Optimal Worst-Case Pricing for a Logit Demand Model with Network Effects

Chenhao Du   William L. Cooper   Zizhuo Wang

Department of Industrial and Systems Engineering
University of Minnesota
111 Church Street SE
Minneapolis, MN 55455, USA

November 10, 2017

Abstract

We consider optimal pricing problems for a product that experiences network effects. Given a price, the sales quantity of the product arises as an equilibrium, which may not be unique. In contrast to previous studies that take a best-case view when there are multiple equilibrium sales quantities, we maximize the seller's revenue assuming that the worst-case equilibrium quantity will arise in response to a chosen price. We compare the best- and worst-case solutions, and provide asymptotic analysis of revenues.

Keywords: Pricing, Choice Model, Network Effect, Revenue Management

Email: duxxx181@umn.edu, billcoop@umn.edu, zwang@umn.edu; Corresponding Author: Cooper

1 Introduction

In this paper, we consider a pricing problem for a product with network effects. A product exhibits network effects if each individual customer's valuation for the product increases in its overall sales. For example, an online multi-player video game may be more exciting and yield greater value to an individual player when there are more total players.

In traditional pricing models for a product *without* network effects, it is common that the sales level (or expected sales level) of the product can be expressed as an explicit single-valued function of its price. The situation is more complicated for a product with network effects. The sales level of such a product influences customers' valuations, and those valuations determine customers' purchasing decisions, which themselves affect the sales level. Consequently, it is natural that, given a price, the sales level for a product with network effects is a solution to a fixed point equation. For a given price, the fixed point equation simply expresses the "equilibrium condition" that the sales level that arises must yield customer valuations that themselves induce that same sales level. If the fixed point equation admits multiple solutions (that is, there are multiple equilibria), then for a given price, there may be one fixed point (that is, equilibrium) that is a high sales level and another that is a low sales level. If this is the case, then sales revenues at those two equilibria will differ. When a pricing decision does not uniquely determine demand (even for a deterministic model), the seller is faced with the question of how to formulate and solve a suitable price optimization problem. In this paper, we study these issues for a seller of a single product.

Our starting point is a multinomial logit (MNL) choice model where each customer picks between two options: (a) buy the product or (b) do not buy the product. We emphasize that there is only one product in question. A textbook treatment of the MNL model can be found in, e.g., [1]. The MNL model may be viewed as a random utility maximization model, where each customer's utility for the product is comprised of an expected-utility term and a random term. We incorporate network effects by modifying the expected-utility term to depend upon sales. With this, customer choice probabilities depend upon sales, and the sales equilibria described above arise as fixed points of the function that, given a price, maps sales levels to choice probabilities.

We will focus on "worst-case" settings where the lowest possible sales equilibrium (that is, the smallest fixed point) is assumed to arise in response to an implemented price. Our motivation for this approach is that a seller may be wise to use a formulation that guards against a particularly undesirable market response to its pricing decision. A similar viewpoint in part underlies the growing literature on robust operational decision making and, more generally, robust optimization. In that literature, a decision maker must select an action but does not know how "nature" (e.g., the market) will respond to the various actions under consideration. For instance, a manager may need to set a price or inventory level when demand is unknown (and there is no probability distribution over demand) or when the probability distribution of demand is unknown (and the manager knows only the set of possible distributions). A robust approach selects an action that guards against responses (demand or demand distribution in the examples just mentioned) from nature that are bad for the decision maker. See, e.g, [2, 4, 13, 14, 15] for examples and many references.

For our pricing problem, the questions we address are as follows. (i) What is the seller's optimal price in the worst-case setting and how does it compare to that in a "best-case" setting where the highest possible sales equilibrium (that is, the largest fixed point) is assumed to arise in response to a price? (ii) What happens if the seller prices in expectation of a best-case equilibrium but the worst-case equilibrium arises? Conversely, what if the seller prices in expectation of a worst-case equilibrium but the best-case equilibrium arises? (iii) How do revenues in these various scenarios depend upon the strength of the network effect? Before we summarize our answers, we remark that question (ii) addresses the broad issue of what happens if a decision maker has incorrect beliefs (in this case about which equilibrium will arise). The answer to (ii) can help inform a seller who is choosing between solving the best- and worst-case formulations, because it will help the seller understand the cost of making an incorrect assumption. Questions about effects of incorrect beliefs (e.g., how bad — or good — are decisions obtained from models founded on incorrect assumptions?) arise in work on robust decision-making and in work on model misspecification. See, e.g., [3, 7, 8] for entry into the literature on model misspecification in operations.

In this paper, we show that the worst-case pricing problem can be solved via a one-dimensional optimization problem with a unimodal objective function. The optimization problem also provides a link between the best- and worst-case formulations, from which we find that the two formulations have the same solution if the network effect is weak but different solutions if the network effect is strong. (In our framework, the "strength" of the network effect depends upon a parameter that governs the extent to which sales affect an individual's expected utility for the product.) In settings where the best- and worst-case problems yield different answers, we find that in the best-case problem the seller sets a higher price and obtains a lower sales level (and higher revenue) than in the worst-case problem. The difference in revenues in the two cases can be large. In fact, in settings with very strong network effects, we prove that the best-case revenue is roughly proportional to the parameter mentioned above, while the worst-case revenue is roughly proportional to its logarithm.

If the seller is "misguidedly optimistic" and sets the price prescribed by the solution of the best-case formulation, but (contrary to the assumption underlying that formulation) the worst-case equilibrium for that price prevails, then the realized revenue may be below what the solution to the best-case formulation suggested it would be and also below what it would have been if the seller had instead implemented the worst-case pricing solution. With a weak network effect, such an issue does not arise because the two formulations have the same solution. However, if the network effect is strong, then the phenomenon is quite pronounced. We prove that the revenue under misguided optimism is roughly proportional to the reciprocal of an expression that is exponential in the parameter that determines the strength of the network effect. Thus, a misguidedly optimistic seller's revenue is almost zero in such settings. If the seller is instead "incorrectly pessimistic" and sets the price prescribed by the solution of the worst-case problem in a setting where the best equilibrium prevails, then a related phenomenon occurs. It turns out that the price obtained from the worst-case formulation yields a unique equilibrium. Nevertheless, in such settings the seller would be better off using the best-case price.

To close this section, we provide a very short literature review. The MNL model and its variants have been widely used in the revenue management literature for problems without network effects. For examples and references, see [9, 11, 12, 17, 18, 19]. To draw distinctions with our work, these papers do not consider network effects, and hence do not need to consider (multiple) sales equilibria.

The papers [10] and [20] use the MNL model — modified as described above — in pricing and assortment planning problems with network effects. Both of these papers contain some results regarding (non-)uniqueness of equilibria, but neither focuses on the issue from a decision-making standpoint. Other research that considers MNL models with network effects includes [5], [16], and Section 7.8 of [1]. These studies address the possibility of multiple equilibria, but their focus is quite different from ours. For additional pointers to the literature on network effects, see [10, 20].

The paper [10] cited in the preceding paragraph is the closest to our current work. It considers pricing for multiple substitute products with network effects, but addresses only the best-case setting; it does not address the worst-case setting. Its main results describe the structure of an optimal solution in the best-case multi-product setting. In particular, it proves that if the products are a priori homogeneous, then the price of one of the products should be set to a low value and the prices of all other products should be set to a single high value. Similar, but more complicated, results are also established for products that are a priori heterogeneous. [10] also presents computational algorithms. Such results about the relationship among prices in a multi-product setting do not have direct bearing on the single-product problem we consider. Aside from the demand model under consideration, the focus of [10] is markedly different from that of the present paper.

The presence of multiple sales equilibria that yield different revenues for the seller is similar to a situation that may arise in Stackelberg games, and more generally, in bilevel programming problems. In a Stackelberg game, the leader makes a decision and the follower responds with its own optimal decision, but there may be multiple possible values for the follower's optimal decision. This is similar in spirit to our pricing problem in which the seller implements a price and the market responds with a sales level, but there may be multiple possible values for that level. This issue in Stackelberg games may be addressed with "optimistic" and "pessimistic" formulations akin to the best-case and worst-case approaches considered herein. For entry into this literature, see [6].

2 The Model

Consider a seller who must set the price p for a single product. Demand is given by a standard logit model, modified to incorporate network effects. To begin, we describe this demand model, which is the same as the one in [10] specialized to a single product. Each individual customer has a valuation $U = v + \epsilon$ for the product where v is constant (given the price) across the population of customers and ϵ varies across the population of customers. We assume $v = y - p + \alpha q$, where y is a constant that depends upon intrinsic properties of the product, q is the sales quantity of the product, and $\alpha \geq 0$ is a network effect sensitivity parameter. The value a customer gets from the product is increasing in q. We may view α as reflecting the strength of the network effect. If α is large, then a customer's valuation is quite sensitive to sales q, and the network effect is strong. If α is small, then a customer's valuation is less sensitive to q, and the network effect is weak. We assume throughout that $y \geq 0$. For discussion of settings where y is negative, see the remark after Lemma 5.1 in Section 5.

Upon defining $v_0 = 0$, we also assume that each individual customer has a valuation $U_0 = v_0 + \epsilon_0$ for the no-purchase option (i.e., for not buying the product) and that ϵ_0 varies across the population of customers. Each customer observes how much (s)he values the product and how much (s)he values the no-purchase option, and then picks the option with the larger value. The assumptions that $y \geq 0$ and $v_0 = 0$ mean that an average customer prefers getting the product for free to getting nothing. If one wishes to consider a situation with $v_0 \neq 0$, then we can simply replace v_0 by $v_0' = 0$ and y by $y' = y - v_0$ in our model. If $y' \geq 0$ then all our results still hold. If $y' < 0$ then some of our main results will hold (Theorems 4.2 and 4.3) and others may not (Theorem 4.1). Here, we again refer to the remark after Lemma 5.1 for further discussion.

We consider a "fluid model" of demand, and with no loss of generality, scale the size of the population of customers to 1. In such a fluid model, the fraction of customers whose ϵ and ϵ_0 are in any particular range (and, in view of the assumption of a population of size 1, also the *number* of customers whose ϵ and ϵ_0 are in that range) is the same as the probability that the ϵ and ϵ_0 of an individual customer are in that range. As in the usual logit model, we assume ϵ and ϵ_0 are independent Gumbel random variables for each customer. It follows from standard results for the MNL model that the probability a typical customer will buy the product when the price is p is

$$P(U > U_0) = \frac{\exp(v)}{1 + \exp(v)} = \frac{\exp(y - p + \alpha q)}{1 + \exp(y - p + \alpha q)} =: F(p, q)\,.$$

From our assumption of a fluid model with a population of size 1, we have $q = P(U > U_0)$. Thus,

$$q = F(p, q)\,. \tag{1}$$

The seller wishes to maximize its revenue $\pi(p, q) = pq$. The seller implements price p, and the market responds with sales quantity q that satisfies (1). The heart of the issue we address is that for a given price p, it is possible that there are multiple quantities that satisfy (1) and the associated revenues may differ greatly. That is, for given p, it is possible that there are $q \neq q'$ such that $q = F(p, q)$ and $q' = F(p, q')$ with (say) $\pi(p, q) \gg \pi(p, q')$. See [10] for an optimistic (best-case) approach where the revenue maximization problem is solved while implicitly assuming that for any price p, the sales quantity that arises is the one with the highest revenue among those that satisfy (1). The best-case assumption in [10] is implicit because that paper does not present a formulation that explicitly differentiates between best and worst cases, but rather presents and solves a formulation that turns out to be equivalent to a best-case formulation. Herein, we mainly focus on a pessimistic (worst-case) setting in which for any price p, the sales quantity that arises is the one with the lowest revenue among those that satisfy (1).

For price p, define $Q(p)$ to be the set of q that satisfy (1), i.e., $Q(p) = \{q \in [0, 1] : q = F(p, q)\}$. With this, we can restate (1) as follows:

$$q \in Q(p)\,. \tag{2}$$

We can now present the best-case and worst-case pricing problems. The best-case pricing problem (in essence studied in [10]) is

$$\begin{aligned} \overline{\pi} &= \sup_p \overline{\pi}(p) \\ \overline{\pi}(p) &= \max_q \{\pi(p,q) : q \in Q(p)\} \ . \end{aligned} \tag{BC}$$

Likewise, the worst-case pricing problem (the main topic of this paper) is

$$\begin{aligned} \underline{\pi} &= \sup_p \underline{\pi}(p) \\ \underline{\pi}(p) &= \min_q \{\pi(p,q) : q \in Q(p)\} \ . \end{aligned} \tag{WC}$$

Lemma 3.1 below establishes that $Q(p)$ is finite for each p. Hence, the maximum over q in (BC) and the minimum over q in (WC) are both attained. As we will see later, there may be no optimal solution to $\sup_p \underline{\pi}(p)$ in (WC). In such cases, we must be satisfied with an ϵ-optimal solution, say p^ϵ, wherein $\underline{\pi}(p^\epsilon) > \sup_p \underline{\pi}(p) - \epsilon$.

3 Preliminary Analysis

In this section, we provide insight into when multiple sales equilibria exist, and also outline an approach from [10] to solve the best-case problem. The approach will also be an ingredient in our procedure for solving the worst-case problem. To begin, for $q \in (0,1)$ let

$$p(q) = y + \alpha q - \log q + \log(1-q). \tag{3}$$

For any given sales quantity $q \in (0,1)$, some algebra shows that $p = p(q)$ is the unique price for which (2) holds. For $q \in (0,1)$, we have that $q \in Q(p)$ if and only if $p(q) = p$. This does not preclude the existence of some other value (say q') such that $p(q)$ and q' also together satisfy (2).

Figure 1 plots $p(q)$ in a case with $\alpha > 4$. (The points BC, WC, and MO are explained later.) As we will see shortly, if $\alpha \le 4$ then $p(q)$ has a simpler structure (i.e., it is decreasing). The (p,q)-pairs that satisfy (2) are simply the points in two-dimensional space on the graph of $p(q)$. Therefore, we can determine the number of sales equilibria for a given price p by counting the number of times a horizontal line at height p intersects $p(q)$. If p is between p^L and p^H in the figure, then there are three q that satisfy (2). The best-case approach assumes sales will be the largest of these three values. If sales instead turn out to be the smallest of the three (which would be consistent with the worst-case assumption), then sales — and revenue — will be much lower. For example, in

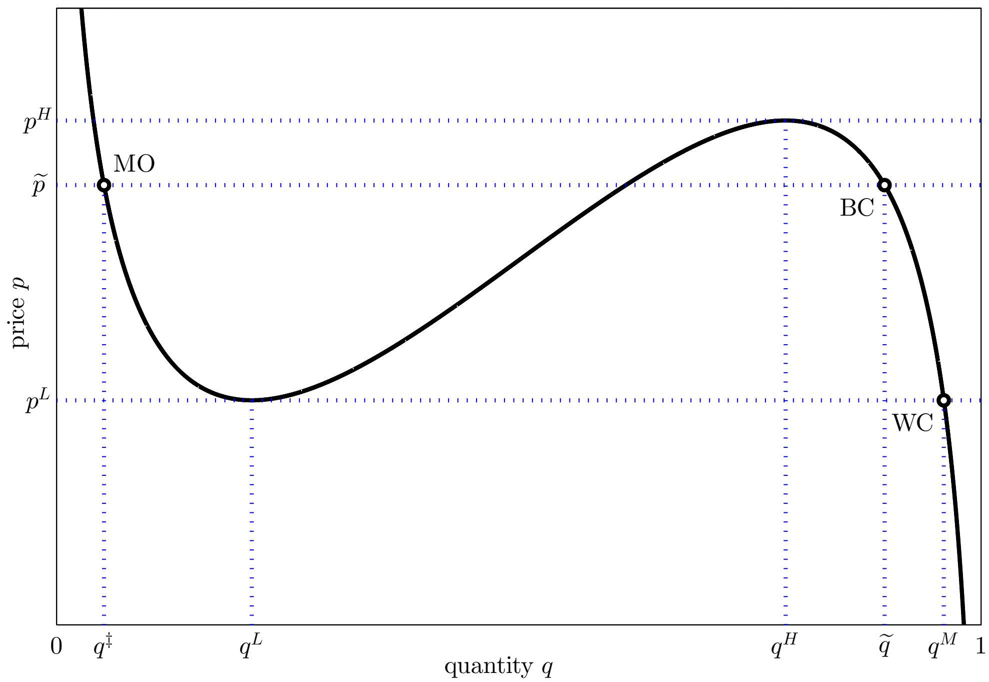

Figure 1: The function $p(q)$ for $\alpha = 6$, $y = 1$.

Figure 1, if the price is $\widetilde{p}$ (which, in this example, is the optimal price in (BC)), then the largest sales quantity that could arise is $\widetilde{q} \approx 0.90$, while the smallest that could arise is $q^\ddagger \approx 0.05$.

The following lemma describes the structure of $p(q)$. In the interest of space, we omit the proof, which follows from (3) and simple calculus.

Lemma 3.1. *The function* $p(q)$ *defined in* (3) *satisfies* $\lim_{q\downarrow 0} p(q) = \infty$ *and* $\lim_{q\uparrow 1} p(q) = -\infty$. *In addition, we have the following.*

1. *Suppose* $\alpha \leq 4$. *Then* $p(q)$ *is decreasing, and for each* p, *there is a unique* q *that satisfies* (2).

2. *Suppose* $\alpha > 4$. *Then* $p(q)$ *has a unique local minimum at* $q^L = 1/2 - \sqrt{1/4 - 1/\alpha}$ *and a unique local maximum at* $q^H = 1/2 + \sqrt{1/4 - 1/\alpha}$. *Also,* $p(q)$ *decreases on* $(0, q^L]$, *increases on* $(q^L, q^H]$, *and decreases on* $(q^H, 1)$. *For* $p^L := p(q^L)$ *and* $p^H := p(q^H)$ *we have:*

   (a) *for each* $p \in (p^L, p^H)$, *there are three* q *that satisfy* (2);

   (b) *for each* $p \in \{p^L, p^H\}$, *there are two* q *that satisfy* (2);

   (c) *for each* $p \notin [p^L, p^H]$, *there is a unique* q *that satisfies* (2).

Lemma 3.1 implies that multiple equilibria may arise only if network effects are strong enough ($\alpha > 4$). In settings with weak network effects ($\alpha \leq 4$), problems (BC) and (WC) are equivalent, because for each price, there is a unique equilibrium. Problem (BC) was already solved in [10]. Hence, we hereafter assume $\alpha > 4$.

We close this section with an approach for solving (BC). Define

$$\widetilde{\pi}(q) = p(q)q = yq + \alpha q^2 - q\log(q) + q\log(1-q), \tag{4}$$

and consider the maximization problem

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

We can summarize our results for (BC) with the following.

Proposition 3.2. *Problems* (BC) *and* (P0) *are equivalent; i.e,* $\overline{\pi} = \widetilde{\pi}^*$. *There is a unique solution* $(\overline{p}, \overline{q})$ *to* (BC)*, there is a unique solution* $\widetilde{q}$ *to* (P0)*, and* $(\overline{p}, \overline{q}) = (p(\widetilde{q}), \widetilde{q})$. *In addition,* $\widetilde{q} > q^H$.

Proofs of this and subsequent results are in Section 5. The essence of the above proposition is that to solve (BC), it suffices to solve the single-dimensional optimization problem (P0) where the decision variable is the quantity. Lemma 5.1 of Section 5 establishes that $\widetilde{\pi}(q)$ is strictly unimodal. Thus, the unique maximizer $\widetilde{q}$ of $\widetilde{\pi}(q)$ can be found efficiently through a bisection search. If there are *no* network effects ($\alpha = 0$), then $\widetilde{\pi}(q)$ is, in fact, concave; for discussion of this result for problems without network effects, we refer to Section 2.1 of [12].

4 Main Results

In this section we solve (WC), and make comparisons with (BC). We then consider what happens if the seller has an incorrect belief about which equilibrium will prevail, and study how the strength of the network effect, as measured by α, affects the seller's revenue in different scenarios.

Let q^M be the larger of the two q for which (p^L, q) satisfies (2); see part 2(b) of Lemma 3.1 and Figure 1. Observe that $q^M > q^H > q^L$ and $p(q^M) = p(q^L) = p^L$. The following, which describes the solution to (WC), is our first main result. Recall that $\widetilde{q}$ is the optimal solution to (P0).

Theorem 4.1.

1. *If* $\widetilde{q} > q^M$*, then the unique optimal solution* $(\underline{p}, \underline{q})$ *to* (WC) *is given by* $(\underline{p}, \underline{q}) = (p(\widetilde{q}), \widetilde{q})$. *Moreover,* $\underline{\pi} = \underline{p} \cdot \underline{q} = p(\widetilde{q})\widetilde{q}$.

2. *If* $\widetilde{q} \leq q^M$*, then there does not exist an optimal solution to* (WC)*. For any* $\epsilon \in (0, p^L]$*, we have that* $(\underline{p}^\epsilon, \underline{q}^\epsilon) := (p^L - \epsilon, q^M + \delta(\epsilon))$ *is an* ϵ*-optimal solution to* (WC)*, where* $\delta(\epsilon) > 0$ *is the unique solution to* $p(q^M + \delta) = p^L - \epsilon$. *Moreover,* $\underline{\pi} = p^L \cdot q^M$.

Proposition 3.2 and Theorem 4.1 reveal a simple relationship between the best- and worst-case problems. If the optimal solution $\widetilde{q}$ to (P0) is larger than q^M, then the two problems have the same solution. On the other hand, if $\widetilde{q}$ is less than q^M, then an optimistic seller will charge more than will a pessimistic seller, and the optimist will expect to obtain a lower sales quantity but higher revenue than will the pessimist. This occurs because the pessimist avoids any price for which there exists a "very low" equilibrium sales quantity, even if there is also a high equilibrium sales quantity associated with that price. This leads the pessimist to charge a lower price than the optimist. The optimist is not concerned with the existence of a low quantity associated with a price so long as there is also high quantity. An example with $\widetilde{q} < q^M$ is depicted in Figure 1. The points labeled BC and WC correspond to the (p, q)-pairs obtained from problems (BC) and (WC). In the interest space, we do not provide a figure that depicts a case with $\widetilde{q} > q^M$. We note, however, that the shape of the function $p(q)$ in such cases does not appear to the naked eye to be markedly different from $p(q)$ shown in Figure 1.

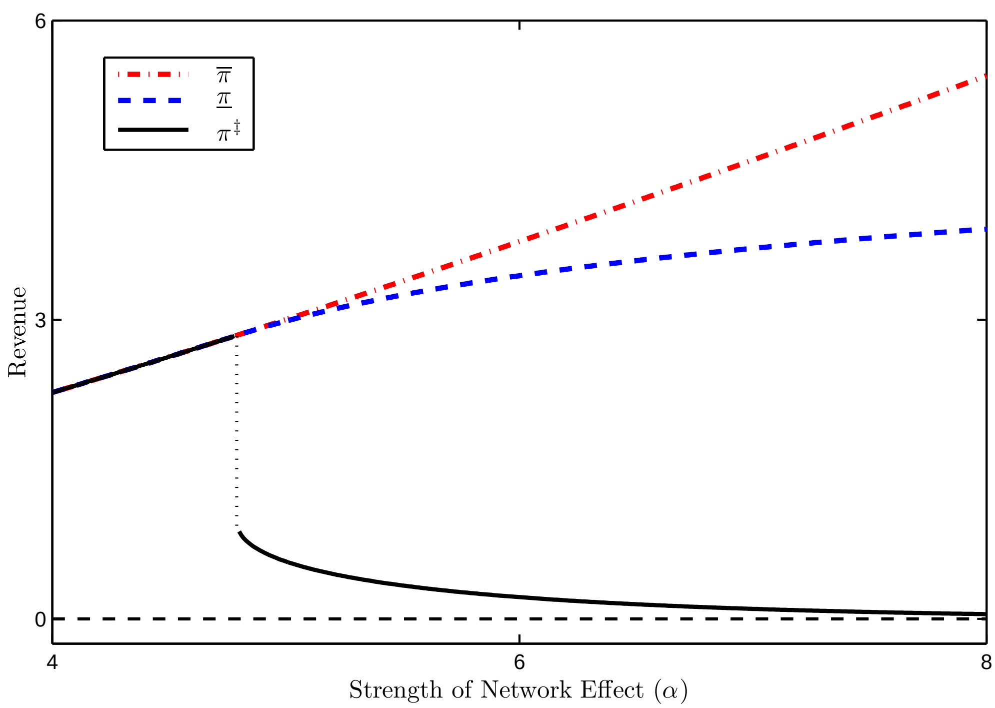

Figure 2: Comparison of revenues $(y = 1)$.

Figure 2 shows how $\underline{\pi}$ and $\overline{\pi}$ vary with the parameter α, which measures the strength of the network effect. For small α, we have $\overline{\pi} = \underline{\pi}$. This can be explained by the fact that for such α we have $\widetilde{q} > q^M$ as in part 1 of Theorem 4.1. On the other hand, $\overline{\pi} > \underline{\pi}$ for large α. For such α, we have $\widetilde{q} \leq q^M$ as in part 2 of the theorem. (These statements about the relationship between $\widetilde{q}$ and q^M are proved in Lemma 5.5.) As α increases, both $\overline{\pi}$ and $\underline{\pi}$ increase because, all else equal, customers value the product more. In both the best and worst cases, larger values of α allow the seller to charge more and also generate higher sales. The figure suggests that $\overline{\pi}$ grows more rapidly than $\underline{\pi}$. In fact, $\overline{\pi}$ is asymptotically proportional to α while $\underline{\pi}$ is asymptotically proportional to $\log \alpha$. This is made precise in the next theorem. In preparation, recall that $f(\alpha) \sim g(\alpha)$ as $\alpha \to \infty$ means $\lim_{\alpha\to\infty} f(\alpha)/g(\alpha) = 1$. Similarly, $f(\alpha) = \Theta(g(\alpha))$ as $\alpha \to \infty$ means $C_1 \leq \liminf_{\alpha\to\infty} f(\alpha)/g(\alpha) \leq \limsup_{\alpha\to\infty} f(\alpha)/g(\alpha) \leq C_2$ for some constants $C_1, C_2 > 0$.

Theorem 4.2. *(i)* $\overline{\pi} \sim \alpha$ *and (ii)* $\underline{\pi} \sim \log \alpha$ *as* $\alpha \to \infty$.

It is easy to see from the expression for q^H in Lemma 3.1 that $q^H \to 1$ as $\alpha \to \infty$. In addition, $q^H < \widetilde{q} < 1$ by Proposition 3.2. Moreover, $q^H < q^M < 1$. Hence $\widetilde{q} \to 1$ and $q^M \to 1$ as $\alpha \to \infty$. It follows that the large asymptotic difference between $\overline{\pi}$ and $\underline{\pi}$ derives from the fact that the seller charges a much higher price in (BC) than in (WC) while obtaining almost the same sales level.

To help build intuition about Theorem 4.2, consider the best-case setting. For simplicity, we take $y = 0$ in this argument. From the proof of the theorem we see that for large α, the optimal price is roughly $p^H = p(q^H)$, which itself is close to $\alpha - \log \alpha$, and the associated best-case quantity is roughly q^H, which itself is close to 1. This yields a revenue of roughly $(\alpha - \log \alpha) \times 1 \sim \alpha$ as in part (i) of the theorem. To understand why the optimal price is roughly $\alpha - \log \alpha$, note that a lower price will increase sales, but not by much because the sales level is already very close to 1. This suggests it is suboptimal to price below $\alpha - \log \alpha$. It can be seen in Figure 1 that pricing higher than p^H yields a very low sales quantity (smaller than q^L, which is itself close to 0 when α is large). Why is it suboptimal to charge a very high price and get very low sales? From (3) we have for $n \geq 1$ that $p(\exp(-\alpha^n)) = \alpha \exp(-\alpha^n) - \log(\exp(-\alpha^n)) + \log(1 - \exp(-\alpha^n)) \approx \alpha^n$. So, a very high price of roughly α^n yields a quantity of roughly $\exp(-\alpha^n)$ and a revenue of roughly $\alpha^n \exp(-\alpha^n)$, which is much lower than the revenue from pricing at $\alpha - \log \alpha$. Increases in price above $\alpha - \log \alpha$ lead to very sharp decreases in the sales quantity, making very high prices suboptimal.

Pricing at $\alpha - \log \alpha$ also yields another equilibrium quantity that is less than q^L (near 0), making it a poor choice for the seller in the worst-case problem. More generally, any prices between p^L and p^H will suffer from a similar issue. Pricing "just below" p^L avoids such a low equilibrium. For large α, it can be checked that $p^L = p(q^L)$ and q^L are roughly equal to $\log \alpha$ (see the proof of Lemma 5.4) and 1 respectively, yielding the worst-case revenue of about $\log \alpha$ in part (ii) of the theorem. Intuitively, the seller must charge a lower price in the worst case than in the best case in order to avoid a "bad equilibrium" near 0. Prices lower than roughly $\log \alpha$ and higher than roughly $\alpha - \log \alpha$ yield a unique equilibrium but are suboptimal in the worst case for the same reasons (described in the previous paragraph) that they are suboptimal in the best case.

To close this section, we address the question: what happens if the seller is "misguidedly optimistic" and solves (BC), but upon implementing the prescribed price, the worst corresponding sales quantity arises? We may similarly ask what if the seller is "incorrectly pessimistic" and sets the price obtained from solving (WC), but the best corresponding sales quantity arises?

We start with the simpler case of incorrect pessimism. Suppose that for any given price, the best equilibrium sales level will actually prevail, but that the seller believes incorrectly that the worst equilibrium will prevail. Let $q^\dagger = \max\{q : q \in Q(\underline{p})\}$ be the *largest* sales quantity that can arise from the worst-case price $\underline{p}$. The incorrectly pessimistic seller implements price $\underline{p}$ and subsequently the sales level $q^\dagger$ arises. The seller obtains revenue $\pi^\dagger := \overline{\pi}(\underline{p}) = \underline{p} \cdot q^\dagger$. In case of non-existence of $\underline{p}$ as in part 2 of Theorem 4.1, we here take $\underline{p} = \lim_{\epsilon \downarrow 0} \underline{p}^\epsilon$. We know that $\underline{p}$ is set to $p(\widetilde{q})$ if $\widetilde{q} > q^M$, and otherwise $\underline{p}$ is set "infinitesimally" below $p(q^M)$ if $\widetilde{q} \leq q^M$. In either case, there is a unique corresponding sales level; see Lemma 3.1 and Figure 1. Hence, the actual sales level will not differ from that predicted by the solution to (WC), and the seller will not realize it was incorrect in its pessimism and will obtain revenue $\underline{\pi}$. That is, $\pi^\dagger = \underline{\pi}$. Thus, the "cost" of incorrect pessimism to the seller (i.e., the loss in comparison to what it could have earned with a correct belief) is $\overline{\pi} - \pi^\dagger = \overline{\pi} - \underline{\pi}$. Note that $\overline{\pi} - \underline{\pi}$ is 0 if $\widetilde{q} \geq q^M$ and is positive otherwise.

Next we turn to the case of misguided optimism. Suppose that for any given price, the worst equilibrium sales level will actually prevail, but the seller believes incorrectly that the best equilibrium will prevail. Let $q^\ddagger = \min\{q : q \in Q(\overline{p})\}$ be the *smallest* sales quantity that can arise from the best-case price $\overline{p}$. The misguidedly optimistic seller implements price $\overline{p}$ and subsequently the sales level $q^\ddagger$ arises. The seller obtains revenue $\pi^\ddagger := \underline{\pi}(\overline{p}) = \overline{p} \cdot q^\ddagger$. If $\widetilde{q} > q^M$, then $\underline{p} = \overline{p} = p(\widetilde{q})$ and $q^\ddagger = \underline{q} = \overline{q}$. Hence, $\pi^\ddagger = \underline{\pi} = \overline{\pi}$, and the cost of misguided optimism $\underline{\pi} - \pi^\ddagger$ is 0. On the other hand, if $\widetilde{q} \leq q^M$ (so $\widetilde{q} \in (q^H, q^M]$), then an appeal to Lemma 3.1 and Figure 1 shows that $\overline{p} = p(\widetilde{q}) \in [p^L, p^H)$ and $q^\ddagger \leq q^L < q^H < \widetilde{q} = \overline{q}$. So, $q^\ddagger$ will be smaller than $\overline{q}$. Consequently, if $\widetilde{q} \leq q^M$, then $\pi^\ddagger = p(\widetilde{q}) \cdot q^\ddagger$ and the cost of misguided optimism is $\underline{\pi} - \pi^\ddagger = p^L \cdot q^M - p(\widetilde{q}) \cdot q^\ddagger$.

Figure 1 shows an example for which $\widetilde{q} \in (q^H, q^M]$. In the figure the (p, q)-pair corresponding to misguided optimism is labeled MO. As indicated above, the solution in the case of incorrect pessimism is the same as WC. The figure shows that the sales quantity ($q^\ddagger \approx 0.05$) obtained in the case of misguided optimism is quite low in comparison to that obtained in both the best-case and worst-case solutions. To help understand the effect misguided optimism has on revenues, Figure 2 plots $\pi^\ddagger$ against the network effect parameter α. Observe that there is a discontinuity in $\pi^\ddagger$ at $\alpha \approx 4.8$. This discontinuity arises at the value of α where $\widetilde{q}$ moves from above q^M to below q^M. When this happens, $q^\ddagger$ shifts from coinciding with $\widetilde{q}$ to being much smaller than q^L. The figure shows that for small α (weak network effect), $\pi^\ddagger$ coincides with $\overline{\pi}$ and $\underline{\pi}$, consistent with the discussion above. For large α, however, $\pi^\ddagger$ is roughly zero. This can be traced to very low quantities that arise in the case of misguided optimism when α is not small. In fact, the figure suggests that $\pi^\ddagger$, when viewed as a function of α, approaches 0 quickly as α grows. This is made precise in Theorem 4.3, which shows that $\pi^\ddagger$ converges to 0 at a rate that is exponential in α.

Theorem 4.3. *$\pi^\ddagger = \Theta(e^{-\alpha}\alpha^2)$ as $\alpha \to \infty$.*

When deciding whether to use (BC) or (WC), a seller may combine Theorems 4.2 and 4.3 to make rough comparisons of the absolute and relative costs of incorrect pessimism and misguided optimism. A seller may also be conservative and simply wish to avoid very low revenues, which would lead it to (WC). In the end, any such decision depends upon the judgement of the seller.

5 Proofs and Auxiliary Results

Proof of Proposition 3.2. For each p we have that (i) for $q \in (0,1)$ we have $q \in Q(p)$ if and only if $p(q) = p$, and (ii) $0, 1 \notin Q(p)$. Therefore, for each p we have $\max_q\{\pi(p,q) : q \in Q(p)\} = \max_{q\in(0,1)}\{pq : p(q) = p\}$.

Lemma 5.1 below establishes that there is a unique optimal solution $\widetilde{q}$ to (P0) and that $\widetilde{q} > q^H$. We next show that $(p(\widetilde{q}), \widetilde{q})$ is an optimal solution (BC). For arbitrary price p, let $q^\dagger(p)$ be the unique maximizer in the problem $\max_q\{pq : p(q) = p\}$. That is, $q^\dagger(p)$ is the largest q such that $p(q) = p$. We have that $q^\dagger(p)$ satisfies $p(q^\dagger(p)) = p$. Therefore,

$$\max_q\{pq : p(q) = p\} = pq^\dagger(p) = p(q^\dagger(p))q^\dagger(p) \leq p(\widetilde{q})\widetilde{q} \tag{5}$$

where the inequality holds because $\widetilde{q}$ maximizes $p(q)q$. For $p = p(\widetilde{q})$, it is apparent that $q^\dagger(p(\widetilde{q})) = \widetilde{q}$, and the weak inequality in (5) becomes an equality. Thus, $(p(\widetilde{q}), \widetilde{q})$ is an optimal solution to (BC).

For uniqueness of the optimal solution to (BC), consider any price $p \neq p(\widetilde{q})$. Then the inequality in (5) must be strict, or else we would get a contradiction that $\widetilde{q}$ is the unique maximizer of $p(q)q$. $\square$

Proof of Theorem 4.1. Let $q^\ddagger(p) = \min\{q : q \in Q(p)\}$. We may write (WC) as $\underline{\pi} = \sup_p pq^\ddagger(p) = \max\{A_1, A_2\}$ where $A_1 = \sup_p \left\{pq^\ddagger(p) : p \in (0, p^L)\right\}$ and $A_2 = \sup_p \left\{pq^\ddagger(p) : p \in [p^L, \infty)\right\}$. By

Lemma 3.1, we have that $A_1 = \sup_q \{\widetilde{\pi}(q) : q \in (q^M, 1)\}$ and $A_2 = \sup_q \{\widetilde{\pi}(q) : q \in (0, q^L]\}$. Lemma 5.1 implies $\widetilde{\pi}(q)$ is increasing on $(0, q^L]$. So, $A_2 = \widetilde{\pi}(q^L)$. Note also that $\widetilde{\pi}(q^L) = p(q^L)q^L = p(q^M)q^L < p(q^M)q^M = \widetilde{\pi}(q^M) = \lim_{\epsilon\downarrow 0} \widetilde{\pi}(q^M + \epsilon)$. Hence, $A_2 < A_1$, and consequently, $\underline{\pi} = A_1$.

If the maximizer $\widetilde{q}$ of $\widetilde{\pi}(q)$ over $(0, 1)$ lies in $(q^M, 1)$, then $\underline{\pi} = p(\widetilde{q})\widetilde{q}$ and the optimal solution to (WC) is $(\underline{p}, \underline{q}) = (p(\widetilde{q}), \widetilde{q})$. By Lemma 5.1, the other possibility is that $\widetilde{q}$ lies in $(q^H, q^M]$. In this case, again by Lemma 5.1 (unimodality of $\widetilde{\pi}(q)$), we have $\underline{\pi} = \widetilde{\pi}(q^M) = p(q^M)q^M = p^L \cdot q^M$. Here, the supremum in (WC) and A_1 is not attained, but an ϵ-optimal solution is given by $(\underline{p}^\epsilon, \underline{q}^\epsilon)$. To see this, note that the set $Q(\underline{p}^\epsilon)$ is the singleton $\{\underline{q}^\epsilon\}$ because $\underline{p}^\epsilon < p^L$. Moreover, $\underline{\pi}(\underline{p}^\epsilon) - \underline{\pi} = \underline{p}^\epsilon \cdot \underline{q}^\epsilon - p^L \cdot q^M = (p^L - \epsilon)(q^M + \delta) - p^L \cdot q^M \geq (p^L - \epsilon)q^M - p^L \cdot q^M > -\epsilon$. □

Lemma 5.1. *Suppose* $\alpha > 4$. *Then* $\widetilde{\pi}(q)$ *is a strictly unimodal function on* $(0, 1)$ *with a unique maximizer* $\widetilde{q} \in (0, 1)$. *In addition,* $\widetilde{q} > q^H$.

Proof. The first and second derivatives of $\widetilde{\pi}(q)$ are

$$\widetilde{\pi}'(q) = y + 2\alpha q - \log(q) + \log(1 - q) - \frac{1}{1 - q}, \tag{6}$$

$$\widetilde{\pi}''(q) = 2\alpha - \frac{1}{q(1 - q)^2}. \tag{7}$$

To prove the lemma it suffices to show (i) $\widetilde{\pi}''(q) < 0$ for $q \in (q^H, 1)$, (ii) $\widetilde{\pi}'(q) > 0$ for $q \in (0, q^H)$, and (iii) $\lim_{q\uparrow 1} \widetilde{\pi}'(q) = -\infty$. Item (iii) follows easily from (6), so we need only establish (i) and (ii).

We begin with (i). Note that $g(q) := \frac{1}{q(1-q)^2}$ is decreasing for $q \in (0, 1/3)$, increasing for $q \in (1/3, 1)$, and attains its minimum of $27/4 = g(1/3)$ at $q = 1/3$. Moreover, $\lim_{q\downarrow 0} g(q) = \lim_{q\uparrow 1} g(q) = \infty$. Recall that $\alpha > 4$. Thus, $2\alpha > 8 > 27/4$, and it therefore follows from (7) that $\widetilde{\pi}''(1/3) > 0$. Hence, $\widetilde{\pi}''(q) = 0$ has exactly two solutions, $q_1 < q_2$. Moreover, $\widetilde{\pi}''(q) < 0$ on $(0, q_1)$ and $(q_2, 1)$, and $\widetilde{\pi}''(q) > 0$ on (q_1, q_2).

Recall from Lemma 3.1 that q^H and q^L are the solutions to $p'(q) = 0$, and therefore $\alpha = \frac{1}{q(1-q)}$ for $q \in \{q^H, q^L\}$. Hence, $\widetilde{\pi}''(q) = 2\alpha - \frac{1}{q(1-q)^2} = 2\alpha - \frac{\alpha}{1-q}$ for $q \in \{q^H, q^L\}$. By Lemma 3.1, we have $q^L < 1/2$ and $q^H > 1/2$. So, $\widetilde{\pi}''(q^L) > 0$ and $\widetilde{\pi}''(q^H) < 0$. This implies $q^L \in (q_1, q_2)$ and $q^H \in (q_2, 1)$, from which (i) now follows.

Next, we prove (ii). For $q \in (q^L, q^H)$, we have $p'(q) > 0$. Therefore, $\widetilde{\pi}'(q) = p'(q)q + p(q) > p(q^L) \geq y + \alpha q^L + \log q^H \geq y + \alpha q^L + 1 - 1/q^H = y + 1$ for $q \in (q^L, q^H)$. So, $\widetilde{\pi}'(q) > 0$ for $q \in (q^L, q^H)$ because $y \geq 0 > -1$. We will complete the proof of (ii) by showing that $\widetilde{\pi}'(q) > 0$ for $q \in (0, 1/2)$. Because $\alpha > 4$, we have from (6) that $\widetilde{\pi}'(q) > y + 8q - \log(q) + \log(1 - q) - \frac{1}{1-q} =: f(q)$. Hence, it suffices to establish that $f(q) > 0$ for $q \in (0, 1/2)$. Note that $f'(q) = 8 - g(q)$ and $g(1/2) = 8$.

Together with the facts about $g(q)$ given above, this implies that $f'(q) = 0$ has exactly one solution $\widehat{q}$ on $(0, 1/2)$ and that solution must lie in $(0, 1/3)$, where $g(q)$ is decreasing. We have $f''(q) = -g'(q)$, so $\widehat{q}$ must be a local minimum of $f(q)$. It is now simple to check that $\widehat{q} = (3 - \sqrt{5})/4$. Moreover, $f(\widehat{q}) = y + 7 - 3\sqrt{5} + \log(2 + \sqrt{5}) > 0$ because $y \geq 0$. It follows that $f(q) > 0$ for all $q \in (0, 1/2)$, which proves (ii). $\square$

Remark. Throughout we have assumed $y \geq 0$. Examination of the proof of Lemma 5.1 above shows that this condition can be relaxed to $y \geq -7+3\sqrt{5}-\log(2+\sqrt{5}) \approx -0.9187$. Hence, Proposition 3.2 and Theorems 4.1, 4.2, and 4.3 all hold under the weaker assumption that $y \geq -7+3\sqrt{5}-\log(2+\sqrt{5})$. If y is somewhat more negative, then Lemma 5.1 may not hold. For example, if $\alpha = 8$ and $y = -4$, then $\widetilde{\pi}$ is not unimodal. For such values of negative y where Lemma 5.1 does not hold, Proposition 3.2 and Theorem 4.1 also do not hold. Nevertheless, it turns out that the asymptotic results in Theorems 4.2 and 4.3 hold for *any* value of y. The key to this observation is the following lemma. Using it in place of Lemma 5.1, we see that Proposition 3.2 and Theorem 4.1 hold if α is sufficiently large. The arguments for Theorems 4.2 and 4.3 then go through virtually unchanged.

Lemma 5.2. *For any $y < 0$, there exists $\alpha(y)$ such that if $\alpha > \alpha(y)$, then $\widetilde{\pi}(q)$ is a strictly unimodal function on $(0, 1)$ with a unique maximizer $\widetilde{q} \in (0, 1)$ and $\widetilde{q} > q^H$.*

Proof. Note that items (i) and (iii) in the proof of Lemma 5.1 hold regardless of y. Hence, we only need to show that item (ii) holds when α is sufficiently large. To do so, consider (6) and observe that for $q \in (0, 1/2)$, we have $\widetilde{\pi}'(q) \geq y - 2 + 2\alpha q + \log(1/q - 1)$ because $1/(1 - q) \leq 2$ for $q \in (0, 1/2)$.

For any y, there exists $q(y) \in (0, 1/2)$ such that $\log(1/q - 1) \geq 2 - y$ for all $q \in (0, q(y))$. Hence, $\widetilde{\pi}'(q) \geq y - 2 + \log(1/q - 1) > 0$ for all $q \in (0, q(y))$. Let $\alpha_1(y) = \frac{2-y}{2q(y)}$. For $\alpha > \alpha_1(y)$ we have $\widetilde{\pi}'(q) \geq y - 2 + 2\alpha q \geq y - 2 + 2\alpha q(y) > 0$ for $q \in [q(y), 1/2)$. It follows that $\widetilde{\pi}'(q) > 0$ for $q \in (0, 1/2)$ for $\alpha > \alpha_1(y)$.

For $q \in [1/2, q^H)$, we have $\widetilde{\pi}'(q) = p'(q)q + p(q) \geq p(1/2) = y + \alpha/2$. Therefore, $\widetilde{\pi}'(q) > 0$ for all $q \in [1/2, q^H)$ if $\alpha > \alpha_2(y) := -2y$.

Item (ii) now follows by taking $\alpha(y) = \max\{\alpha_1(y), \alpha_2(y)\}$. $\square$

Proof of Theorem 4.2. To prove part (i), note that $\overline{\pi}/\alpha = \widetilde{\pi}(\widetilde{q})/\alpha = p(\widetilde{q})\widetilde{q}/\alpha$. By Lemma 3.1 and Proposition 3.2, we have $q^H = \frac{1}{2} + \sqrt{\frac{1}{4} - \frac{1}{\alpha}}$ and $q^H < \widetilde{q} < 1$. It follows that $q^H \to 1$ and $\widetilde{q} \to 1$ as $\alpha \to \infty$. To complete the proof of part (i), we next show that $p(\widetilde{q})/\alpha \to 1$. To do so, observe that Lemma 5.3 below and the monotonicity of $p(q)$ on $(q^H, 1)$ imply that $G(\alpha) < p(\widetilde{q}) < p(q^H)$ for sufficiently large α where $G(\alpha) := y + \alpha - \frac{1}{2} - \log(2\alpha - 1)$. So

$$\frac{G(\alpha)}{\alpha} < \frac{p(\widetilde{q})}{\alpha} < \frac{p(q^H)}{\alpha}. \tag{8}$$

We have $G(\alpha)/\alpha \to 1$. Note that $1 - q^H = q^L$ and $q^L = 1/(\alpha q^H)$. Therefore, by (3),

$$\frac{p(q^H)}{\alpha} = \frac{y}{\alpha} + q^H + \frac{1}{\alpha}\left[\log q^L - \log q^H\right] = \frac{y}{\alpha} + q^H - \frac{1}{\alpha}\left[\log \alpha + 2\log q^H\right] \to 1.$$

From (8), we now have $p(\widetilde{q})/\alpha \to 1$, which completes the proof of part (i).

Next, we turn to part (ii). Lemma 5.5 below establishes that $\widetilde{q} < q^M$ for α sufficiently large. Theorem 4.1 implies that $\underline{\pi} = \widetilde{\pi}(q^M)$ for such α. Part (ii) now follows from Lemma 5.4 below. □

Lemma 5.3. *For $\alpha > \frac{1}{2}(1 + \exp(y - 1))$ we have $p(\widetilde{q}) > y + \alpha - \frac{1}{2} - \log(2\alpha - 1)$.*

Proof. Consider $q^+ = 1 - \frac{1}{2\alpha}$. From (6), we have $\widetilde{\pi}'(q^+) = y - \log(2\alpha - 1) - 1 < 0$, where the inequality holds for $\alpha > \frac{1+\exp(y-1)}{2}$. The unimodality of $\widetilde{\pi}(q)$ implies $q^H < \widetilde{q} < q^+$ for such α. The function $p(q)$ is decreasing on $(q^H, 1)$ so $p(\widetilde{q}) > p(q^+) = y + \alpha - \frac{1}{2} - \log(2\alpha - 1)$. □

Lemma 5.4. *$\widetilde{\pi}(q^M) \sim \log \alpha$ as $\alpha \to \infty$.*

Proof. Recall that $p(q^M) = p(q^L)$ by definition. So,

$$\widetilde{\pi}(q^M) = q^M p(q^L) = q^M\left[y + \alpha q^L + \log q^H - \log q^L\right] = q^M\left[y + \frac{1}{q^H} + 2\log q^H + \log \alpha\right]$$

where the final equality above uses $q^L = 1/(\alpha q^H)$. Therefore,

$$\frac{\widetilde{\pi}(q^M)}{\log \alpha} = q^M\left[\frac{y}{\log \alpha} + \frac{1}{q^H \log \alpha} + \frac{2\log q^H}{\log \alpha} + 1\right]. \tag{9}$$

Note that $y/\log \alpha \to 0$. In addition, $1/(q^H \log \alpha) \to 0$ and $2\log q^H/\log \alpha \to 0$ because $q^H \to 1$. Finally, we also have $q^M \to 1$ because $q^H < q^M < 1$. In view of (9), this completes the proof. □

Proof of Theorem 4.3. To begin, we derive an upper bound on $q^\ddagger$. Recall that $p(q^\ddagger) = p(\widetilde{q})$, so by Lemma 5.3 and the definition of $p(q)$, we have $y + \alpha q^\ddagger + \log\left(1/q^\ddagger - 1\right) > y + \alpha - \frac{1}{2} - \log(2\alpha - 1)$ for sufficiently large α. It follows that

$$\log\left(1/q^\ddagger - 1\right) > \alpha - \frac{1}{2} - \log(2\alpha - 1) - \alpha q^\ddagger.$$

Note that $q^\ddagger < q^L = \frac{1}{2} - \sqrt{\frac{1}{4} - \frac{1}{\alpha}} \le \frac{2}{\alpha}$ because $\left(\frac{1}{2} - \frac{2}{\alpha}\right)^2 \le \frac{1}{4} - \frac{1}{\alpha}$, where the final inequality holds because $\alpha > 4$. Thus,

$$\log\left(1/q^\ddagger - 1\right) > \alpha - \frac{5}{2} - \log(2\alpha - 1).$$

Hence,

$$1/q^{\ddagger} \;>\; \frac{\exp(\alpha - 5/2)}{2\alpha - 1} + 1.$$

Therefore,

$$q^{\ddagger} \;<\; \left[1 + \frac{\exp(\alpha - 5/2)}{2\alpha - 1}\right]^{-1} \;\leq\; C\alpha e^{-\alpha}$$

where $C = 2e^{5/2}$. Also, $p(q^H) = y + \alpha q^H + \log(1/q^H - 1) \leq y + \alpha + \log(1/q^H - 1) \leq y + \alpha$ because $\frac{1}{2} < q^H < 1$. For $\alpha \geq y$ we now have $p(q^H) \leq 2\alpha$. Consequently $\pi^{\ddagger} = p(\widetilde{q})q^{\ddagger} \leq p(q^H)q^{\ddagger} \leq C_2\alpha^2 e^{-\alpha}$ where $C_2 = 4e^{5/2}$ for all α sufficiently large.

To finish the proof, we will next use a similar argument to establish an (asymptotic) lower bound on $\pi^{\ddagger}$. Let $q^{\circ} = 1 - 1/\alpha$. Then $\widetilde{\pi}'(q^{\circ}) = y - \log(\alpha - 1) + \alpha - 2 > 0$ for α sufficiently large. Hence, $\widetilde{q} > q^{\circ}$. We also have $q^H = \frac{1}{2} + \sqrt{\frac{1}{4} - \frac{1}{\alpha}} < q^{\circ}$. Thus, $p(\widetilde{q}) < p(q^{\circ})$ because $p(q)$ is decreasing on $(q^H, 1)$. Therefore, $p(q^{\ddagger}) = p(\widetilde{q}) < p(q^{\circ}) = y + \alpha - 1 - \log(\alpha - 1)$, from which we obtain $y + \alpha q^{\ddagger} + \log\left(1/q^{\ddagger} - 1\right) < y + \alpha - 1 - \log(\alpha - 1)$. For $B := e/4$, steps similar to those above now yield

$$q^{\ddagger} > \left[1 + \frac{\exp(\alpha - 1)}{\alpha - 1}\right]^{-1} > \left[\frac{2\exp(\alpha - 1)}{\alpha} + \frac{2\exp(\alpha - 1)}{\alpha}\right]^{-1} = B\alpha e^{-\alpha},$$

where the second inequality holds because (i) $\alpha/2 < \alpha - 1$ (since $\alpha > 4$) and (ii) $\frac{2\exp(\alpha-1)}{\alpha} > 1$. Lemma 5.3 implies that $p(\widetilde{q}) > \alpha/2$ for α sufficiently large. So for $C_1 := B/2$, we now have $\pi^{\ddagger} = p(\widetilde{q})q^{\ddagger} \geq C_1\alpha^2 e^{-\alpha}$ for all α sufficiently large. This completes the proof. □

Lemma 5.5. *Given y, there exist α', α'' such that $\widetilde{q} > q^M$ for $\alpha < \alpha'$ and $\widetilde{q} < q^M$ for $\alpha > \alpha''$.*

Proof. Let α' be such that $q^M = 3/4$ for $\alpha = \alpha'$ (see Lemma 5.6 below). By Lemma 5.1, $\widetilde{\pi}(q)$ is unimodal and $\widetilde{q}$ is the unique solution to $\widetilde{\pi}'(q) = 0$. Thus, to prove $\widetilde{q} > q^M$ for $\alpha < \alpha'$, it suffices to prove $\widetilde{\pi}'(q^M) > 0$ for $\alpha < \alpha'$. By (6), $\widetilde{\pi}'(q^M) > 8q^M + \log(1/q^M - 1) - 1/(1 - q^M)$. It is easy to check that $8q + \log(1/q - 1) - 1/(1 - q) > 0$ for $q \in (0, 3/4)$. By Lemma 5.6, $q^M < 3/4$ for $\alpha < \alpha'$. Hence, we have established that $\widetilde{\pi}'(q^M) > 0$ for $\alpha < \alpha'$, and therefore $\widetilde{q} > q^M$ for $\alpha < \alpha'$.

Next we prove the existence of α''. It is easy to check that $\widetilde{\pi}(q^H) \sim \alpha$. Lemma 5.4 shows that $\widetilde{\pi}(q^M) \sim \log\alpha$. Consequently, $\widetilde{\pi}(q^H) - \widetilde{\pi}(q^M) \to \infty$. So, there exists α'' such that $\widetilde{\pi}(q^H) > \widetilde{\pi}(q^M)$ for $\alpha > \alpha''$. Lemma 5.1 implies $\widetilde{\pi}(q)$ increases on $[q^H, \widetilde{q}]$. Therefore, $\widetilde{q} < q^M$ for $\alpha > \alpha''$. □

Lemma 5.6. *q^M is continuous and increasing in α. In addition, there exists α such that $q^M = 3/4$.*

Proof. Let $C(\alpha) := p^L = p(q^L)$. By (3), the derivative of $C(\alpha)$ with respect to α is $C'(\alpha) = 1/2 - \sqrt{1/4 - 1/\alpha} = q^L$. From its definition, q^M is the unique solution to $p(q) - C(\alpha) = 0$ on the domain $q \in (q^H, 1)$. By the Implicit Function Theorem, q^M is continuous and differentiable in α. Differentiating $p(q^M) - C(\alpha) = 0$ with respect to α (and writing q_M rather than q^M for readability) gives us $\left(\alpha - [q_M(1-q_M)]^{-1}\right) q'_M = C'(\alpha) - q_M$. We have $C'(\alpha) - q_M < 0$ because $C'(\alpha) = q^L$ and $q_M > q^H$. In addition, $\alpha - [q(1-q)]^{-1} < 0$ for $q \in (q^H, 1)$. Therefore, $q'_M > 0$.

We have that $q_M \uparrow 1$ as $\alpha \to \infty$. Hence, in view of the continuity of q_M, the proof will be complete if we show that there exists α for which $q_M \leq 3/4$. To do so, it suffices to show there exists α for which $p(3/4) \leq p(q^L)$. We have $p(3/4) - p(q^L) = 3\alpha/4 - \log 3 - \alpha q^L - \log(q^H/q^L)$. As $\alpha \downarrow 4$, the expression for $p(3/4) - p(q^L)$ approaches $1 - \log 3 < 0$ because $q^L \to 1/2$ and $q^H \to 1/2$. It follows that $p(3/4) < p(q^L)$ for α sufficiently close to 4. □

Acknowledgment

This material is based upon work supported by the National Science Foundation under Grant Number CMMI 1462676.

References

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

[2] M. O. Ball and M. Queyranne. Toward robust revenue management: Competitive analysis of online booking. *Operations Research*, 57(4):950–963, 2009.

[3] 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.

[4] Ş. İ. Birbil, J. Frenk, J. A. Gromicho, and S. Zhang. The role of robust optimization in single-leg airline revenue management. *Management Science*, 55(1):148–163, 2009.

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

[6] B. Colson, P. Marcotte, and G. Savard. An overview of bilevel optimization. *Annals of Operations Research*, 153(1):235–256, 2007.

[7] W. L. Cooper, T. Homem-de Mello, and A. J. Kleywegt. Learning and pricing with models that do not explicitly incorporate competition. *Operations Research*, 63(1):86–103, 2015.

[8] A. V. Den Boer and D. D. Sierag. Decision-based model selection. Working paper, University of Amsterdam, 2016.

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

[10] 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.

[11] 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.

[12] 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.

[13] A. E. Lim and J. G. Shanthikumar. Relative entropy, exponential utility, and robust dynamic pricing. *Operations Research*, 55(2):198–214, 2007.

[14] G. Perakis and G. Roels. Robust controls for network revenue management. *Manufacturing & Service Operations Management*, 12(1):56–76, 2010.

[15] P. Rusmevichientong and H. Topaloglu. Robust assortment optimization in revenue management under the multinomial logit choice model. *Operations Research*, 60(4):865–882, 2012.

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

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

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

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

[20] R. Wang and Z. Wang. Consumer choice models with endogenous network effects. *Management Science*, 63(11):3944–3960, 2017.