

Optimal Pricing for a Multinomial Logit Choice Model with Network Effects

Chenhao Du William L. Cooper Zizhuo Wang

Department of Industrial and Systems Engineering
University of Minnesota
Minneapolis, MN 55455

Appeared in *Operations Research*, 64(2), pp. 441–455, 2016

Abstract

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 boosts the sales of the single low-price product in comparison to the sales of all other products. We also obtain comparative statics results that describe how optimal prices and sales levels vary with a parameter that determines the strength of the network effects. We extend our analysis to settings with heterogeneous products, and establish 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.

Keywords: Pricing, Choice Models, Network Effects, Revenue Management

1 Introduction

A product is said to exhibit network effects if individual consumers value it more when more other consumers purchase it. For example, in multi-player online video games, individuals may obtain greater enjoyment (and thus greater value) if there are more other players and hence more others to play with or against. The expansion of social media and internet technology has led to the emergence of new products that display network effects (e.g., multi-player online video games and group buying deals) and has also served to strengthen network effects for traditional products such as movies, television programs, and books by allowing people to easily participate in communities built around those products. In the meantime, advances in information technology have enabled firms to collect large amounts of data in hopes of better understanding how customers make choices among potential purchases. 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 paper: (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, in the online appendix 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 paper. 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 remainder of this section, we review related literature, focusing on two main streams of research: the study of multi-product pricing problems and the study of network goods. In a typical

multi-product pricing problem, a seller offers a menu of products and must determine the prices of those products. Customers choose among the products (or decide not to purchase) according to a consumer choice model. Perhaps the most widely used such model is the multinomial logit (MNL) model, comprehensive reviews of which can be found in, e.g., Anderson et al. (1992) or Ben-Akiva and Lerman (1985).

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. In a model that does not include costs, this means that all products are priced identically and have identical sales quantities. To the best of our knowledge, a version of 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. For work of this nature, see Rayfield et al. (2015) and Gallego and Wang (2014) and references therein. Some of the computational approaches we consider are related to these search methods. 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 our paper 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.

The literature dealing with products with network effects 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 Qu  rou (2012), and references therein for studies of local network effects. In our paper, 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 (1996). 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) consider 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. They consider only symmetric Nash equilibria wherein all firms necessarily set the same price. As we prove later, even in a completely homogeneous case, the revenue maximizing decisions of a single seller may involve different prices across products. Hence, the questions, conclusions, and methodology in our paper are quite distinct from those of Anderson et al. Brock and Durlauf (2002) study a multinomial-logit choice model of neighborhood effects, wherein decisions of individuals are influenced by decisions of others in their “neighborhood”. They show that there could be multiple choice equilibria and also discuss the econometric implementation of the model. Starkweather (2003) also studies pricing and product compatibility problems based upon an MNL model with network effects. He shows that there could be multiple demand equilibria. To draw distinctions with our paper, 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, again our paper is quite different from that of Starkweather. In a recent paper, Wang and Wang (2014) 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 paper.

The rest of the paper is organized as follows. Section 2 describes the MNL model with network effects. Section 3 focuses on the homogeneous-products case and contains two theorems that describe properties of optimal solutions in that case. Section 4 focuses on settings with heterogeneous products. Section 5 describes results of numerical studies. Section 6 concludes the paper. Proofs and extensions are contained in the online appendices.

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 $i \in \mathcal{N} = \{1, \dots, n\}$ 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 $\epsilon_0, \epsilon_1, \dots, \epsilon_n$ 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

$$v_i = y_i - \gamma_i p_i + \alpha_i x_i, \tag{1}$$

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 Online Appendix B.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 Online Appendix B.2 we allow a more general form of network effects by replacing $\alpha_i x_i$ in (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 Online Appendix B.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, \dots, x_n and prices p_1, \dots, p_n , then the probability that a customer purchases product $i \in \mathcal{N}$ is

$$q_i = P(u_i = \max_{j \in \{0, 1, \dots, n\}} u_j) = \frac{\exp(v_i)}{1 + \sum_{j=1}^n \exp(v_j)} = \frac{\exp(y_i - \gamma_i p_i + \alpha_i x_i)}{1 + \sum_{j=1}^n \exp(y_j - \gamma_j p_j + \alpha_j x_j)}. \quad (2)$$

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

$$x_i = M q_i. \quad (3)$$

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)–(3) now reduce to

$$q_i = F_i(\mathbf{q}) \text{ for all } i \in \mathcal{N} \quad \text{where} \quad F_i(\mathbf{q}) = \frac{\exp(y_i - \gamma_i p_i + \alpha_i q_i)}{1 + \sum_{j=1}^n \exp(y_j - \gamma_j p_j + \alpha_j q_j)}. \quad (4)$$

Note that sales quantities affect choice probabilities, which themselves affect sales quantities. Hence, the preceding expression (4) may be viewed as an equilibrium condition. In equilibrium, a vector of sales quantities $\mathbf{q} = (q_1, \dots, q_n)$ 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 \mathbf{q} . We can obtain a slightly different justification of (4) by interpreting x_1, \dots, x_n on the right side of (2) as customers' perceptions of sales levels, in which case (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 $\mathbf{p} = (p_1, \dots, p_n)$, the function $F(\mathbf{q}) = (F_1(\mathbf{q}), \dots, F_n(\mathbf{q}))$ is a continuous function of $\mathbf{q} = (q_1, \dots, q_n)$ from $[0, 1]^n$ to $[0, 1]^n$. By the Brouwer fixed point theorem, there exists at least one solution to (4). In general, given prices \mathbf{p} , there could be more than one \mathbf{q} that satisfies (4). Nevertheless, given any \mathbf{q} satisfying $q_1, \dots, q_n > 0$ and $\sum_{i \in \mathcal{N}} q_i < 1$, there is a unique $\mathbf{p} = (p_1(\mathbf{q}), \dots, p_n(\mathbf{q}))$ defined by

$$p_i(\mathbf{q}) = \frac{1}{\gamma_i} \left(\alpha_i q_i - \log q_i + \log \left(1 - \sum_{j=1}^n q_j \right) + y_i \right) \quad (5)$$

such that (4) holds for \mathbf{q} . For some \mathbf{q} , if the seller charges prices determined by (5), there may be solutions to (4) other than \mathbf{q} . We will discuss the issue of potential existence of multiple equilibria

in Section 5.1 and Online Appendix A. For now, we assume that the seller has the capability to choose the sales \mathbf{q} by implementing prices $(p_1(\mathbf{q}), \dots, p_n(\mathbf{q}))$. This allows us to use $\mathbf{q} = (q_1, \dots, q_n)$ as the decision variables.

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

$$\pi(\mathbf{q}) = \sum_{j=1}^n q_j p_j(\mathbf{q}) = \sum_{j=1}^n \frac{\alpha_j}{\gamma_j} q_j^2 + \sum_{j=1}^n \frac{q_j}{\gamma_j} \log \left(1 - \sum_{j=1}^n q_j \right) + \sum_{j=1}^n \frac{q_j}{\gamma_j} (y_j - \log q_j). \quad (6)$$

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

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

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

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 $\mathbf{q} = (q_1, \dots, q_n)$, the expression (5) for prices becomes

$$p_i(\mathbf{q}) = \alpha q_i - \log q_i + \log \left(1 - \sum_{j=1}^n q_j \right) + y, \quad (7)$$

and the total revenue (6) becomes

$$\pi(\mathbf{q}) = \alpha \sum_{j=1}^n q_j^2 + \sum_{j=1}^n q_j \left(y + \log \left(1 - \sum_{j=1}^n q_j \right) \right) - \sum_{j=1}^n q_j \log q_j. \quad (8)$$

Note that the objective function (8) is symmetric in \mathbf{q} . Thus we can assume $q_1 \geq q_2 \geq \dots \geq q_n$ in (P0) without loss of optimality. Then problem (P0) becomes

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

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 $\mathbf{q}^* = (q_1^*, \dots, q_n^*)$ as

$$\mathbf{q}^* = \{\mathbf{q} \mid \mathbf{q} \in Q^* \text{ and } \mathbf{e}^T A \mathbf{q} \geq \mathbf{e}^T A \mathbf{q}' \text{ for all } \mathbf{q}' \in Q^*\}, \tag{9}$$

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

$$A_{ij} = \begin{cases} \frac{1}{\sqrt{n}} & i = 1 \\ \frac{1}{\sqrt{(i-1)i}} & i \geq 2, i > j \\ -\frac{i-1}{\sqrt{(i-1)i}} & i \geq 2, i = j \\ 0 & \text{otherwise.} \end{cases} \tag{10}$$

Later we will show that \mathbf{q}^* defined above is unique. That is, the set on the right side of (9) contains only a single element. A key step in our analysis is to make a change of variables $\mathbf{s} = A \mathbf{q}$ in (P1). We will establish that after the change of variables, the objective function will be supermodular in (\mathbf{s}, α) . 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 \mathbf{q} -space (the feasible region of (P1)) is still a sublattice in the \mathbf{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

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

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 Online Appendix C.

Theorem 3.1. *There exists $\hat{\alpha}$ (which depends upon y and n) such that*

- (a) *if $\alpha \leq \hat{\alpha}$, then $q_1^* = q_2^* = \dots = q_n^*$ and q_i^* increases in α for all $i \in \mathcal{N}$;*
- (b) *if $\alpha > \hat{\alpha}$, then $q_1^* > q_2^* = \dots = q_n^*$ 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, $\lim_{\alpha \rightarrow \infty} q_1^* = 1$ and $\lim_{\alpha \rightarrow \infty} q_i^* = 0$ for $i = 2, \dots, n$.*

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 3.1, consider the objective function (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 \mathbf{q} while $-\sum_{j=1}^n q_j \log q_j$ is concave in \mathbf{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 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 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 4.1 that allows us to solve the problem using a one-dimensional search.

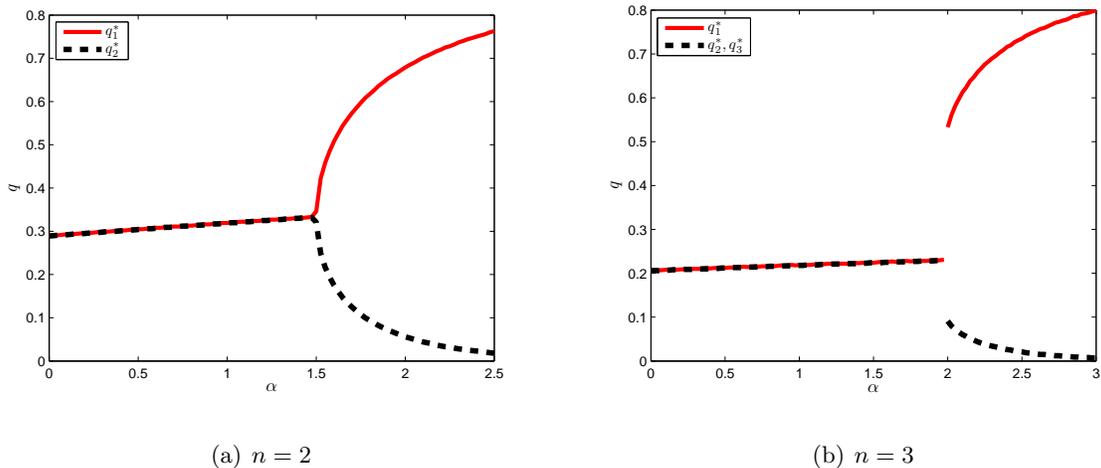


Figure 1: Optimal quantities \mathbf{q}^* as a function of α

Next, we will present two examples to illustrate the above results. In both examples we set $y = 2$. In Figure 1(a) we take $n = 2$, and in Figure 1(b) we take $n = 3$. Both figures show how the optimal \mathbf{q}^* changes with α . As one can see, when α is small, all the entries in \mathbf{q}^* are equal and increase monotonically in α . When α passes a certain threshold, q_1^* becomes the single largest entry of \mathbf{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 3.1 describes this pattern. In Figure 1(a), the threshold is $\alpha = 1.5$, which is precisely the solution to $R(\alpha) = 0$ at $y = 2, n = 2$. In Figure 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 3.1.

Observe in Figure 1(b) that \mathbf{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 $q_1 = \dots = q_n$ and of the form $q_1 > q_2 = \dots = q_n$. 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 $q_2 = \dots = q_n$.

(So we can solve (P1) by maximizing $\Pi(d)$ over d .) Figures 2(a) and 2(b) show $\Pi(d)$ plotted against d for two distinct values of α in the example from Figure 1(b). In both Figure 2(a) and Figure 2(b), one can see two peaks corresponding to the aforementioned two local optima. In Figure 2(a) in which $\alpha = 1.98$, the left peak achieves a higher objective value, and thus the entries of \mathbf{q}^* are all equal ($d = 0$). In Figure 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 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 \mathbf{q}^* is always continuous in α , and thus the discontinuity does not appear in Figure 1(a). The proof is given in Online Appendix C.)

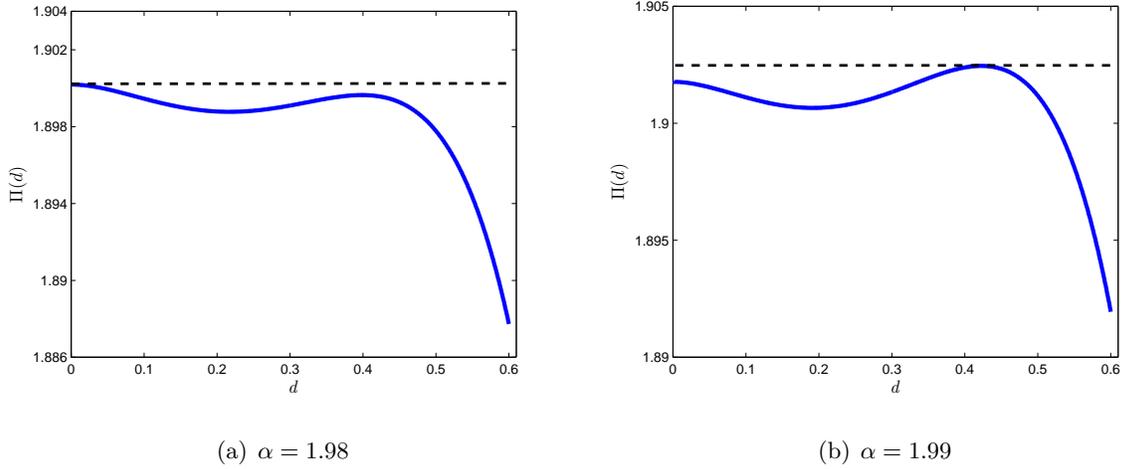
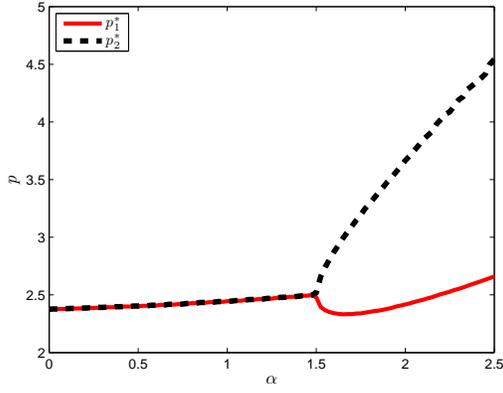


Figure 2: Two local optima of $\Pi(d)$

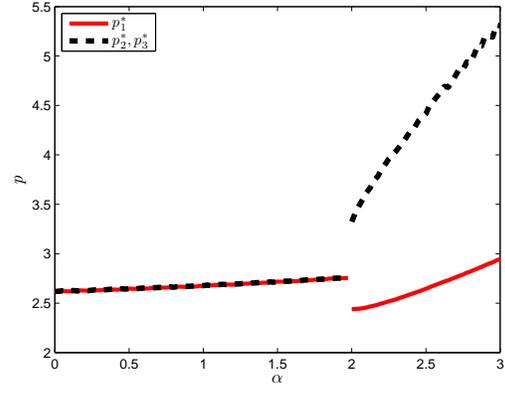
In addition to the sales \mathbf{q}^* , we are also interested in the optimal prices $\mathbf{p}^* = (p_1(\mathbf{q}^*), \dots, p_n(\mathbf{q}^*))$ induced by (7). We have the following theorem describing the behavior of \mathbf{p}^* for different values of α . The proof, which relies heavily on Theorem 3.1, is in Online Appendix C.

Theorem 3.2. *Let $\hat{\alpha}$ be defined as in Theorem 3.1.*

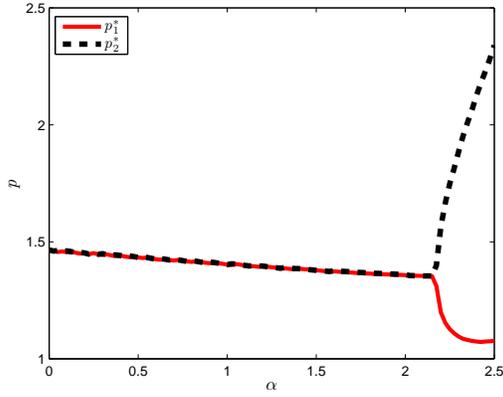
- (a) *If $\alpha \leq \hat{\alpha}$, then $p_1^* = p_2^* = \dots = p_n^* = p^*$ and one of the following three scenarios holds.*
- (i) *p^* increases monotonically in α , or*
 - (ii) *p^* decreases monotonically in α , or*
 - (iii) *p^* first decreases and then increases in α .*
- (b) *If $\alpha > \hat{\alpha}$, then $p_1^* < p_2^* = \dots = p_n^*$ and p_2^* increases in α . Moreover, $\lim_{\alpha \rightarrow \infty} p_1^* = \infty$ and $\lim_{\alpha \rightarrow \infty} p_2^* = \infty$.*



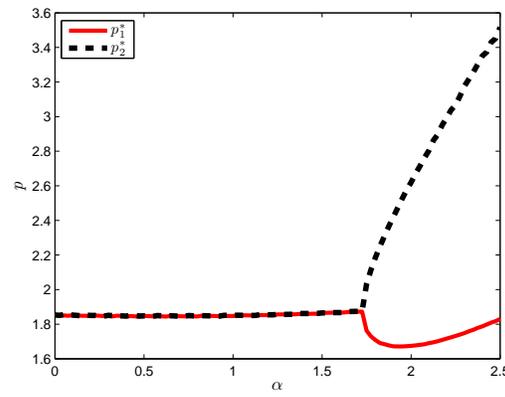
(a) $n = 2, y = 2$



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



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



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

Figure 3: Optimal prices \mathbf{p}^* as a function of α

In Figure 3, we use four examples to illustrate the behavior of \mathbf{p}^* in Theorem 3.2. Figures 3(a) and 3(b) correspond to the same two cases as Figures 1(a) and 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 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 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 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 \mathbf{p}^* in response to network effects is more complicated than that of the optimal sales quantities \mathbf{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 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 3.1. (The complete proof is provided in the online appendix.) We start with a proposition. The conditions in (11) below are derived from the first-order optimality conditions for (P1).

Proposition 3.3. *There are at most two distinct entries in $\mathbf{q}^* = (q_1^*, \dots, q_n^*)$, and*

$$2\alpha q_i^* - \log q_i^* = C \left(\sum_{j \in \mathcal{N}} q_j^* \right) \quad \text{for all } i \in \mathcal{N}, \quad (11)$$

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

To provide some insight into Theorem 3.1, a key point is that the expression on the right side in (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 \mathbf{q}^* as indicated in Theorem 3.1. In the online appendix, we show that among the entries of \mathbf{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 \mathbf{q}^* vary with α , we apply the variable transformation mentioned earlier. Let $\mathbf{s} = A\mathbf{q}$, where A is given by (10). Then $\mathbf{s} = (s_1, \dots, s_n)$ can be written in terms of \mathbf{q} as follows

$$s_i = \begin{cases} \frac{1}{\sqrt{n}} \sum_{j=1}^n q_j & i = 1 \\ \frac{1}{\sqrt{(i-1)i}} \sum_{j=1}^{i-1} q_j - \frac{i-1}{\sqrt{(i-1)i}} q_i & i = 2, \dots, n. \end{cases} \quad (12)$$

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, $\mathbf{q} = A^{-1}\mathbf{s} = A^T\mathbf{s}$, and thus q_i can be written as

$$q_i = \begin{cases} \frac{s_1}{\sqrt{n}} + \sum_{j=2}^n \frac{s_j}{\sqrt{(j-1)j}} & i = 1 \\ \frac{s_1}{\sqrt{n}} - \frac{(i-1)s_i}{\sqrt{(i-1)i}} + \sum_{j=i+1}^n \frac{s_j}{\sqrt{(j-1)j}} & i = 2, \dots, n. \end{cases} \quad (13)$$

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

$$\tilde{\pi}(s_1, \dots, s_n) = \alpha \sum_{i=1}^n s_i^2 + \sqrt{n} s_1 (y + \log(1 - \sqrt{n} s_1)) - \sum_{i=1}^n q_i \log q_i, \quad (14)$$

where q_i is defined in (13). The constraints of (P1) are equivalent to

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

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

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

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

Proposition 3.4. *$\tilde{\pi}(s_1, \dots, s_n)$ is supermodular in (\mathbf{s}, α) and the feasible region of (P2) is a sublattice in \mathbb{R}^n . There exists an entry-wise maximal optimal solution $\mathbf{s}^* = (s_1^*, \dots, s_n^*)$ to (P2) such that $s_i^* \geq s_i$ for $i = 1, \dots, n$ for any other optimal $\mathbf{s} = (s_1, \dots, s_n)$. Moreover, \mathbf{s}^* increases monotonically in α .*

The uniqueness of \mathbf{q}^* in (9) follows from Proposition 3.4. To see this, observe that if there are multiple optimal \mathbf{q} for (P1), then (9) picks among them by selecting those for which the sum of the entries of $\mathbf{s} = A\mathbf{q}$ is largest. The proposition establishes that there is an entry-wise maximal optimal solution \mathbf{s}^* to (P2). Of course, \mathbf{s}^* must also have the property that the sum of its entries is strictly greater than that of any other optimal \mathbf{s} . So \mathbf{q}^* is unique and $\mathbf{q}^* = A^{-1}\mathbf{s}^*$.

Using the monotonicity properties of \mathbf{s}^* from the proposition, we can prove the monotonicity properties of \mathbf{q}^* that are stated in Theorem 3.1. Details appear in Online Appendix C.

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 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 3.1 and 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 4.1. *Any optimal solution $\mathbf{q}^\dagger = (q_1^\dagger, \dots, q_n^\dagger)$ to (P0) must satisfy*

$$2\alpha_i q_i^\dagger - \log q_i^\dagger = C_i(\mathbf{q}^\dagger) \quad \text{for all } i \in \mathcal{N}, \quad (15)$$

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

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 (15).

Lemma 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, $\underline{q}^c < \bar{q}^c$. Furthermore, $\underline{q}^c \in (0, 1/2\alpha)$ and \underline{q}^c decreases in c , while $\bar{q}^c \in (1/2\alpha, \infty)$ and \bar{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 Online Appendix D. For $i \in \mathcal{N}$ and $c \geq 1 + \log(2\alpha_i)$, define $\underline{q}_i^c \leq \bar{q}_i^c$ to be the solutions to $h_{\alpha_i}(q) = c$. (For $c = 1 + \log(2\alpha_i)$, we have $\underline{q}_i^c = \bar{q}_i^c = 1/2\alpha_i$.) It is easy to find \underline{q}_i^c and \bar{q}_i^c 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 4.3. *Any optimal solution \mathbf{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) $q_i^\dagger = \bar{q}_i^{C_i}$ and $q_j^\dagger = \underline{q}_j^{C_j}$ for all $j \in \mathcal{N} \setminus \{i\}$ for some single $i \in \mathcal{N}$

where $C_k = C_k(\mathbf{q}^\dagger)$ for $k \in \mathcal{N}$ and $C_k(\cdot)$ is defined in Lemma 4.1.

From the proposition, we see that for an optimal solution, there is at most one sales quantity that takes the ‘‘high value’’ $\bar{q}_i^{C_i}$, 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 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 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 (6) can be rewritten as $\pi(\mathbf{q}) = \sum_{j=1}^n \pi_j(q_j)$ 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 \mathbf{q} there are two entries q_i and q_j that lie in the regions of convexity. Consider $\mathbf{q}^\epsilon = \mathbf{q} + \epsilon \mathbf{e}_{ij}$ and $\mathbf{q}^{-\epsilon} = \mathbf{q} - \epsilon \mathbf{e}_{ij}$ where \mathbf{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 $\pi(\mathbf{q}^\epsilon) + \pi(\mathbf{q}^{-\epsilon}) > 2\pi(\mathbf{q})$, implying at least one of \mathbf{q}^ϵ and $\mathbf{q}^{-\epsilon}$ performs better than \mathbf{q} , which is a contradiction with the optimality of \mathbf{q} . A more rigorous proof of Proposition 4.3 is provided in Online Appendix D.

In view of the non-concavity of the objective function (6), it is important to develop a specialized computational approach for (P0). We take this up next. Given a $\mathbf{c} = (c_1, \dots, c_n)$, consider the vectors \mathbf{q} for which

$$2\alpha_i q_i - \log q_i = c_i \quad \text{for all } i \in \mathcal{N}. \quad (16)$$

Note that there are at most 2^n such \mathbf{q} . For each such \mathbf{q} , if $C_i(\mathbf{q}) = c_i$ for all $i \in \mathcal{N}$ then we have a \mathbf{q} that satisfies (15) and that is an initial candidate to be an optimal solution to (P0). A key insight is that by Proposition 4.3, given a \mathbf{c} , we can a priori eliminate all but $(n + 1)$ of the vectors \mathbf{q} that satisfy (16). By searching through vectors \mathbf{c} that can potentially appear on the right side of (15), we arrive at a set of “non-eliminated” candidate solutions \mathbf{q} each of which satisfies (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 \mathbf{c} is an n -dimensional vector. However, it turns out that the search through \mathbf{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 $C_i(\mathbf{q}^\dagger)$ in (15) is given by $C_i(\mathbf{q}^\dagger) = \gamma_i K_2 / (1 - K_1) - \log(1 - K_1) + 1 - y_i$. Hence, it suffices to search over values

of K_1 and K_2 to find sales vectors that satisfy the necessary condition (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 paper. We summarize the above procedures in Algorithm 1. After running the algorithm, we can also compute the optimal prices using (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}^{c_i} and \bar{q}^{c_i} that satisfy $2\alpha_i q - \log q = c_i$ for $i \in \mathcal{N}$.
 - (c) Let $\underline{\mathbf{q}} = (\underline{q}_1^{c_1}, \underline{q}_2^{c_2}, \dots, \underline{q}_n^{c_n})$ and for $i = 1, \dots, n$ let $\bar{\mathbf{q}}^i = (\underline{q}_1^{c_1}, \dots, \underline{q}_{i-1}^{c_{i-1}}, \bar{q}_i^{c_i}, \underline{q}_{i+1}^{c_{i+1}}, \dots, \underline{q}_n^{c_n})$. Here, $\bar{\mathbf{q}}^i$ is the same as $\underline{\mathbf{q}}$ except that the i -th entry is replaced by $\bar{q}_i^{c_i}$. Let $R = \{\underline{\mathbf{q}}, \bar{\mathbf{q}}^1, \bar{\mathbf{q}}^2, \dots, \bar{\mathbf{q}}^n\}$.
 - (d) For each $\mathbf{q} \in R$: if both $K_1 = \sum_{j=1}^n q_j$ and $K_2 = \sum_{j=1}^n q_j / \gamma_j$ hold then let $Q = Q \cup \{\mathbf{q}\}$.
 3. End For-Loop over (K_1, K_2) .
 4. Evaluate (6) at each element in Q to find \mathbf{q}^\dagger that maximizes (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.

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 $\alpha_1 > \dots > \alpha_n$, 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 (6)

becomes

$$\pi(\mathbf{q}) = \sum_{j=1}^n \alpha_j q_j^2 + \sum_{j=1}^n q_j \left(y + \log \left(1 - \sum_{j=1}^n q_j \right) \right) - \sum_{j=1}^n q_j \log q_j. \quad (17)$$

We refer to the revenue maximization problem (P0) with the objective function specialized to (17) as (P3). The following proposition summarizes properties of the optimal solution to (P3). The proof is in Online Appendix D.

Proposition 4.4. *For any optimal solution $\mathbf{q}^\dagger = (q_1^\dagger, \dots, q_n^\dagger)$ to (P3)*

1. \mathbf{q}^\dagger must satisfy $q_1^\dagger > \dots > q_n^\dagger$ and

$$2\alpha_i q_i^\dagger - \log q_i^\dagger = C \left(\sum_{j \in \mathcal{N}} q_j^\dagger \right) \quad \text{for all } i \in \mathcal{N}, \quad (18)$$

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

2. For \mathbf{q}^\dagger , we have $C \geq 1 + \log(2\alpha_1)$, where $C := C \left(\sum_{j \in \mathcal{N}} q_j^\dagger \right)$. In addition, $q_1^\dagger \in \{\underline{q}_1^C, \bar{q}_1^C\}$ and $q_i^\dagger = \underline{q}_i^C$ for $i = 2, \dots, n$.
3. The optimal prices $\mathbf{p}^\dagger = (p_1^\dagger, \dots, p_n^\dagger) = (p_1(\mathbf{q}^\dagger), \dots, p_n(\mathbf{q}^\dagger))$ obtained from \mathbf{q}^\dagger and (5) must satisfy $p_1^\dagger < \dots < p_n^\dagger$.

From the proposition, we see that an optimal solution to problem (P3) follows the similar structure as in Proposition 4.3, i.e., there is at most one entry in \mathbf{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 (18) depends just upon the sum of the entries of \mathbf{q} . Proposition 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 \mathbf{q} that satisfy $2\alpha_i q_i - \log q_i = c$ are left to be evaluated. Specifically, we only need to consider the \mathbf{q} where $q_i = \underline{q}_i^c$ for $i = 2, \dots, n$ and either (a) $q_1 = \underline{q}_1^c$ or (b) $q_1 = \bar{q}_1^c$. Algorithm 2 below outlines an approach to solve (P3) based upon these observations.

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 \bar{q}_1^c and \underline{q}_i^c for $i \in \mathcal{N}$ that satisfy $2\alpha_i q - \log q = c$.
 - (c) Let $\underline{\mathbf{q}} = (\underline{q}_1^c, \underline{q}_2^c, \dots, \underline{q}_n^c)$ and $\bar{\mathbf{q}}^1 = (\bar{q}_1^c, \underline{q}_2^c, \dots, \underline{q}_n^c)$. Let $R = \{\underline{\mathbf{q}}, \bar{\mathbf{q}}^1\}$.
 - (d) For each $\mathbf{q} \in R$: if $C(\sum_{j \in \mathcal{N}} q_j) = c$ holds then let $Q = Q \cup \{\mathbf{q}\}$.
 2. End For-Loop over σ .
 3. Evaluate (17) at each element in Q to find \mathbf{q}^\dagger that maximizes (6).
-

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 $\gamma_1 < \dots < \gamma_n$. In this setting, the revenue function (6) becomes

$$\pi(\mathbf{q}) = \alpha \sum_{j=1}^n \frac{q_j^2}{\gamma_j} + \sum_{j=1}^n \frac{q_j}{\gamma_j} \left(y + \log \left(1 - \sum_{i=1}^n q_i \right) \right) - \sum_{j=1}^n \frac{q_j}{\gamma_j} \log q_j. \quad (19)$$

We refer to (P0) with the objective function specialized to (19) as (P4). The next proposition describes the optimal solution to (P4). The proof is in Online Appendix D.

Proposition 4.5. *For any optimal solution $\mathbf{q}^\dagger = (q_1^\dagger, \dots, q_n^\dagger)$ to (P4)*

1. \mathbf{q}^\dagger must satisfy $q_1^\dagger > \dots > q_n^\dagger$ and

$$2\alpha q_i^\dagger - \log q_i^\dagger = C_i(\mathbf{q}^\dagger) \quad \text{for all } i \in \mathcal{N}, \quad (20)$$

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

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

The preceding structure is similar to what is described in Proposition 4.4, and is also useful for simplifying the computation of \mathbf{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 4.5 can be easily modified to accommodate heterogeneous intrinsic utilities that are ordered in the direction opposite of the γ_i , i.e., $y_1 \geq y_2 \geq \dots \geq y_n$.

The computations in the case of heterogeneous price sensitivities considered in this section are more difficult than those in subsection 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 \bar{q}^{c_1} and \underline{q}^{c_i} that satisfy $2\alpha q - \log q = c_i$ for $i \in \mathcal{N}$.

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

5 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.

5.1 The Issue of Multiple Equilibria

As discussed earlier, one issue with the network MNL model is that given a price vector \mathbf{p} , there are possibly multiple \mathbf{q} that satisfy the equilibrium condition (4). We are particularly interested in this issue at prices $\mathbf{p}^* = \mathbf{p}(\mathbf{q}^*)$, which are the optimal prices obtained from our model. Online Appendix A addresses the question of uniqueness and stability of solutions to (4). In this section, we study whether the sales vector will converge to \mathbf{q}^* when prices are set at \mathbf{p}^* if customers repeatedly adjust their purchases according to market conditions or their perceptions thereof. Such issues are of interest even if \mathbf{q}^* is the unique equilibrium at prices \mathbf{p}^* . In addition, if we find that there is convergence to \mathbf{q}^* from all tested starting points of the adjustment process, then this provides some

evidence that \mathbf{q}^* is the unique equilibrium (or at least the unique stable equilibrium) associated with prices \mathbf{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 \mathbf{q}^* . Then we calculate $\mathbf{p}^* = \mathbf{p}(\mathbf{q}^*)$ from (5).
2. From a starting sales vector \mathbf{q}^0 , we iteratively compute \mathbf{q}^t using the following dynamics:

$$q_i^t = \frac{\exp(y_i - \gamma_i p_i^* + \alpha_i q_i^{t-1})}{1 + \sum_{j=1}^n \exp(y_j - \gamma_j p_j^* + \alpha_j q_j^{t-1})}. \quad (21)$$

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

The dynamics in (21) 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 \mathbf{q}^0 , which would suggest to the customers that sales levels would be \mathbf{q}^1 . Realizing this, customers would then hypothesize sales levels \mathbf{q}^1 , which would then suggest that sales levels would be \mathbf{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 \mathbf{q} . That vector will be an equilibrium (that is, satisfy (4)) for prices \mathbf{p}^* . If the seller initially announces that sales levels will be \mathbf{q}^* and if the population believes the seller so that $\mathbf{q}^0 = \mathbf{q}^*$, then the iterative procedure will converge immediately to quantities \mathbf{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 A.2 in Online Appendix A, the sales vector \mathbf{q}^t must converge to \mathbf{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 $n = 1, \dots, 5$ and for each n , we test 1000 randomly generated problems. To generate a problem, we sample α_i , y_i and γ_i for $i = 1, \dots, n$ from uniform distributions over $[0, 4]$, $[0, 2]$, and $[1, 2]$ respectively. In this first set of experiments, we focus on the starting point $\mathbf{q}^0 = \mathbf{0} := (0, \dots, 0)$, which is a natural choice if initially no customers buy any of the products. For each n , we compute the average of $\|\mathbf{q}^t - \mathbf{q}^*\| = (\sum_{i=1}^n (q_i^t - q_i^*)^2)^{1/2}$ over the 1000 experiments for different values of t . The results are plotted in Figure 4. The figure shows that for each n , the average of $\|\mathbf{q}^t - \mathbf{q}^*\|$ converges to 0 as t increases. In fact, in all of these experiments

\mathbf{q}^t converges to \mathbf{q}^* starting from $\mathbf{q}^0 = \mathbf{0}$, and in most cases, \mathbf{q}^t becomes very close to \mathbf{q}^* in fewer than 20 steps.

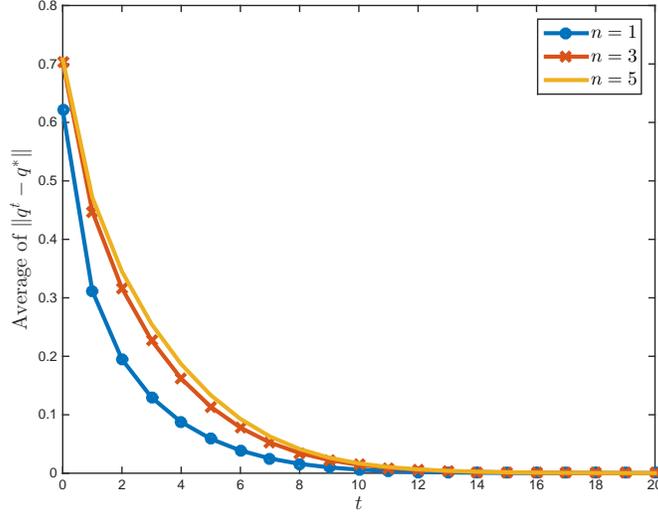


Figure 4: Average of $\|\mathbf{q}^t - \mathbf{q}^*\|$ versus t

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

From Table 1, we can see that when starting from $\mathbf{q}^0 = \mathbf{0}$, the vector of sales levels \mathbf{q}^t converges to \mathbf{q}^* for $\alpha \leq 4$, but not for $\alpha \geq 5$. In fact, for $\alpha \leq 4$, there is convergence to \mathbf{q}^* from all sampled starting points. However, for $\alpha \geq 5$, the vector \mathbf{q}^t converges to \mathbf{q}^* only when \mathbf{q}^0 is “close enough” to \mathbf{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 \mathbf{q}^* . For $\alpha = 7$,

α		0	1	2	3	4	5	6	7
$n = 5$	$\mathbf{q}^0 = \mathbf{0}$	Yes	Yes	Yes	Yes	Yes	No	No	No
	Random \mathbf{q}^0	100%	100%	100%	100%	100%	11.8%	1.8%	0.2%
	π_5^*	1.9445	2.0343	2.1295	2.4413	3.1210	3.8807	4.6843	5.5178
$n = 1$	$q^0 = 0$	Yes	Yes	Yes	Yes	Yes	No	No	No
	Random q^0	100%	100%	100%	100%	100%	64.0%	41.7%	33.4%
	π_1^*	1.0000	1.3216	1.8013	2.4142	3.1170	3.8801	4.6843	5.5178
$(\pi_5^* - \pi_1^*)/\pi_5^*$		48.57%	35.03%	15.41%	1.11%	0.13%	0.02%	<0.01%	<0.01%

Table 1: Convergence results and revenues for different values of α .

just two of the 1000 samples give us convergence to \mathbf{q}^* .

When $\alpha \geq 5$, the optimal \mathbf{q}^* follows the “one high, others low” pattern described in part (b) of Theorem 3.1. In the cases for which there is not convergence to \mathbf{q}^* from most initial conditions (the columns on the right in the table), the low value in \mathbf{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 $\mathbf{q}^0 = \mathbf{0}$, then \mathbf{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 \mathbf{q}^* . Rather, a lack of convergence to \mathbf{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 \mathbf{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 \mathbf{q}^* from $\mathbf{q}^0 = \mathbf{0}$ 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 \mathbf{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 \mathbf{q}^* . Since the low value in \mathbf{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 \mathbf{q}^* starting from a larger set of initial conditions.

In the rows of Table 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 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 Online Appendix A for references) and more generally in problems involving strategic agents. As we note in the online 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 $\mathbf{p}(q^*)$ and announce the anticipated sales vector \mathbf{q}^* . If the customers believe the seller’s announcement, then the sales will actually be \mathbf{q}^* . If the customers believe that the sales will be in a neighborhood of \mathbf{q}^* then adjustments as described above will lead sales to \mathbf{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 \mathbf{q}^* may prevail when prices $\mathbf{p}(q^*)$ 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 $\sum_{i=1}^n p_i q_i^{\text{LR}}(\mathbf{p})$ over \mathbf{p} where $\mathbf{q}^{\text{LR}}(\mathbf{p}) = (q_1^{\text{LR}}(\mathbf{p}), \dots, q_n^{\text{LR}}(\mathbf{p}))$ is the vector of quantities that yields the lowest revenue among those \mathbf{q} that satisfy (4) for prices \mathbf{p} . Such a maximization is beyond the scope of this paper and would require a different analysis. It could be the subject of another paper.

5.2 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 \mathbf{q}^* . Then we use price vector $\mathbf{p}^* = \mathbf{p}(\mathbf{q}^*)$, which gives revenue $\pi(\mathbf{q}^*) = \sum_{i=1}^n p_i^* q_i^*$.
2. We consider the optimal symmetric solution. That is, we impose an additional constraint that $q_1 = \dots = q_n$ in (P1). Let the optimal solution to this problem be \mathbf{q}^u and the corresponding price be $\mathbf{p}^u = \mathbf{p}(\mathbf{q}^u)$. 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 \mathbf{q}° denote the optimal solution, and let $\mathbf{p}^\circ(\mathbf{q})$ denote (7) with $\alpha = 0$. We compute prices $\mathbf{p}^\circ = \mathbf{p}^\circ(\mathbf{q}^\circ)$. Note that \mathbf{q}° will not be an equilibrium for these prices (except when α is actually 0) because the equilibrium condition (4) uses the actual value of α . If there are multiple sales vectors \mathbf{q} that satisfy the equilibrium condition for prices \mathbf{p}° , then in our comparison we select the one (which we will call $\bar{\mathbf{q}}^\circ$) 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. In the table, ℓ^u represents the percentage of the revenue $\pi(\mathbf{q}^*)$ that is lost if we use prices \mathbf{p}^u . Similarly, ℓ° is the percentage lost if we use prices \mathbf{p}° .

α	0	1	2	3	4	5
$\pi(\mathbf{q}^*)$	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: Comparison of revenues for different pricing decisions.

From Table 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 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
$\bar{\ell}^\circ$	10.31%	12.30%	13.03%	13.82%	12.45%	11.47%

Table 3: Comparison of revenues for different pricing decisions in heterogeneous settings

In Table 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.

5.3 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 $\tilde{\mathbf{q}}$. Then, using (7) with $\tilde{\alpha}$ in place of α , the seller obtains prices $\tilde{\mathbf{p}}(\tilde{\mathbf{q}})$. 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 (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 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 \mathbf{p}^* or $\tilde{\mathbf{p}}(\tilde{\mathbf{q}})$. Let \mathbf{q}' denote the resulting equilibrium sales vector when the seller implements prices $\tilde{\mathbf{p}}(\tilde{\mathbf{q}})$ in the problem that has actual network parameter α . Note that $\mathbf{q}' \neq \tilde{\mathbf{q}}$ when $\tilde{\alpha} \neq \alpha$. On the other hand, if $\tilde{\alpha} = \alpha$, then the seller does not make an estimation error and therefore $\mathbf{q}' = \tilde{\mathbf{q}}$. Let $\tilde{\pi} = \sum_{i=1}^5 \tilde{p}_i(\tilde{\mathbf{q}})q'_i$ be seller's revenue from implementing prices $\tilde{\mathbf{p}}(\tilde{\mathbf{q}})$ and let $\tilde{\ell}$ be the percentage revenue loss in comparison to the optimal solution (without estimation errors); i.e., $\tilde{\ell} = 100 \times (\pi(\mathbf{q}^*) - \tilde{\pi})/\pi(\mathbf{q}^*)$. Results are summarized in Table 4.

$\tilde{\alpha}$		2	2.2	2.4	2.6	2.8	3
$\tilde{\mathbf{p}}(\tilde{\mathbf{q}})$	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 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 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.

6 Concluding Remarks

In this paper, 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 paper can serve as a building block for much future research.

There are many possible directions for additional study. For instance, one could incorporate network effects into variations of the MNL model such as the mixed logit or nested logit models and study the resulting pricing problems. It also could be of interest to address questions centering on customers' perceptions of product variety. As established in part (b) of Theorem 3.1, all products but one will have low sales quantities for large values of α . If those low sales quantities are perceived by customers as a low degree of product variety, then this could possibly affect customers' purchase behaviors. This suggests research questions about how to formulate and solve pricing models that account for such possibilities. Going in yet another direction, it would also be of interest to further study settings in which there are multiple solutions to $\mathbf{q} = F(\mathbf{q})$ at $\mathbf{p} = \mathbf{p}^*$ and try to find prices that maximize the worst-case revenue (over all possible equilibria for any given prices). Network effects may also be created by sellers themselves (e.g., group buying), so it would be interesting to study how a seller could design a mechanism to stimulate network effects and how such a mechanism might affect pricing decisions.

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 (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 (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.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant Number CMMI 1462676. The authors are grateful to the three referees and associate editor for their comments and suggestions, which helped to improve this paper. The authors especially appreciate the associate editor's insightful suggestions about how to approach the general heterogeneous case found in Section 4.

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.
- M. Ben-Akiva and S. R. Lerman. *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press, 1985.
- F. Bloch and N. Qu erou. 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.
- O. Candogan, K. Bimpikis, and A. Ozdaglar. Optimal pricing in networks with externalities. *Operations Research*, 60(4):883–905, 2012.

- 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.
- N. Economides. The economics of networks. *International Journal of Industrial Organization*, 14(6):673–699, 1996.
- J. Farrell and G. Saloner. Standardization, compatibility, and innovation. *The RAND Journal of Economics*, 16(1):70–83, 1985.
- 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. Katz and C. Shapiro. Network externalities, competition, and compatibility. *The American Economic Review*, 75(3):424–440, 1985.
- 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.
- 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.
- 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.
- K. Talluri and G. van Ryzin. Revenue management under a general discrete choice model of consumer behavior. *Management Science*, 50(1):15–33, 2004.

- 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.
- 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. Working Paper, Johns Hopkins University, 2014.
- Z. Xue and J. Song. Demand management and inventory control for substitutable products. Working paper, Duke University, 2007.

Author Biographies

Chenhao Du is a Ph.D. student in the Department of Industrial and Systems Engineering at the University of Minnesota. His research interests include pricing and revenue management problems, especially those involving novel models of customer behavior.

William L. Cooper is a Professor in the Department of Industrial and Systems Engineering at University of Minnesota. His research interests include revenue management and applied probability.

Zizhuo Wang is an assistant professor in the Department of Industrial and Systems Engineering at the University of Minnesota. His research interests include pricing and revenue management, operations management, and data-driven decision making problems.