Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Quality-Based Competition, Profitability, and Variable Costs

Chester Chambers, Panos Kouvelis, John Semple,

MANAGEMENT SCIENCE
Vol. 52, No. 12, December 2006, pp. 1884–1895
ISSN 0025-1909 | EISSN 1526-5501 | 06 | 5212 | 1884

DOI 10.1287/mnsc.1060.0581

Quality-Based Competition, Profitability, and Variable Costs

Chester Chambers
Cox School of Business, Southern Methodist University, Dallas, Texas 75275, cchamber@mail.cox.smu.edu

Panos Kouvelis
Olin School of Business, Washington University, One Brookings Drive, St. Louis, Missouri 63130, kouvelis@wuslt.edu

John Semple
Cox School of Business, Southern Methodist University, Dallas, Texas 75275, jsemple@mail.cox.smu.edu

We consider the impact of variable production costs on competitive behavior in a duopoly where manufacturers compete on quality and price in a two-stage game. In the pricing stage, we make no assumptions regarding these costs—other than that they are positive and increasing in quality—and no assumptions about whether or not the market is covered. In the quality stage, we investigate a broad family of variable cost functions and show how the shape of these functions impacts equilibrium product positions, profits, and market coverage. We find that seemingly slight changes to the cost function's curvature can produce dramatically different equilibrium outcomes, including the degree of quality differentiation, which competitor is more profitable (the one offering higher or lower quality), and the nature of the market itself (covered or uncovered). Our model helps to predict and explain the diversity of outcomes we see in practice—something the previous literature has been unable to do.

Key words: game theory; operations strategy; quality competition
History: Accepted by Sunil Chopra, operations and supply chain management; received April 28, 2003. This paper was with the authors 1 year and 4 months for 3 revisions.

1. Introduction

Consumers encounter numerous occasions where they must choose between competing products or services based on their respective combination of quality and price. Because the financial performance of the firm hinges on the overall attractiveness of its offering, companies invest considerable time and effort deciding on the combination they believe will achieve their goals in the marketplace. Insights into this problem have often been obtained through economic models of vertical product differentiation. In these models, profit-maximizing firms compete on two dimensions: product quality and product price. Most of these models assume that variable costs are independent of product quality and that market coverage (i.e., whether or not the whole market is served) is predetermined. The former assumption is troubling given the number of real-world examples where higher quality can only be obtained at significantly higher variable cost. The latter assumption is problematic because it is impossible to guarantee that a preselected market structure would logically result from profit-maximizing competition without additional ad hoc restrictions.

As a case in point, consider the "piano example" introduced by Gabszewicz and Thisse (1979). In that work, a customer considers the purchase of a single piano from either manufacturer A or B, with A offering the piano of higher quality. As Garvin (1990) observes, this parallels the situation existing between two major manufacturers of grand pianos, Steinway and Yamaha. Steinway, which is perceived as the high-quality producer, relies on a labor-intensive process utilizing expensive, highly skilled craftsmen who build each unit separately to strict material specifications. Yamaha uses a more automated process involving less skilled labor, less stringent material specifications, and fewer labor hours. Consequently, the variable costs differ significantly, with Steinway pianos being more expensive to produce. Neither Steinway nor Yamaha can ignore the dependence of variable costs on product quality when setting quality levels and prices. Moreover, while it might be sensible in the case of grand pianos to assume that the market is uncovered, this assumption is not justified in other situations.

To develop insights about the implications of the cost-quality relationship, we develop and analyze a simple vertical differentiation model of a duopoly. Following Ronnen (1991), we use H to represent the higher-quality provider and L to represent the lower-quality provider. Their respective quality levels are denoted by q_H and q_L (with $q_L < q_H$) and are selected from some interval $[0, q_{\max}]$ (the *quality domain* or *quality spectrum*). The associated variable costs for qualities q_L and q_H are c_H and c_L, respectively. We will show that the equilibrium price for H equals the variable cost of production plus a premium that increases with the quality differential $(q_H - q_L)$. Additionally, we show that the equilibrium profit for H is precisely $(q_H - q_L)$, which makes H's quality ambitions clear. If a pure-strategy equilibrium (in qualities) exists, H provides a high-quality/high-price product appealing to the most price insensitive segment of the market. In contrast, L's equilibrium price always equals c_H, and therefore L's profit depends in large part on the difference in variable costs. Moreover, it is the quality level offered by L that determines whether the market is covered (all customers buy) or uncovered (some customers do not buy).

For the case in which there is an industrywide cost-quality curve obeying mild properties (described in §4), we establish sufficient condition for the existence of a pure-strategy equilibrium in qualities. In this equilibrium, H always takes the highest possible quality, whereas L takes a position that depends on both the *curvature* of the variable cost function (defined formally in §4) and the *range* of variable costs. We find that increasing curvature pressures L to increase quality. This seems intuitively clear because increasing curvature of the variable cost function means higher qualities are available at lower costs. However the range of variable costs creates two types of markets, and L's response to increasing curvature is different in each type. If the market is such that the most price-sensitive customers are unwilling to pay even the variable cost of production for the highest-quality good, then L's quality will converge to H's. In this case, L will ultimately be more profitable than H, and the market will be uncovered. In contrast, if the market is such that the most price-sensitive customers are willing to pay at least the variable cost of production for the highest-quality good, then L's quality will not converge to H's. In this case, there is a boundary in the quality space (below H's position) that L will not cross, even if the marginal cost of increasing quality is negligible. This is because crossing the boundary sets off a price war with H, and L is much worse off for doing so. L therefore stops increasing quality at the boundary, and the resulting market will be covered. The profitability of the players depends on the position of this boundary, which is related to H's variable cost in an extremely simple way.

Our model therefore explains how simple properties of the variable cost function impact the configuration of quality positions in an industry. We prove that it is the cost function that (a) justifies the assumption that there will be at least two profitable positions in the market, (b) determines the separation in quality levels, (c) determines which quality position is most profitable, and (d) determines whether or not the market is covered. We demonstrate that subtle changes in the variable cost function can lead to very different market outcomes, thus serving as a cautionary tale about the robustness of results based on a narrow functional form.

2. Literature

Researchers in the fields of economics and marketing have investigated the manner in which products of different quality levels compete in the marketplace. In these works, the term "quality" refers to the level of some attribute or some scalar metric representing a vector of attributes (e.g., variety, functionality, reliability, etc.). The defining characteristic of quality in these treatments is that the marketplace consists of individuals who all agree that a higher level is always preferable to a lower level. For example, all other things being equal, higher-octane gasoline is preferable to lower-octane gasoline, more pixels are preferable to fewer pixels in issues of visual resolution, and more feedback from physicians is preferable to less feedback in health-care settings. Another interpretation of "quality" would be a weighted score of the item based on its attributes, such as *Consumer Reports* might provide. (See Tirole 1988 for an extended review.)

Mussa and Rosen (1978) consider a monopolist choosing quality positions to serve a market of heterogeneous customers. This basic model is extended to consider oligopolists competing on quality by Gabszewicz and Thisse (1979), Tirole (1988), and Donnenfeld and Weber (1992). Each of these works assumes that variable production costs are independent of product quality and that the market is covered, meaning that every potential customer receives positive utility from the purchase of at least one product offered. We relax these assumptions because of the abundance of real-world examples where neither claim holds, such as luxury automobiles, grand pianos, or plasma screen televisions.

The work of Moorthy (1988) has been particularly influential in this area. He appears to be the first to explicitly include variable costs in his framework by assuming $c(q) = \gamma q^2$, where q is product quality and γ is a positive constant. (Also, see extensions by Rhee 1996, Villas-Boas 1998, and Desai 2001.) Each of these works assumes that the market is uncovered and that variable costs are quadratic in product quality. We assume that the market can be covered or uncovered and that variable production costs are increasing and convex in quality (as argued by Moorthy 1988). We note that our model leads to quite different conclusions.

The works of Ronnen (1991) and Lehmann-Grube (1997) each assume that the market is uncovered. Ronnen (1991) ignores variable costs for the development of his main results, but briefly discusses how its inclusion affects his conclusion that increasing a minimum quality standard results in reducing the ratio of price to quality that the customers face in the marketplace. Ronnen's result stems largely from the fact that maximum differentiation is always optimal in his formulation when no variable costs are present. Our work will explain that maximal differentiation is often an irrational response for the producer of the lower-quality good. Lehmann-Grube (1997) ignores variable costs but does consider "development costs" that increase in quality but are sunk once production begins. The inclusion of such an expense would have no impact on most of our results (see Theorems 1–3).

The need to address this general problem with fewer restrictive assumptions was noted by Wauthy (1996). He appears to be the first to point out the sizable impact these assumptions have on the models' results. In his words, "the transition from uncovered market structures to covered ones is not smooth. The nature of competition changes when the market is covered. This is so because price competition becomes a pure battle for market shares" (p. 352). Moreover, he concludes that strategies deemed optimal under one market coverage assumption could be suboptimal once this assumption is relaxed. We point out that Wauthy made these observations while ignoring the presence of production costs, but our work is able to substantiate, clarify, and deepen them even in the presence of such costs. Table 1 summarizes this and other works.

In short, no previous analysis of this topic covers as much ground as ours, primarily because all prior works involve restrictive assumptions. Upon loosening these restrictions, we find that the results of prior works are not robust in that they do not hold for a more general setting. We proceed by characterizing each player's profit function, best-response function, and the resulting price equilibrium in §3. We then investigate the existence and properties of the quality equilibrium in §4 assuming an industrywide cost curve. We discuss our results and findings in §5.

Table 1 Classification of the Vertical Differentiation Literature

	Covered market	Uncovered market	Zero variable cost	Quadratic variable cost
Tirole (1988)	✓		✓	
Gabszewicz and Thisse (1979)	✓		✓	
Moorthy (1988)		✓		✓
Donnenfeld and Weber (1992)	✓		✓	
Rhee (1996)	✓			✓
Ronnen (1991)		✓	✓	
Lehmann-Grube (1997)		✓	✓	

3. Model Formulation and Development

We consider a heterogeneous market served by two producers whose products are differentiated by their respective quality levels. We model competition via a two-stage game. In the first stage, the producers choose their quality levels; in the second stage, they set prices. We ignore scale economies and assume that variable production costs $c(q)$ are strictly increasing in q. Each player has perfect information about his rival. We also assume that all potential customers will purchase at most one unit and have perfect information about product qualities and prices. The problem is analyzed by deriving equilibrium prices (given quality levels), then calculating equilibrium quality levels for players, anticipating the subsequent price competition.

3.1. Model Assumptions

Without loss of generality, we assume that the respective quality parameters associated with positions H and L are q_H and q_L and satisfy $q_H > q_L > 0$. Because there is typically an upper bound on quality due to the laws of physics, the state of technology, or the ability to improve the processes and procedures that improve quality, we assume a maximum quality level $q_{\max}$.

The prior literature has established a number of ways to calculate consumer surplus (or net utility) for a heterogeneous market. For example, several early works capture market heterogeneity using the net utility function $U(q, p; \alpha) = \alpha q - p$, where α is a *taste parameter* uniformly distributed on $[a, b]$ (see Moorthy 1988). Later works (e.g., Ronnen 1991 and Lehmann-Grube 1997) normalize the problem so that $\alpha \in [0, 1]$.

In our approach, we define surplus utility as $U(q, p: \theta) = q - \theta p$, where θ is uniformly distributed on $[0, 1]$. We call θ the *price sensitivity* parameter. With respect to previous works, our utility specification is equivalent to $U(q, p; \alpha) = \alpha q - p$ if one assumes that the taste parameter has the density $f(\alpha) = 1/\alpha^2$ for $\alpha \in [1, \infty)$. (See the online appendix on the *Management Science* website at http://mansci.pubs.informs.org.ecompanion.html for details.) Compared to the uniform distribution, this downward-sloping density implies that the majority of buyers have lower reservation prices and thus are more sensitive to price. This would appear to be a more reasonable assumption when considering discretionary purchases, such as luxury cars, boats, and grand pianos.

Our utility specification offers some unique advantages regarding the study of market coverage. The traditional specification $U(q, p; \alpha) = \alpha q - p$ for $\alpha \in [0, 1]$ implies that the market is uncovered for all variable cost functions because a buyer with $\alpha = 0$ will not buy any product at any price. The specification $U(q, p: \theta) = q - \theta p$ for $\theta \in [0, 1]$ does not imply that the market is uncovered, as even the extreme-point buyers find finite quality-price combination of positive utility. However, there are situations where the market will be strictly uncovered due to the shape of the variable cost function. For example, it is clear that the market must be uncovered if $q - 1 \cdot c(q) < 0$ for all q (equivalently, the average variable cost satisfies $c(q)/q > 1$ for all q). This is equivalent to stating that there are some markets where the relative price of quality rises so quickly that some customers will never buy. This is a strong condition that does not hold for many consumer goods and markets of interest.

3.2. Pricing Behavior of Player H

In this section, we assume that $q_H > q_L$ are given quality parameters, and we derive the best price response function for H. For fixed values of q_H, q_L, p_H, and p_L, customer utility U is a linear function of θ. Given two products (L and H), each customer prefers the product associated with the higher of these two linear functions or the horizontal axis. The horizontal axis represents the option of buying nothing, which we label the null product. Consequently, there are three situations that can occur. These are shown in Figures 1(a), (b), and (c).

These figures show that product H is preferred by all customers with a θ value between zero and some upper bound. This upper bound is determined by either the 0-quality/0-price null product (Figure 1(a)), the product L (Figure 1(b)), or the market's total size (=1) (Figure 1(c)). Thus, the market share for H is $\min(q_H/p_H, (q_H - q_L)/(p_H - p_L), 1)$. We define c_L as $c(q_L)$ and c_H as $c(q_H)$. Thus, we can write the profit for H (given p_L) as

$$\pi_H = \min\left(\frac{q_H}{p_H}, \frac{q_H - q_L}{p_H - p_L}, 1\right)(p_H - c_H). \tag{1}$$

Figure 1 suggests that a variety of quality-price combinations could occur. However, an analysis of (1) helps eliminate dominated price possibilities (all proofs are included in the appendix unless otherwise noted).

Proposition 1. *Given $q_H > q_L$ and price p_L, the optimal price response from H, $p_H = p_H^*(p_L)$, must satisfy*

Condition 1. $p_H \geq p_L(q_H/q_L)$.

Condition 2. $p_H \geq p_L + q_H - q_L$.

Figure 1 **Market Shares for H and L**

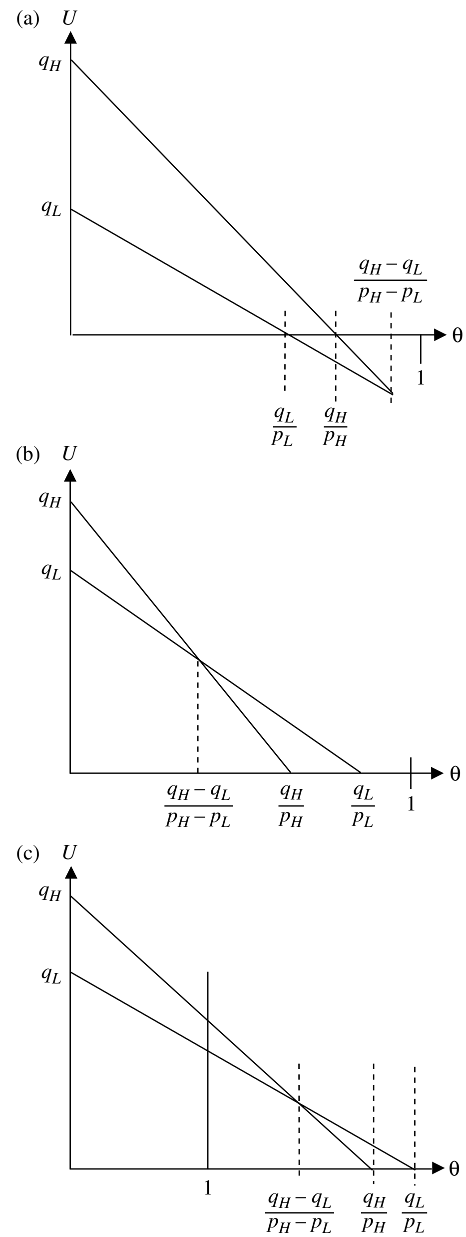

Condition 1 implies that H will never price his product so that $U_H(\theta) = U(p_H, q_H; \theta)$ intersects $U_L(\theta) = U(p_L, q_L; \theta)$ in the region $U < 0$, as in Figure 1(a). Condition 2 implies that H will never price his product so that the two functions intersect for $\theta > 1$, as in Figure 1(c). Consequently, the only possible situation is for the functions to intersect when $U \geq 0$ and $\theta \in [0, 1]$. In this case, the market share for H is $(q_H - q_L)/(p_H - p_L)$ as depicted in Figure 1(b). While we have implicitly assumed that L is in the market, we will show shortly that this is always the case.

We call a pair of prices *feasible* if Conditions 1 and 2 are met and both prices equal or exceed the respective product's variable cost. The set of feasible prices

Figure 2 Set of Feasible Prices

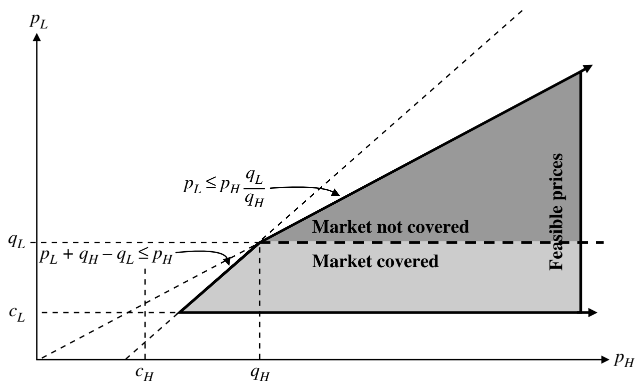

for the situation where $q_L > c_L$ is illustrated in Figure 2. Observe that Proposition 1 establishes a lower bound on H's price response based on the lowest conceivable price for L ($=c_L$). This lower bound for H is $\max(c_H, c_L + q_H - q_L, c_L \cdot q_H/q_L)$, which we refer to as H's *price floor*. We may assume that H always prices at or above this floor.

The conclusion that Conditions 1 and 2 must hold at equilibrium guarantees that the market share for H is precisely $(q_H - q_L)/(p_H - p_L)$. Thus, Equation (1) simplifies to

$$\pi_H = \left(\frac{q_H - q_L}{p_H - p_L}\right)(p_H - c_H). \tag{2}$$

We note that

$$\frac{\partial \pi_H}{\partial p_H} = \left[\frac{q_H - q_L}{(p_H - p_L)^2}(c_H - p_L)\right].$$

This term is strictly positive when $p_L < c_H$ and negative when $p_L > c_H$. Thus, if $p_L < c_H$, then H should increase his price. Conversely, if $p_L > c_H$, then H should lower his price until reaching some lower bound, such as his price floor. Using this insight, one can show that the best response for H (see the appendix) is

$$R_H(p_L) \equiv \begin{cases} \text{if } p_L < c_H, \quad p_H = \infty, \\ \text{if } p_L > c_H, \\ \quad p_H = \max\left(p_L\left(\dfrac{q_H}{q_L}\right), p_L + q_H - q_L\right), \\ \text{if } p_L = c_H, \\ \quad p_H \ge \max\left(c_H + q_H - q_L, c_H\dfrac{q_H}{q_L}\right). \end{cases} \tag{3}$$

Observe that when $p_L < c_H$, $p_H = \infty$. The infinite price is due to the existence of a completely price-insensitive customer ($\theta = 0$). If we eliminate this possibility, say by insisting $\theta \in [\varepsilon, 1]$ for some infinitesimal $\varepsilon > 0$, then H simply employs a large but finite price. Because this aspect of H's best response does not affect the equilibrium prices, it suffices to treat the infinite price as a symbolic response, e.g., "H sets a very high price." In either case, it is H's indifference to competitors who price below c_H that ensures L will enter the marketplace. This line of reasoning is formally stated in the following result.

Proposition 2. *Given $q_H > 0$, there exists $q_L > 0$ such that, if a price equilibrium exists, both H and L realize positive market shares and receive positive profits.*

Proposition 2 implies that at least two product positions are tenable because H will find it impossible to cover the market with a single product, and he will find it suboptimal to price his offering to exclude all market share from L.

3.3. Pricing Behavior of Player L

The description of the best response for H is simplified by the fact that a customer with $\theta = 0$ always prefers product H. Only the highest value of θ for which H is preferable is unknown. The situation for L is more complex because both the lower and upper bounds on the range of θ values of customers who prefer product L are setting specific. In the event that the quality-price combination that L selects covers the market ($q_L - p_L \ge 0$), L's market share is given by $(1 - (q_H - q_L)/(p_H - p_L))$. On the other hand, if $(q_L - p_L < 0)$, L's market share is $(q_L/p_L - (q_H - q_L)/(p_H - p_L))$. For all feasible prices (given qualities), the profit for L is therefore

$$\pi_L = \left\{\min\left(1, \frac{q_L}{p_L}\right) - \frac{q_H - q_L}{p_H - p_L}\right\}(p_L - c_L). \tag{4}$$

If we knew a priori that π_L would be maximized by covering the market, then it can be shown (see the appendix) that the optimal price for L would be $p_L^{\#}(p_H) = p_H - \sqrt{(p_H - c_L)(q_H - q_L)}$. This is possible only if $p_L^{\#}(p_H) < q_L$. Conversely, if we knew that π_L would be maximized by not covering the market, then the optimal price would be

$$p_L^{\#\#}(p_H) = p_H \cdot \frac{\sqrt{\dfrac{c_L q_L}{(q_H - q_L)(p_H - c_L)}}}{1 + \sqrt{\dfrac{c_L q_L}{(q_H - q_L)(p_H - c_L)}}}.$$

This may occur only if $p_L^{\#\#}(p_H) > q_L$. A detailed and somewhat lengthy analysis of L's price response leads to the following response function (see the appendix for complete derivation):

$$R_L(p_H) = \begin{cases} p_L = p_L^{\#}(p_H) & \text{if } p_L^{\#}(p_H) < q_L, \\ p_L = p_L^{\#\#}(p_H) & \text{if } p_L^{\#\#}(p_H) > q_L, \\ p_L = q_L & \text{otherwise.} \end{cases} \tag{5}$$

One can show that $p_L^{\#}(p_H) > p_L^{\#\#}(p_H)$ and that both $p_L^{\#}(p_H)$ and $p_L^{\#\#}(p_H)$ are increasing functions of p_H, provided H prices above his price floor and both players

Figure 3 Best Responses

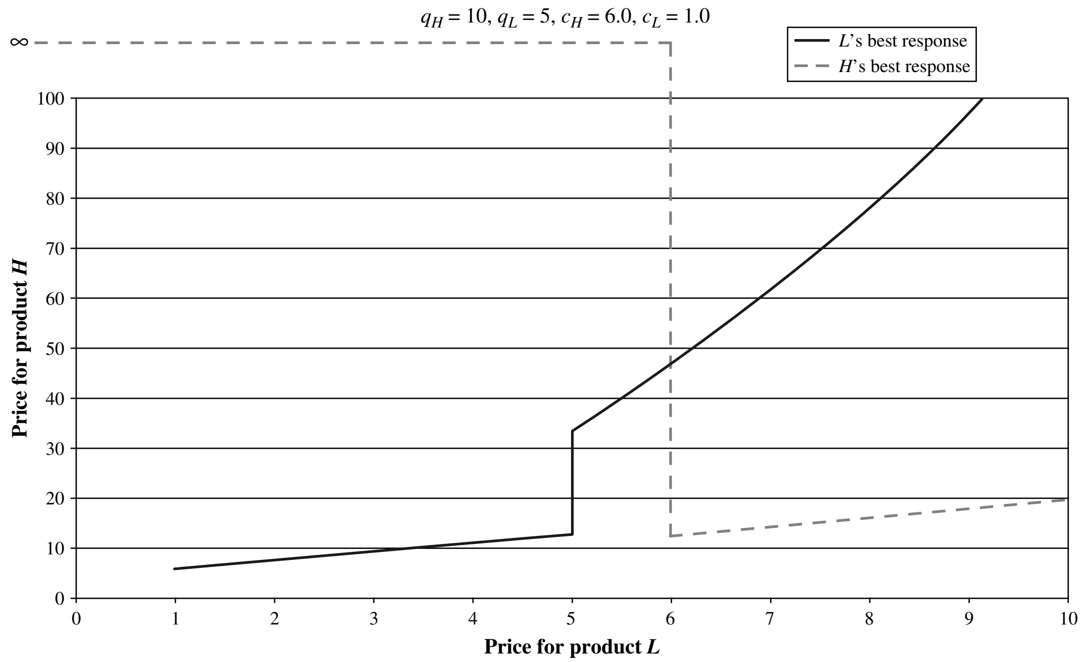

have strictly increasing cost functions. Consequently, $R_L(p_H)$ is well defined.

Figure 3 illustrates the best-response curves for the case where $q_H = 10$ and $q_L = 5$. For this example, we used the piecewise-linear, convex function $c(q) = 0.2q$ for $0 \le q \le 5$ and $c(q) = q - 4$ for $5 \le q \le 10$. This implies that $c_H = 6$ and $c_L = 1$. Because $p_L^{\#}$ and $p_L^{\#\#}$ are increasing functions of p_H, at most one of the following conditions can hold for a given p_H: either (i) $p_L^{\#} < q_L$, (ii) $p_L^{\#\#} > q_L$, or (iii) $p_L^{\#\#} \le q_L \le p_L^{\#}$. In the first case, L's best response is $p_L^{\#}$, and the market is covered. This is illustrated by the portion of L's response curve lying in the region to the left of $p_L = q_L = 5$. In the second case, L's best response is $p_L^{\#\#}$, and the market is uncovered. This is illustrated by the portion of L's response curve lying in the region to the right of $p_L = q_L = 5$. In the third case, L's best response is $p_L = q_L$. This is illustrated by the vertical segment of L's response curve that coincides with the line $p_L = 5$. The existence of this vertical segment implies that H would have to significantly increase his price to motivate L to raise his price above 5, which is the price above which L leaves the market uncovered.

3.4. Price Equilibrium

In general, the intersection of the best-response functions determines the Nash equilibrium in prices. The intersection of the response functions (3) and (5) leads to the following general result.

Theorem 1 (Price Equilibrium). *Let q_L and q_H be given quality levels satisfying $q_L < q_H$, and let the corresponding variable costs be c_L and c_H with $c_L < c_H$. The Nash equilibrium in prices is described as follows*:

Case 1. *If $c_H < q_L$, then*

$$p_L^* = c_H \quad \text{and}$$

$$p_H^* = c_H + \tfrac{1}{2}(q_H - q_L) + \sqrt{\tfrac{1}{4}(q_H - q_L)^2 + (c_H - c_L)(q_H - q_L)}. \tag{6}$$

Case 2. *If $c_H > q_L$, then*

$$p_L^* = c_H \quad \text{and}$$

$$p_H^* = c_H + \frac{(q_H - q_L)(c_H)^2}{2c_L q_L} + \sqrt{\frac{(q_H - q_L)^2 (c_H)^4}{4(c_L q_L)^2} + \frac{(q_H - q_L)(c_H)^2}{(c_L q_L)}(c_H - c_L)}. \tag{7}$$

Case 3. *If $c_H = q_L$, then $p_L^* = c_H$ and p_H^* is any price between those computed using* (6) *or* (7).

Observe that in all cases, L's best response to the equilibrium price chosen by H is the variable cost of H. H has numerous best responses for this price; these are represented by the vertical segment extending up to infinity in Figure 3. However, the price that drives L to his equilibrium price is determined by either (6) or (7). For the example in Figure 3 (recall that $q_L = 5$, $q_H = 10$, $c_L = 1$, and $c_H = 6$), Theorem 1 implies $p_L^* = 6$ and $p_H^* = \$46.45$. In this example, (7) was used to determine the equilibrium price because $q_L < c_H$. In other situations, (6) would be the appropriate formula. Straightforward algebra demonstrates that the equilibrium price for H set using (7) is substantially higher than that set using (6). Consequently, there is a discontinuity in H's equilibrium price as L's quality level passes through the value c_H. This discontinuity is important in explaining the nature of competition in this vertically differentiated duopoly.

To clarify the implications of this discontinuity, we return to our previous numerical example and vary L's quality position while holding H's position fixed at $q_H = 10$ (this assumption holds for all subsequent figures as well). We assume that equilibrium pricing is used by both players, and we plot values of q_L on the horizontal axis and market share values on the vertical axis. Figure 4(a) shows market share values for three products; product L, product H, and the option of buying nothing, which we label the "null product."

If q_L is close to 0, then H can offer his product at a very high price, serve a small segment of the market, and profit from the most price-insensitive customers. In this case, most of the market buys nothing because H's product is extremely expensive and L's product is "junk." For example, if $q_L = 1$, H's equilibrium price is \$1,632, L's equilibrium price is \$6, L serves 16% of the market, while H serves only 0.55% of the market. As q_L increases, player L gains market share, and this motivates H to drop his price. In the region between $q_L = 0$ and $q_L = 6$, both players gain market share as the portion of the market that is unserved shrinks. It is important to note that the market share for H is increasing in q_L over this range. This occurs because H has a great deal of room to drop his price, making it more attractive to customers that would buy L if H were to leave his price too high. For example, when $q_L = 5$, H's equilibrium price is only \$46.45; L serves 71% of the market, while H serves 12.4% of the market.

When $q_L = c_H$, the market is covered, and the market share for the null product is zero. From this point onward, the only way for L to gain market share is to take it from H. However, if L raises quality so that $q_L > c_H$, H sets his price using (6) instead of (7). One can think of the significantly lower price used by H as a "price war" response to L's quality ambitions. For example, comparing a scenario with $q_L = 5.99$ to one with $q_L = 6.01$, we see H's price drop from \$21.20 to \$12.50. Note that this drop in price for H corresponds to a jump in market share, from roughly 26% to 62%. As q_L rises above 6, the market share values become stable as H drops his price almost linearly, approaching a price of \$6 at $q_L = 10$.

3.5. Optimal Profits

The following result describes the profits for H assuming that H prices at or above his price floor.

Theorem 2. *Let q_L and q_H be given quality levels with $q_L < q_H$, and let the corresponding variable costs be c_L and c_H with $c_L < c_H$. At the Nash equilibrium in prices, we have*

$$\pi_H^* = q_H - q_L. \tag{8}$$

Note that Theorem 2 applies whether the market is covered or not.

The situation for the lower-quality manufacturer is more complex and involves two possibilities. If L can and does cover the market, his profit is

$$\pi_L^* = (c_H - c_L)\left\{\frac{-1+\sqrt{1+4(c_H-c_L)/(q_H-q_L)}}{1+\sqrt{1+4(c_H-c_L)/(q_H-q_L)}}\right\}. \tag{9}$$

In the region $q_L < c_H$, L chooses to leave the market uncovered, and his profit is

$$\pi_L^* = (c_H - c_L)\left\{\frac{q_L}{c_H} - \frac{q_L}{c_H} \cdot\left[\frac{2c_L/c_H}{1+\sqrt{1+4c_L q_L(c_H-c_L)/[(c_H)^2(q_H-q_L)]}}\right]\right\}. \tag{10}$$

We observe that the first factor in (9) and (10) is L's profit margin (per unit). The rest is L's market share. We also observe that the profit function (10) is units invariant with respect to costs. This implies that the shape of the variable cost function is a primary driver of profits.

Figure 4(b) plots the cost and profit functions for L (using equilibrium prices) for the numerical example described in Figure 4(a). Note that the profit-maximizing position is the "elbow" of the cost curve. When q_L crosses the boundary at 6, we enter the region where the market is covered. We see a discontinuity in L's profit function as we move from the setting depicted in (9) to the one depicted in (10). This is caused by H switching from (7) to (6) when setting equilibrium prices.

While the concept of an elbow in the cost curve is helpful in explaining some outcomes, it is not sufficient in general. Consider the example shown in Figure 5(a). Here, L can increase product quality with very little increase in production cost as long as q_L lies below 8. However, we see that L's profit is actually maximized when $q_L = 4$, which corresponds to the production cost for player H. This implies that the discontinuity in the profit function for L is another likely candidate for L's profit-maximizing position. However, we point out that in many cases, neither of these points (the elbow or the discontinuity) will be optimal for L. For example, consider the setting depicted in Figure 5(b). In this instance, L can increase quality very inexpensively up to $q_L = 2.5$, and the market is not covered unless $q_L > 8$, but the optimal position lies in between these two points.

Figure 4 (a) Market Shares for L, H, and Null Products; (b) Costs and Profits for L as a Function of q_L

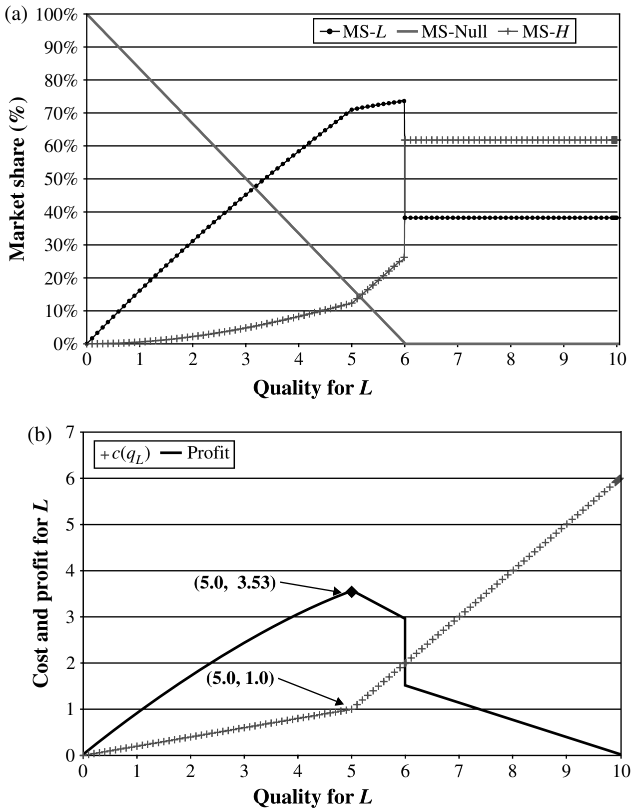

Figure 5 (a) Costs and Profits for L as a Function of q_L; (b) Costs and Profits for L as a Function of q_L

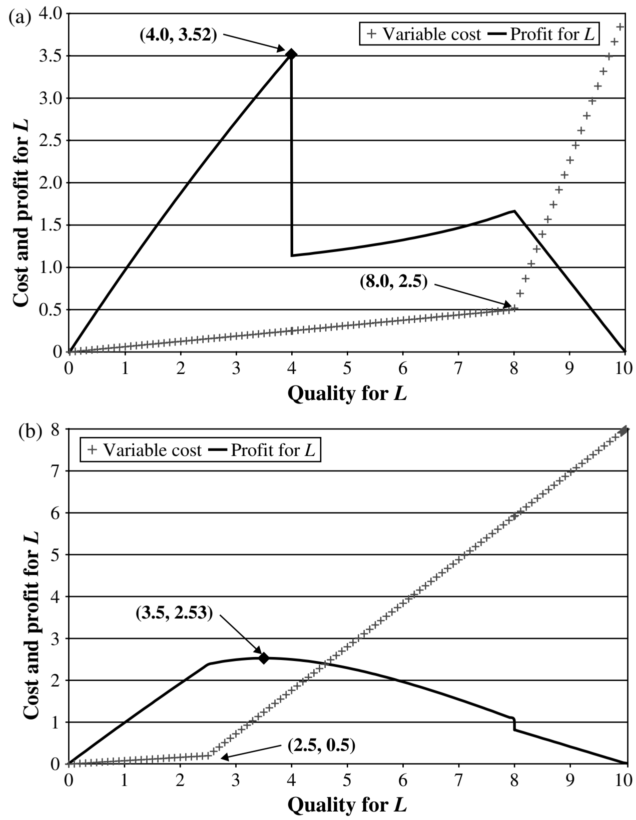

4. Production Costs and Quality Equilibrium

We note that Theorems 1 and 2, along with the derivations of Equations (9) and (10), assume only that $c(q)$ is nonnegative and nondecreasing in q. The application of these results facilitates the determination of the profit-maximizing position for players H and L in the general case. However, if we can be more specific in the description of the problem setting, we can develop additional, generalizable results.

4.1. Definitions

If each player has a different cost function, then the equilibrium quality levels (if they exist) are problem specific. Because our goal is to derive generalizable results, we now consider, consistent with previous literature, cases in which the industry is characterized by a single variable cost function, and some basic properties of that cost function are known. We label the class of functions that we consider Ω, and list its properties below.

Definition. Ω is the class of variable cost functions having the following properties:

Property 1. $c(0) = 0$.

Property 2. $c(q)$ is strictly increasing and twice differentiable on $[0, q_{\max}]$.

Property 3. $c(q)$ is convex on $[0, q_{\max}]$.

Property 4. The average variable cost function $c(q)/q$ is convex and log-concave on $[0, q_{\max}]$.

These four properties are satisfied by all of the variable cost functions used in the literature thus far (constant, linear, or quadratic). However, Ω also includes all functions of the form $c(q) = \gamma q^r e^{\beta q}$ for $\gamma > 0$, $\beta \geq 0$, $r \in \{1, [2, \infty)\}$[1] as well as polynomials of the form $c(q) = \gamma q \cdot \prod_{i=1}^{n}(\alpha_i + q)$ with $\alpha_i > 0$. This represents a broad cross section of convex increasing shapes. Moreover, it is straightforward to show that Ω is closed under multiplication.

All of our previous examples reflect functions that satisfy Properties 1–3. Property 4 is slightly more restrictive. A log-concave function f has the property that f'/f is nonincreasing, which implies that the derivative cannot change more rapidly than f. Thus, Property 4 forces the average variable cost function (and by extension the variable cost function) to exhibit a certain level of smoothness, even if $c(q)$ is rising rapidly.

We used piecewise-linear functions in our earlier examples to make the idea of an elbow clear. In these

[1] A somewhat larger class of functions includes the powers $1 < r < 2$ and ensures that all of our subsequent results hold. (See Chambers et al. 2004.) This class is more difficult to describe and thus is not presented here.

cases, the elbow is the point of maximum curvature. While Property 4 eliminates piecewise-linear functions, each of the variable cost functions used in our previous examples can be well approximated using functions in Ω. Thus, we would like to generalize the notion of an elbow for functions within Ω. To do this, we need a formal definition of curvature that considers the entire quality spectrum.

Definition. Suppose that $c(q)$ is a nonnegative, strictly increasing, convex function. Without loss of generality, suppose that $c(0)=0$. The curvature of $c(q)$ over $[0, q_{\max}]$, denoted by $K_{[0,\,q_{\max}]}(c(q))$, is

$$K_{[0,\,q_{\max}]}(c(q)) = \max_{q\in[0,\,q_{\max}]} \frac{q}{q_{\max}} - \frac{c(q)}{c(q_{\max})}. \tag{11}$$

The expression $q/q_{\max} - c(q)/c(q_{\max})$ measures the difference between $c(q)$ and the secant line passing through $(0,0)$ and $(q_{\max}, c_{\max})$, expressed as a fraction of $c_{\max}$. Consequently, $K_{[0,\,q_{\max}]}(c(q))$ is units invariant and $0 \le K_{[0,\,q_{\max}]}(c(q)) < 1$. Observe that the value of q that optimizes (11) is identical to the elbow in Figures 4(a), 5(a), and 5(b).

4.2. Structural Results

We can use the definitions stated in §4.1 to prove some fairly general results regarding L's quality position relative to H, including the following.

Theorem 3. *Suppose that $c(q)\in\Omega$. For $q_L\in[0, q_H]$, L obtains higher profits by selecting quality positions that satisfy $q_L \le c_H$.*

Observe that because L's equilibrium price is $p_L^* = c_H$, markets where $q_L \le c_H$ are uncovered unless L takes the position $q_L = c_H$. Thus, Theorem 3 only allows L to cover the market in a minimal sense, i.e., such that the most price-sensitive buyer has zero net utility. This is due to the price war response (6) of H. The significance of Theorem 3 is that as long we know that the cost function is in Ω, we may ignore Equation (9). This simplifies the development of additional results and leaves the possibility that all customers may still be served if L selects $q_L = c_H$. Given Theorem 3, we can safely state that L's best quality response to H is determined by solving the following mathematical program:

$$\max_{q_L\ge 0}\ \pi_L = (c_H - c(q_L))\left(\frac{q_L}{c_H}\right)\left[1 - \frac{2c(q_L)}{c_H} \cdot\left(1+\sqrt{1+4\frac{q_L\cdot c(q_L)}{(c_H)^2}\cdot\frac{(c_H-c(q_L))}{(q_H-q_L)}}\right)^{-1}\right] \tag{12}$$

$$\text{s.t. } q_L \le q_H,\quad q_L \le c_H.$$

Note that we can define $c_H = c(q_H)$ and $c_L = c(q_L)$, and that the functional form of (12) ensures that $\min(c_H, q_H)$ provides an upper bound on the optimal value of q_L.

To guarantee that the proposed equilibrium is unique, we need to show that $q_H = q_{\max}$ is optimal for H when he is not limited to positions above q_L. In many settings, it is simply impractical for firm H to take a position below that of firm L due to the differences in managerial styles, production systems, and the firms' competencies. However, in other settings such behavior may be feasible. To understand when this might occur, we need to compare profits for H over two regions: the regions above and below L's position. This is especially difficult because no closed-form expression for the optimal solution to (12) exists for the general case. However, our definition of curvature will prove useful in understanding the basic problem.

To make some general claims, we first need to determine an upper bound on L's position. This can be accomplished by decomposing L's profit function into the product of two functions, $f(q_L)$ and $g(q_L)$, where

$$f(q_L) = (c_{\max} - c(q_L))\left(\frac{q_L}{c_{\max}}\right) \quad \text{and}$$

$$g(q_L) = 1 - \frac{2c(q_L)/c_{\max}}{1+\sqrt{1+4\dfrac{q_L c(q_L)}{c_{\max}^2}\cdot\dfrac{c_{\max}-c(q_L)}{q_{\max}-q_L}}}. \tag{13}$$

If $c(q)$ is convex, then $f(q_L)$ is concave. We observe that the point maximizing $f(q_L)$ is the same point that defines the curvature (the "elbow") of $qc(q)$ over the interval $[0, q_{\max}]$. Because this point is of special importance in our subsequent developments, we denote it by q_K^*, i.e.,

$$q_K^* = \arg\max_{0\le q\le q_{\max}}\left(\frac{q}{q_{\max}} - \frac{qc(q)}{q_{\max}c(q_{\max})}\right). \tag{14}$$

We can now derive upper bounds on L's optimal response along with sufficient conditions that ensure the existence of a pure-strategy equilibrium.

Theorem 4. *Given $c\in\Omega$ and $0 < q_{\max} < \infty$,*

(a) *The solution to* (12) *satisfies $q_L^* \le q_K^*$, and the equilibrium profit for H is bounded below by $q_{\max} - q_K^*$.*

(b) *If $q_{\max} - q_K^* > c(q_K^*)$, then the only pure-strategy equilibrium is given by taking $q_H = q_{\max}$ and $q_L = q_L^*$, where q_L^* solves* (12).

(c) *If $qc'(q)/c(q)$ is nondecreasing on $(0, q_{\max}]$ and $q_{\max} - q_K^* > q_K^* K_{[0,\,q_{\max}]}(qc(q))$, then the only pure-strategy equilibrium is given by taking $q_H = q_{\max}$ and $q_L = q_L^*$, where q_L^* solves* (12).

We remark that Property 4 implies that $g(q_L)$ as defined in (13) is nonincreasing on $[0, q_{\max}]$, which, in turn, implies that the solution to (12) satisfies the upper bound $q_L^* \le q_K^*$ stated in part (a). This, in turn, establishes a lower bound of $q_{\max} - q_K^*$ for the profit of H on the region above L. Part (b) insists that the lower bound for H's profit in the region above L exceeds the upper bound $(=c(q_K^*))$ on H's profit in the region below L. When this condition is satisfied, a pure equilibrium exists (and it is given by $q_H^* = q_{\max}$ and $q_L = q_L^*$). However, the conditions for (b) are not scale invariant with respect to cost. Part (c) provides a second set of sufficient conditions that is scale invariant with respect to costs. The same type of profit comparison used in (b) is done in part (c), except this time the upper bound on H's profit in the region below L is $q_K^* K_{[0,\, q_{\max}]}(q \cdot c(q))$. The additional assumption on $qc'(q)/c(q)$ is needed to ensure that $K_{[0,\, q_H]}(q \cdot c(q))$ is a nondecreasing function of q_H, which implies $K_{[0,\, q_{\max}]}(q \cdot c(q)) \ge K_{[0,\, q_L^*]}(q \cdot c(q))$. This simplifies the upper bound on H's profit in the region below L. The nondecreasing assumption on $qc'(q)/c(q)$ is not particularly onerous; it holds for all of the functional forms discussed for Ω in §4.1.

Observe that Theorem 4 does not say anything about a lower bound for L's position in the pure-strategy equilibrium. To derive a lower bound, we again take $q_H = q_{\max}$. If $c_{\max} \ge q_{\max}$, then the market is necessarily uncovered. One may then insert the value q_K^* from (14) into (10) to obtain a lower bound on L's optimal profit, π_L^*. It can then be shown that

$$\pi_L^* > \frac{q_{\max}^2}{q_K^*}\big[K_{[0,\, q_{\max}]}(q \cdot c(q))\big]^2 > q_{\max}\big[K_{[0,\, q_{\max}]}(q \cdot c(q))\big]^2.$$

We also observe from (10) that $\pi_L(q_L) < q_L$, from which it follows that

$$q_L^* > q_{\max}\big[K_{[0,\, q_{\max}]}(q \cdot c(q))\big]^2. \qquad (15)$$

This implies that increasing the curvature of $c(q)$ eventually pressures L to increase quality.

In the case where $c_{\max} < q_{\max}$, L could conceivably cover the market. Inserting $q_L = c_{\max}$ into (10) and again using the inequality $\pi_L(q_L) < q_L$ yields the inequality

$$q_L^* > \pi_L^* > \left(1 - \frac{c(c_{\max})}{c_{\max}}\right)^2 c_{\max}. \qquad (16)$$

Because increasing the curvature of $c(q)$ drives the term $c(c_{\max})$ to zero, (16) implies that L would eventually be pressured to raise quality up to $c_{\max}$. When this occurs, the market is covered and H will chose a point above q_L because the profit available to H by choosing a point below L is bounded by $c(c_{\max})$ (see (12)), whereas the profit available to H by staying at $q_{\max}$ is bounded below by $q_{\max} - c_{\max}$. For sufficiently high curvature, $c(c_{\max}) < q_{\max} - c_{\max}$.

Figure 6 (a) Costs and Profits for L as a Function of q_L; (b) Costs and Profits for L as a Function of q_L

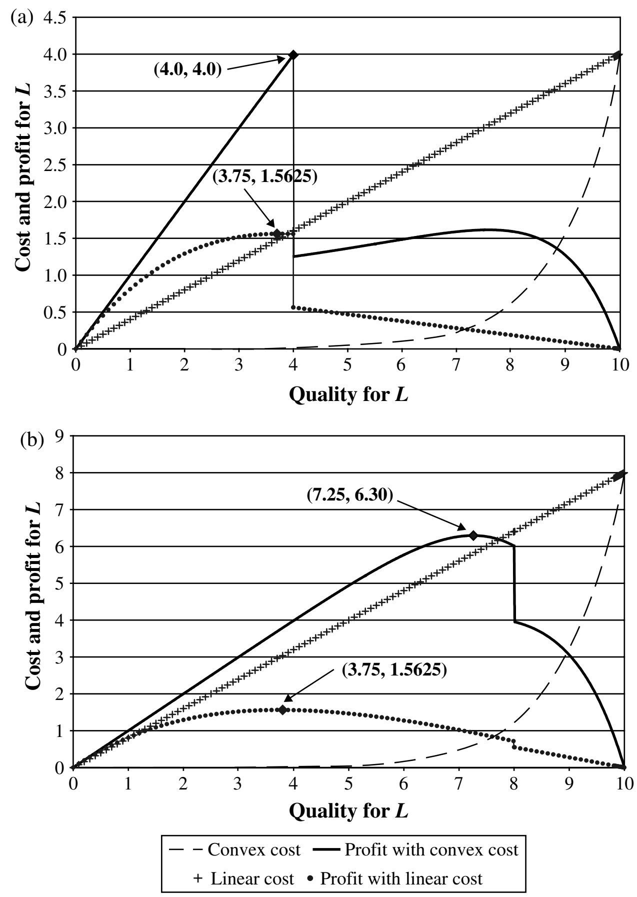

4.3. Curvature and Profitability

To make the curvature-profitability connection concrete, we consider two related examples, each comparing two cost functions with $q \in [0, 10]$. In Figure 6(a), we consider a linear cost function $(c(q) = 0.4q)$ versus the function $c(q) = 1 \cdot 10^{-5} q^3 e^{0.6q}$. Both functions are members of Ω, but the latter approximates the cost function in Figure 5(a) and exhibits significantly greater curvature than the linear function. The pure quality equilibrium is $q_H^* = 10$ (for both cases), while $q_L^* = 3.75$ with the linear cost function and $q_L^* = 4$ given the convex cost function. Observe that in this example, $c_{\max} = 4$ plays a critical role in constraining L's quality level and neutralizing the effects of curvature. This protects H's high-quality position and ensures that H has superior profits (6.25 for linear costs, 6.00 for convex costs). Nevertheless, higher curvature is still better for L ceteris paribus. For linear costs (no curvature), L's profit is 1.56; for convex costs (greater curvature), L's profit rises to 4.0. Note that the market is covered given convex costs, whereas it is uncovered for linear costs.

Figure 6(b) compares results when ($c(q) = 0.8q$) or $c(q) = 2 \cdot 10^{-5} q^3 e^{0.6q}$ and $c_{\max} = 8$. Note that the curvature is unchanged because each cost function is simply multiplied by two. However, lifting the maximum cost to eight (increasing the range of $c(q)$) allows the impact of curvature to be more fully realized. Comparing the case with linear costs to the case with convex costs, we see that L's equilibrium profit increases from 1.5625 to 6.3, while H's profit drops from 6.25 to 2.75. Thus, the effect of increasing curvature and increasing $c(q)$'s range has made L the more profitable position. In doing so, the market has become uncovered for both linear costs and convex costs. We also note that at least one of the two sufficient conditions in parts (b) and (c) of Theorem 4 holds for each cost function used in Figures 6(a) and 6(b).

5. Summary and Conclusions

We have investigated a duopoly where competitors compete on quality and price in a two-stage game. The analysis of this game has led us to several observations. First, we observe that given quality levels $q_L < q_H$, there exist equilibrium prices that allow both positions to profitably coexist. These results do not require any assumption other than that the lower-quality product costs less to produce than the higher-quality product $c_L < c_H$. Moreover, we find that the equilibrium price for L always equals the variable production cost for H. The equilibrium price for H equals the variable production cost for H plus a premium that depends on the separation in quality. This premium follows different formulas depending on whether L covers the market or not. This leads to a discontinuity in equilibrium prices and profits. Under loose conditions detailed in Theorem 3, L will never attempt to strictly cover the market.

Additionally, we find that both high- and low-quality positions benefit from product differentiation, but in different ways. The high-quality position benefits because H's profits are completely determined by the separation in quality. Consequently, H always takes the highest possible quality position. The low-quality position benefits because L's profits are primarily driven by the separation in variable costs. Variable cost functions exhibiting low curvature keep competitors' qualities far apart. As curvature increases, L can (and will) increase quality while maintaining the cost differential, thus gaining market share and improving profits. If the market is such that the most price-sensitive customer is willing to pay the variable cost of production for the highest-quality product, then rising curvature allows L to increase quality until the market is covered. In this case, the competitors' quality levels cannot converge, and the impact to H's profits may be small. On the other hand, if this is not the case, then the market is necessarily uncovered and rising curvature allows L to box H into a quality corner. This implies that their quality positions converge and H's profits are dramatically reduced.

There is some anecdotal evidence to support this description. For instance, returning to our grand piano example, we notice that Yamaha prices its units significantly below the prices for comparably sized Steinways. Yamaha's pricing and reputation for quality suggests that relatively high-quality levels are possible at relatively low production costs. (See Gourville and Lassiter 1999 for supporting price and production information.) On the other hand, Steinway achieves the industry's highest-quality levels for its grand

Figure 7 Quality Positions in the Piano Example

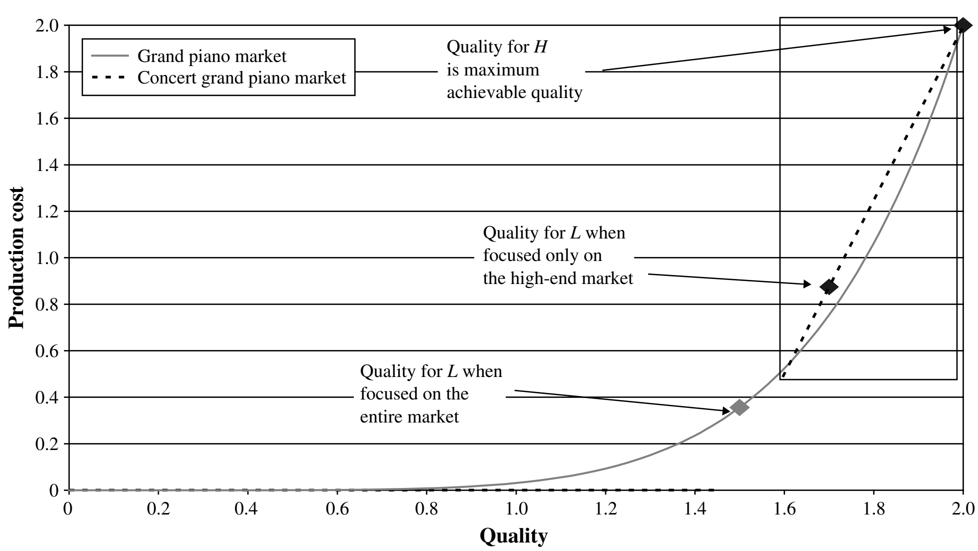

pianos by incurring higher labor and material costs (e.g., Steinway reports that it must discard half of the wood it purchases for their units because after an extensive curing period, the wood's density is unsuitable for further use).

These facts collectively suggest that there is significant curvature in the variable cost curve, as shown in Figure 7. Given such a relationship, our model suggests that the products for players L and H would be approximately positioned as shown in the figure. In this market, our model predicts that L (Yamaha) should be more profitable than H (Steinway) because of its larger market share and significant variable cost differential. This prediction is consistent with industry information. For example Gourville and Lassiter (1999) report that in 1994, Steinway delivered roughly 2,700 grand pianos, while Yamaha delivered roughly 20,000.

On the other hand, if we restrict attention to the market for concert grand pianos, the situation is quite different. For this market segment, Gourville and Lassiter (1999, p. 4) report that Yamaha mimics Steinway's design and production techniques. The variable cost curve should exhibit less curvature because the production technology is roughly proportional to the labor required (i.e., linear variable costs). This situation is better represented by the rectangular box in Figure 7. Here, our model predicts that H will be more profitable. Industry analysts indicate that Steinway holds this position and is clearly the most profitable in this space because over 90% of all classical music concerts featuring a piano soloist utilize a Steinway concert grand piano.

In either market, our model offers a reasonable explanation for the relative advantages and disadvantages of different quality positions as consequences of the variable cost of quality.

An online supplement to this paper is available on the *Management Science* website (http://mansci.pubs.informs.org/ecompanion.html).

References

Chambers, C., P. Kouvelis, J. Semple. 2004. Quality-based competition, profitability, and variable costs. Working paper, Cox School of Business, Southern Methodist University, Dallas, TX.

Desai, P. S. 2001. Quality segmentation in spatial markets: When does cannibalization affect product line design? *Marketing Sci.* **20**(3) 265–283.

Donnenfeld, S., S. Weber. 1992. Vertical product differentiation with entry. *Internat. J. Indust. Organ.* **10** 449–472.

Gabszewicz, J. J., J.-F. Thisse. 1979. Price competition, quality, and income disparities. *J. Econom. Theory* **20** 340–359.

Garvin, D. A. 1990. Steinway and Sons. Teaching note, Harvard Business School Publishing, Boston, MA.

Gourville, J. T., Joseph B. Lassiter, III. 1999. Steinway and Sons: Buying a legend. Teaching note, Harvard Business School Publishing, Boston, MA.

Lehmann-Grube, U. 1997. Strategic choice of quality when quality is costly. *Rand J. Econom.* **28**(2) 372–384.

Moorthy, K. S. 1988. Product and price competition in a duopoly. *Marketing Sci.* **7**(2) 141–168.

Mussa, M., S. Rosen. 1978. Monopoly and product quality. *J. Econom. Theory* **18** 301–317.

Rhee, Byong-Duk. 1996. Consumer heterogeneity and strategic quality decisions. *Management Sci.* **42**(2) 157–172.

Ronnen, U. 1991. Minimum quality standards, fixed costs, and competition. *RAND J. Econom.* **22** 490–504.

Tirole, J. 1988. *The Theory of Industrial Organization*. MIT Press, Cambridge, MA.

Villas-Boas, J. M. 1998. Product line design for a distribution channel. *Marketing Sci.* **17**(2) 156–169.

Wauthy, X. 1996. Quality choice in models of vertical differentiation. *J. Indust. Econom.* **44** 345–353.