alpha_1)$. In addition, recall that $\alpha_1 > \cdots > \alpha_n$. For each $i \geq 2$, we have $\alpha_1 > \alpha_i$ and $h_{\alpha_1}(q) > h_{\alpha_i}(q)$ for all $q > 0$. By part 1, $h_{\alpha_1}(q_1^\dagger) = h_{\alpha_i}(q_i^\dagger) = C$. Suppose for a contradiction that $q_i^\dagger = \overline{q}_i^C$. Then we have $q_i^\dagger > 1/2\alpha_i > 1/2\alpha_1$. By Proposition 2.4.4 part 1, $q_1^\dagger \geq q_i^\dagger$, and because $h_{\alpha_i}(q)$ increases on $q > 1/2\alpha_i$, then $h_{\alpha_1}(q_1^\dagger) > h_{\alpha_i}(q_1^\dagger) \geq h_{\alpha_i}(q_i^\dagger) = C$ which contradicts $h_{\alpha_1}(q_1^\dagger) = C$. This completes the proof of part 2.

For part 3, from (2.5), we know that $p_i^\dagger - p_j^\dagger = \alpha_i q_i^\dagger - \alpha_j q_j^\dagger - (\log q_i^\dagger - \log q_j^\dagger)$ for any $i \neq j$. And from Proposition 2.4.4 part 1, we know that $\log q_i^\dagger - \log q_j^\dagger = 2(\alpha_i q_i^\dagger - \alpha_j q_j^\dagger)$. Thus $p_i^\dagger - p_j^\dagger = \alpha_i q_i^\dagger - \alpha_j q_j^\dagger - 2(\alpha_i q_i^\dagger - \alpha_j q_j^\dagger) = -(\alpha_i q_i^\dagger - \alpha_j q_j^\dagger)$. Since $q_i^\dagger > q_j^\dagger$ and $\alpha_i > \alpha_j$ for $i < j$, it follows that $\alpha_i q_i^\dagger > \alpha_j q_j^\dagger$ and therefore $p_i^\dagger < p_j^\dagger$. $\square$

Proof of Proposition 2.4.5. Similarly, (2.20) follows from the first-order condition. Next we prove $q_1^\dagger > q_2^\dagger > \cdots > q_n^\dagger$ by contradiction. Suppose $\boldsymbol{q} = (q_1, q_2, \ldots, q_n)$ is an optimal solution. In the following, we show that $q_1 > q_2$, the rest will follow from exactly the same argument. We consider another solution $\tilde{\boldsymbol{q}}$ such that $\tilde{q}_1 = q_2$, $\tilde{q}_2 = q_1$, and $\tilde{q}_i = q_i$ for $i \geq 3$. Now we consider $\pi(\boldsymbol{q}) - \pi(\tilde{\boldsymbol{q}})$, we have

$$\begin{aligned}\pi(\boldsymbol{q}) - \pi(\tilde{\boldsymbol{q}}) = &\left(\frac{1}{\gamma_1} - \frac{1}{\gamma_2}\right) \times \\ &\left(\alpha(q_1^2 - q_2^2) + (q_1 - q_2)\left(y + \log\left(1 - \sum_{j=1}^n q_j\right)\right) + q_2 \log q_2 - q_1 \log q_1\right).\end{aligned}$$

Since $\boldsymbol{q}$ is optimal, $\pi(\boldsymbol{q}) - \pi(\tilde{\boldsymbol{q}}) \geq 0$, and therefore,

$$\alpha(q_1^2 - q_2^2) + (q_1 - q_2)\left(y + \log\left(1 - \sum_{j=1}^n q_j\right)\right) + q_2 \log q_2 - q_1 \log q_1 \geq 0.$$

Also by (2.20), $2\alpha q_1 - \log q_1 < 2\alpha q_2 - \log q_2$, therefore $-q_1 \log q_1 < 2\alpha q_1 q_2 - 2\alpha q_1^2 - q_1 \log q_2$. Thus, we must have

$$\begin{aligned}
0 &\leq \alpha(q_1^2 - q_2^2) + (q_1 - q_2)(y + \log(1 - \sum_{j=1}^{n} q_j)) + q_2 \log q_2 - q_1 \log q_1 \\
&< -\alpha(q_1 - q_2)^2 + (q_1 - q_2)(y + \log(1 - \sum_{j=1}^{n} q_j)) + (q_2 - q_1)\log q_2 \\
&= (q_1 - q_2)(-\alpha(q_1 - q_2) + y + \log(1 - \sum_{j=1}^{n} q_i) - \log q_2).
\end{aligned}$$

Again by (2.18), we have

$$\begin{aligned}
&\alpha(q_2 - q_1) + y + \log(1 - \sum_{j=1}^{n} q_j) - \log q_2 \\
&= \alpha(q_2 - q_1) + 1 - 2\alpha q_2 + \frac{\gamma_2 \sum_{j=1}^{n} q_j/\gamma_j}{1 - \sum_{j=1}^{n} q_j} \\
&= -\alpha(q_1 + q_2) + 1 + \frac{\gamma_2 \sum_{j=1}^{n} q_j/\gamma_j}{1 - \sum_{j=1}^{n} q_j} \\
&\geq -\alpha(q_1 + q_2) + 1 + \frac{q_1 + q_2}{1 - \sum_{j=1}^{n} q_j}
\end{aligned}$$

where the last inequality holds because $\gamma_1 < \gamma_2$.

Now it is easy to see that when $1 - \sum_{j=1}^{n} q_j \leq 1/\alpha$ or $\alpha \leq 1$, the right hand side is positive. Thus $q_1 > q_2$. Now it remains to consider the case when $1 - \sum_{j=1}^{n} q_j > 1/\alpha$ and $\alpha > 1$.

In this case, we rewrite the difference $\pi(\boldsymbol{q}) - \pi(\tilde{\boldsymbol{q}})$ in the following way:

$$\begin{aligned}
&\pi(\boldsymbol{q}) - \pi(\tilde{\boldsymbol{q}}) = \\
&\left(\frac{1}{\gamma_1} - \frac{1}{\gamma_2}\right)\left(\alpha q_1^2 + q_1(y + C) - q_1 \log q_1 - (\alpha q_2^2 + q_2(y + C) - q_2 \log q_2)\right)
\end{aligned}$$

where $C = \log(1 - \sum_{j=1}^{n} q_j) > -\log \alpha$. Now define $f(x) = \alpha x^2 + x(y+C) - x \log x$. Next we show that $f(x)$ is strictly increasing in x on $[0, 1]$ for any C. If this is the case, in order for $\pi(\boldsymbol{q}) \geq \pi(\tilde{\boldsymbol{q}})$, we must have $q_1 \geq q_2$. By (2.20), $q_1 \neq q_2$. Consequently, we must have $q_1 > q_2$.

To show $f(x)$ is increasing, we have $f'(x) = 2\alpha x + y + C - 1 - \log x$. Note that this function is convex and achieves minimum on $[0, 1]$ at $x = 1/2\alpha$ (remember in this case, $\alpha > 1$). The minimum value of $f'(x)$ is $C + \log 2\alpha + y \geq \log 2 > 0$. Therefore, $f'(x) > 0$ for all x on $[0, 1]$. And thus the part 1 is proved. Part 2 follows exactly the same as in proof of Proposition 2.4.4 part 2. $\square$