WORKING PAPERS

Better Product at Same Cost,
Lower Sales and Lower Welfare

David J. Balan
George Deltas

WORKING PAPER NO. 312

June 2012

BUREAU OF ECONOMICS
FEDERAL TRADE COMMISSION
WASHINGTON, DC 20580

Better Product at Same Cost, Lower Sales and Lower Welfare*

David J. Balan
Bureau of Economics
Federal Trade Commission
600 Pennsylvania Ave NW
Mail Drop NJ-4264
Washington, DC 20580
dbalan@ftc.gov

George Deltas
Department of Economics
University of Illinois
at Urbana-Champaign
1407 W. Gregory
Urbana, IL 61801
deltas@illinois.edu

June 2012

Abstract

We analyze the effect of product quality on the output of a high-quality dominant firm facing a low-quality competitive fringe. Using a standard vertical differentiation model, we show that profit maximizing output decreases with product quality when the dominant firm's marginal cost is lower than that of the fringe, is independent of quality when marginal cost is the same for all firms, and is increasing in quality when the dominant firm's marginal cost is higher than that of the fringe. The driving force behind this result is that an increase in product quality does not cause a parallel shift in the dominant firm's residual demand, but rather causes it to pivot. This, in turn, causes the dominant firm's marginal revenue curve to rotate, rather than shift outwards, resulting in inwards movement around the equilibrium output when the dominant firm's marginal cost is lower than the fringe's. Equally strikingly, higher quality at the original marginal cost may result in all consumers being weakly worse off, with some being strictly worse off. Similar results can be obtained without a competitive fringe, but only under some more restrictive conditions.

JEL Classification Codes: L15, L13.

Keywords: Product innovation, vertical differentiation, dominant firm, competitive fringe.

*We would like to thank Roberto Burguet for insightful discussion. We also thank Patrick DeGraba, Ian Gale, Daniel Hosken, Sonia Jaffe, Justin Johnson, David Meyer, Lars Persson, Michael Salinger, David Schmidt, Brown Bag participants at the FTC, and seminar participants at the 2009 International Industrial Organization Conference for helpful comments. Any errors are ours. The views expressed in this article do necessarily represent those of the Federal Trade Commission or of any individual Commissioner.

1 Introduction

It is now understood that it can be profitable for a firm to take an action that increases the willingness-to-pay of its more likely customers, even at the cost of decreasing the willingness-to-pay of its less likely customers. Such an action effectively rotates the firm's demand curve through an interior point. The firm may be trading off fewer sales with a higher profit margin per sale (Johnson and Myatt 2006). In this paper, we obtain a similar but more striking result: under one quite common competitive environment, and using the canonical model of consumer preferences for vertically differentiated products, we show that an action that increases the willingness-to-pay for all of a firm's consumers, but does not increase its marginal cost, results in a reduction in that firm's sales. Moreover, it is possible that no consumer is made better off.

The competitive environment that we consider includes one single-product dominant firm and a competitive fringe. What makes the firm "dominant" is that it produces a higher quality product at equal or lower marginal cost (it may have higher fixed costs, such as R&D expenditures). We follow standard practice by assuming that the small rivals are price takers characterized by marginal cost pricing, while the dominant firm is strategic and behaves as a monopolist with respect to the residual demand.[1] We consider an increase in the quality of the dominant firm's product that is not accompanied by any change in its marginal cost, such as from a technological innovation (quality increases that are associated with marginal cost increases are less interesting to analyze, as it will become clear below). This quality increase causes the willingness-to-pay for the dominant firm's product to increase for all consumers. The quality of the fringe's product is assumed to remain unchanged.

Our first main result is that such an increase in willingness-to-pay has no effect on the dominant firm's profit-maximizing quantity if the marginal cost of the dominant firm is equal to that of the fringe, and causes quantity to decrease when the dominant firm's cost is below that of the fringe. What is surprising about this result is not that higher quality *may* cause the dominant firm's quantity to decrease, but rather that under this standard model, it *must* cause quantity to (weakly) decrease.[2]

[1]A high-quality producer may co-exist with a low-quality/high-cost fringe, because even though it could drive the fringe out of the market by charging a low price, it makes higher profits from a higher price and fewer sales.

[2]It is trivial to obtain such a result in a highly artificial environment. For example, suppose a monopolist optimally sells to half of the potential consumers. Next suppose that the willingness-to-pay of consumers in the top quartile of the demand goes up by a large amount, E, while the willingness-to-pay of the remaining consumers goes up a small amount,

Moreover, a variant of our model dispenses with the competitive fringe and still obtains the same result, albeit requiring some more stringent assumptions on consumer preferences. To highlight the mechanism that generates this outcome, we also examine a set of consumer preferences under which this main result no longer generally holds. However, even in this case, the effect we identify in the baseline model is still present and may still predominate.

Our second main result is that a quality increase of the dominant firm's product has an ambiguous effect on aggregate consumer welfare, even when it comes at a zero cost. Moreover, such a quality increase makes some consumers strictly worse off, and can possibly make *all* consumers weakly worse off. In contrast, a quality increase of the product offered by the fringe firms not only increases aggregate consumer welfare, but it makes all consumers strictly better off.

Some markets are characterized by a dominant firm selling a high-quality product competing with a number of much smaller rivals that sell a lower quality product but have equal or higher production costs. Indeed, this is often the reason for a firm's market dominance. Good examples include innovative consumer electronics products, many of which have attributes that most consumers regard as high quality, but that are not the result of the use of higher cost inputs. Rather, the high quality is the result of higher levels of R&D effort, or simply of superior ability at product design. Another example is any firm whose product enjoys patent protection and which competes against a fringe that sells imperfect substitutes to the patented product, such as a branded pharmaceutical company selling a latest-generation patented drug that competes against generic versions of older substitute drugs. In settings like this, where the dominant firm's superior quality comes from its innovative product characteristics or from patent protection, rather than from the use of more expensive inputs, there is no reason why a dominant firm must have higher marginal costs than the fringe firms. In fact, these costs may well be lower, especially if the dominant firm can reap economies of scale or has stronger bargaining power than its rivals in negotiating terms with input suppliers.

There is a substantial literature dealing with competition between firms whose products are vertically differentiated by quality. Key early papers in the literature include Gabszewicz and Thisse (1979) and Shaked and Sutton (1983) which show that, unlike in models of horizontal differentiation, the number of firms in the industry does not generally get arbitrarily large as fixed costs approach

---

ϵ. For a sufficiently high value of E, the monopolist will raise his price and serve only the consumers in the top quartile of demand.

zero.[3] This literature provides the vertical quality differentiation framework that we use in this paper.

The most relevant paper to ours is Johnson and Myatt (2006). That paper explores the effects of demand "rotation," by which the authors mean a change in a product's attributes (or in consumer perception of those attributes) that makes some consumers like the product more and others like it less. The paper then analyzes when it is profitable for firms to induce such rotations (i.e., when the benefits of having some consumers like the product more outweigh the costs of other consumers liking it less). In their framework, optimal firm strategies consist of attempting to widen product appeal and go for market share, or increase appeal to a consumer niche and go for high prices. Unlike their model, ours is one of pure vertical differentiation; all consumers value a high-quality product more than a low-quality product, but they differ in how much more.

2 The Model

2.1 Modeling environment

Consider a standard vertically differentiated product category, in which a product is fully described by a single important attribute, that we refer to as "quality." This could be capacity (for the case of jump drives, RAM memory, or hard-drives), speed (for computing devices), fuel efficiency (for heaters and furnaces), the perceived therapeutic value of a drug, battery life for a hand-held portable device, etc. In these cases, the product attribute is essential, i.e., its complete absence from a product would make that product worthless to all consumers. In other cases, the attribute may be important but not essential, and the product without it may be of nearly equal functionality. For example, the resolution of a cell-phone camera may be valuable but not essential as cell-phones with no camera (zero resolution) are of positive value and are good substitutes for making phone calls. The product attribute takes on a numerical value, where a higher value means the product is more desirable.[4]

[3] Subsequent papers develop this idea further. Choi and Shin (1992) consider whether the high and low quality firms will between them "cover the market." Frascatore (1999) considers the case where the inputs necessary to produce a higher quality product are in fixed supply and so must be competed for. Lehmann-Grube (1997) and Motta (1993) analyze the case where the cost of providing quality is significant. Noh and Moschini (2006) analyze how quality might be strategically chosen to deter entry.

[4] Some of the prior literature provides specific micro-foundations for the nature of the product attribute, distinguishing between those that are perfect substitutes for quantity (e.g., a razor that can do more shaves than another razor), those that allow a product to remain in use for a longer calendar time (i.e., to delay obsolescence), those that increase the reliability of the good, and others (see Beath and Katsoulacos 1991 for a detailed discussion of this literature). In this

There is a unit mass of consumers, who differ in the marginal willingness-to-pay for the attribute. In particular, the preferences of consumer i for product j are described by the indirect utility function

$$U_{ij} = V_i + \theta_i g(x_j) - P_j, \tag{1}$$

where V_i is the willingness of consumers to pay for the product in the absence of the attribute, θ_i is the marginal willingness of consumer i to pay for a unit increase in the attribute, x_j is the value of the attribute for product j, $g(\cdot)$ is a continuously differentiable and monotonically increasing function, and P_j is the price of product j. V_i is distributed with some (possibly degenerate) marginal distribution $H(V)$ on the interval $[V_{MIN}, V_{MAX}]$ (note that in most other papers on vertical differentiation, the value of V_i is the same for all consumers or even set to zero). The parameter θ_i is distributed with marginal distribution $F(\theta)$ with support $[\theta_{MIN}, \theta_{MAX}]$. The value of θ_{MIN} could be as low as 0, while the value of θ_{MAX} could be arbitrarily high. The dispersion in θ could be driven by differences in consumer income or by differences in the direct utility function.[5] The correlation or joint distribution of V_i and θ_i need not be specified as it has no bearing on the results. In what follows, we never compute the profit-maximizing level of the product attribute. Rather, we consider the effect of changes in that level regardless of the source of the change, whether exogenous or endogenous, as long as they don't affect the firm's marginal cost. Consumers have the option of making no purchase and earning a utility of zero.

A dominant firm sells a product of quality x_1, and faces a perfectly competitive fringe which sells products of a lower quality x_0 at a price equal to their (constant) marginal cost c_0.[6] Assumption 1, which is formally stated below, ensures that in equilibruim all consumers with values of θ_i and V_i such that $V_i + \theta_i g(x_0) - c_0 > 0 \Rightarrow \theta_i > \frac{c_0 - V_i}{g(x_0)}$ purchase some version of the product, and all those

paper, we follow the bulk of the recent literature in simply treating the attribute as something that consumers are willing to pay for, without being explicit about its nature or the way that it affects the product's use.

[5]Much of the early literature on vertical differentiation assumes that consumers have the same preferences but different incomes. However, even in that early literature it was clear that differences in income could be reinterpreted as differences in preferences (Gabszewicz and Thisse 1979, Gabszewics, Shaked, and Sutton, 1986), and that a combination of income and preference differences would generally yield the same results (Shaked and Sutton 1983).

[6]Perfectly competitive pricing follows trivially in our model if firms choose prices, given that the products of the fringe firms are perfect substitutes. More generally, the assumption that small firms are non-strategic is standard in models where a dominant firm faces a competitive fringe, and approximates the solution to a game between a firm that is large (in equilibrium) and many smaller (in equilibrium) rivals. It is straightforward to show this, for example, under Cournot competition between a firm with $MC = q$ and N rivals with $MC = Nq$, where N is large.

with lower values of θ_i (and V_i) do not. We now analyze the effect of a change in the dominant firm's quality x_1, holding its cost c_1 constant. This can be thought of temporally, with the dominant firm starting at an initial quality level and then improving, or equivalently as comparing a dominant firm with a particular quality level to an alternative situation in which its quality is even higher. In what follows, we perform comparative statics with respect to increases in x_1, holding the dominant firm's marginal cost constant at c_1. We focus on the case of $c_1 \leq c_0$, i.e., in the case where the dominant firm is dominant in both product quality and marginal cost, but also discuss for completeness the case of $c_1 > c_0$. We do not consider simultaneous changes in production costs and the product attribute for the simple reason that the partial effect of increases in the dominant firm's marginal cost is well understood and always leads to reduced output. By holding production costs fixed and isolating the effect of increased quality on output, we pinpoint the existence of a solely demand-induced reduction in output even when the product in question improves. We now turn to the derivation of the market equilibrium and the comparative statics.

2.2 The Effect of Higher Dominant Firm Quality on Output

Denote the dominant firm's price by P_1. Given that the price of the competitive fringe is equal to the marginal cost c_0, the critical value θ_c such that the corresponding consumer is indifferent between purchasing from the dominant firm and purchasing from the competitive fringe is

$$\begin{aligned} V_i + \theta_c g(x_0) - c_0 &= V_i + \theta_c g(x_1) - P_1 \Rightarrow \\ \theta_c(g(x_1) - g(x_0)) &= P_1 - c_0 \Rightarrow \\ \theta_c &= \frac{P_1 - c_0}{g(x_1) - g(x_0)}. \end{aligned} \tag{2}$$

Note that the value of V_i does not affect which variant of the product is chosen by consumers as long as the consumers for whom $\theta_i = \theta_c$ strictly prefer purchasing either of the two variants to purchasing neither for any value of V_i . A sufficient condition for this to be true is:

Assumption 1. The solution P_1^* to the dominant-firm's profit maximization problem satisfies the condition $\frac{P_1^* - c_0}{g(x_1) - g(x_0)} > \frac{c_0 - V_{MIN}}{g(x_0)}$.

Note that if the inequality in Assumption 1 is satisfied for $V_i = V_{MIN}$, it is also satisfied for all higher values of V_i. Also note that this condition implies that the fringe has a positive market share for consumers of every value of V_i. Assuming this condition is met, the demand function of the dominant firm is equal to $Q = 1 - F((P_1 - c_0)/(g(x_1) - g(x_0)))$, and the dominant firm chooses P_1 to maximize

$\pi = (P_1 - c_1)[1 - F(\frac{P_1 - c_0}{g(x_1) - g(x_0)})]$.[7] Rather than solve this maximization problem, we find that it provides more insight to recast the problem as one of optimal choice of output. The two approaches are equivalent since the dominant firm is the only strategic player and there is a one-to-one mapping between its price and the quantity it sells (a brute force proof of our main result that is based on first-order conditions of profit maximization with respect to price was used in earlier versions of the paper and this approach is used in the proof of Proposition 2 below). Solving the (residual) demand function of the dominant firm for P_1 yields the inverse demand function

$$P_1 = c_0 + (g(x_1) - g(x_0))F^{-1}(1 - Q). \tag{3}$$

Note that the demand intercept is $c_0 + (g(x_1) - g(x_0))\theta_{MAX}$ and is increasing in x_1. We assume that the MR function associated with this demand function is differentiable and monotonically decreasing, i.e., that $F^{-1}(1 - Q) + Q\frac{dF^{-1}(1-Q)}{dQ}$ is monotonically decreasing in Q. An increase in x_1 causes the residual inverse demand curve to pivot about some point. As the result below shows, this pivoting of the inverse demand curve causes the MR curve to rotate, because the demand intercept increases and the demand slope gets steeper.[8] Moreover, the height of the rotation point of the MR curve is the fringe marginal cost c_0.

Lemma 1 *An increase in the product quality of the dominant firm, x_1, causes a rotation of the marginal revenue curve, with the height of the rotation point equal to c_0. Marginal revenue is increasing for output levels to the left of the rotation point and decreasing for output levels to the right.*

Proof. Multiplying the RHS of (3) by Q and differentiating, we obtain marginal revenue

$$MR(Q) = c_0 + [g(x_1) - g(x_0)] \left[F^{-1}(1 - Q) + Q\frac{dF^{-1}(1 - Q)}{dQ}\right]. \tag{4}$$

---

[7] The equilibrium can in fact entail co-existence of a low cost dominant firm with a high cost fringe, as assumed in the analysis. For example, if $g(\cdot)$ is the identity function, $F(\cdot)$ is uniform on the $[0, 1]$ interval and $H(\cdot)$ is degenerate with $V_i = 0$, then profit maximization by the dominant firm entails $P_1 = (x_1 - x_0 + c_0 + c_1)/2$ and $\theta_c = (x_1 - x_0 + c_1 - c_0)/2(x_1 - x_0)$. The value of θ at which a consumer is indifferent between purchasing from the fringe and not purchasing at all is $\theta_{c'} = c_0/x_0$. If $x_1 = 2$ and $x_0 = 1$, then $\theta_c > \theta_{c'}$ (and thus the fringe has positive sales) as long as $c_0 < (1 + c_1)/3$, i.e., if the cost disadvantage of the fringe is not too high.

[8] We use the term "pivot" to describe the effect on demand, and the term "rotate" to describe the effect on marginal revenue. Though both terms describe a rotational movement, the former refers to a movement in which all price-quantity points move outwards in a clockwise direction (or stay fixed), while the latter refers to a movement in which some price-quantity points move outwards while others move inwards in a clockwise direction (with one point remaining fixed).

Note that

$$\frac{dMR(Q)}{dx_1} = g'(x_1)\left[F^{-1}(1-Q) + Q\frac{dF^{-1}(1-Q)}{dQ}\right]. \tag{5}$$

Evaluating at $Q = 0$, we obtain $\frac{dMR(0)}{dx_1} = g'(x_1)\theta_{MAX} > 0$, i.e., MR is increasing in x_1 for sufficiently low output levels. Substituting (5) back into (4) gives

$$MR(Q) = c_0 + \frac{g(x_1) - g(x_0)}{g'(x_1)}\frac{dMR(Q)}{dx_1}. \tag{6}$$

Since the quantity at which MR rotates must satisfy $dMR(Q)/dx_1 = 0$, we see that the height of the point about which MR rotates is equal to c_0. Given that the MR curve is assumed to be downward sloping, given that there is a one-to-one relationship between MR and dMR/dx_1 (from equation 6), and given that MR is increasing in x_1 for output $Q = 0$, the MR rotation implies that MR is constant in x_1 for the output level that corresponds to $MR = c_0$, is increasing in x_1 for lower values of Q, and is increasing in x_1 for higher values of Q. □

We now turn to the main question of interest. How does the dominant firm's quantity depend on the quality of its product? One might expect that it would go up. This prediction arises from models with horizontal product differentiation and consumers who value quality equally (e.g., Deltas, Harrington and Khanna, 2010). However, in our purely vertical framework, this is not the case if the dominant firm's marginal cost is at or below that of the fringe firms, as our main result below states.

Proposition 1 *Holding costs constant, the equilibrium quantity of the dominant firm is decreasing in its product quality x_1 when $c_1 < c_0$, is invariant to x_1 when $c_1 = c_0$, and is increasing in x_1 otherwise.*

Proof. The dominant firm's profit maximizing quantity is determined by the intersection of MR and c_1 (the firm's marginal cost). Denote by Q_0^* the optimal output level if $c_1 = c_0$. Given that MR is downward sloping, both Q_0^* and the intersection of MR with c_1 define unique output levels. Lemma 1 shows that an increase in x_1 increases MR for all quantity levels lower than Q_0^*, decreases them for all quantity levels higher than Q_0^*, and leaves them unchanged for $Q = Q_0^*$. The result is obtained by observing that the optimal output Q^* is less than Q_0^* if $c_1 > c_0$, is equal to Q_0^* if $c_1 = c_0$, and is greater than Q_0^* if $c_1 < c_0$. □

A simple way to see the intuition behind our main result is as follows. Suppose for the moment that $V_{MIN} > c_0$ so that all consumers buy some version of the product and that $\theta_{MIN} = 0$ so that there is some consumer who does not value quality at all. In this case the residual inverse demand faced by the dominant firm is determined by how much consumers value a product of quality x_1 when the alternative is to buy (from the fringe) a product of quality x_0 at a price c_0. Any consumer for whom $\theta > 0$ will have a valuation for the dominant firm's product higher than c_0, but a consumer for whom $\theta = 0$ will have a valuation equal to c_0. This consumer regards both products as equally good, and so is willing to pay c_0 for the dominant firm's product when the alternative is to buy from the fringe at c_0. An increase in the dominant firm's quality from x_1 to x_1' causes the residual inverse demand curve faced by the dominant firm to pivot, not to shift parallel, because the increase in each consumer's willingness-to-pay depends on how much they value quality. Defining $\bar{Q}$ as the quantity corresponding to a consumer for whom $\theta = 0$, the increase in quality causes the dominant firm's inverse demand curve to pivot about the point $(\bar{Q}, c_0)$.[9] This is depicted in Figure 1, in which the distribution of θ is uniform. Lemma 1 above shows that the height of the rotation point of the marginal revenue curve is also c_0, which is indicated in Figure 1 and leads directly to Proposition 1.

Now we relax the assumption that everyone buys some version of the product and allow for the possibility that consumers with sufficiently low θ and/or V_i do not buy at all. Now there are two notional inverse demand curves that the dominant firm might face: one where consumers' preferred alternative is to buy from the fringe at a price c_0, and one where the alternative is not to buy at all. A quality increase causes the latter inverse demand curve to pivot (and its MR curve to rotate), but about a point whose height is other than c_0. But as long as Assumption 1 is satisfied, the relevant inverse demand curve for the dominant firm is the former one, and so allowing the possibility that consumers buy nothing has no effect on its conduct. The only change is that now some consumers buy nothing instead of buying from the fringe.

3 Consumer Surplus and Total Welfare

We consider the welfare effects of an increase in the quality of the dominant firm's product from x_1 to x_1' when $c_1 \leq c_0$, with associated equilibrium prices of P_1 and P_1', starting with evaluation of the consumer surplus (the less interesting case of $c_1 > c_0$ can be analyzed in a similar manner). Since each consumer has three possible choices (buy nothing, buy from the fringe, buy from the dominant firm) both before and after the quality increase, there are nine choice pair possibilities. Given our

[9] Note that if zero were not in the support of θ, $\bar{Q}$ would be obtained from a demand that would result from hypothetically assuming the existence of consumers with $\theta = 0$ and extrapolating the demand to that value of θ.

assumptions, five of these nine can be ruled out.[10]

The remaining four possible types are illustrated in Figure 2. First are consumers with values of θ_i low enough that $V_i+\theta_i g(x_0)-c_0<0$, i.e., $\theta_i \in (\theta_{MIN}, \frac{c_0-V_i}{g(x_0)})$. These consumers do not buy any version of the product either before or after the quality increase, and so their welfare is unchanged. Second are consumers with values of θ_i such that $\theta_i \in (\frac{c_0-V_i}{g(x_0)}, \frac{P_1-c_0}{g(x_1)-g(x_0)})$. These consumers buy the product from the fringe both before and after the increase in the quality of the dominant firm's product. Since the fringe's price and quality are unaffected by the dominant firm's quality increase, the welfare of these consumers is unaffected as well. Third are consumers with values of $\theta_i \in (\frac{P_1-c_0}{g(x_1)-g(x_0)}, \frac{P_1'-c_0}{g(x_1')-g(x_0)})$. These consumers buy from the dominant firm before the quality increase, but switch to the fringe after it. Since the utility of consuming the fringe's product did not change, by revealed preference these consumers must be worse off following the quality improvement. Fourth are consumers for whom $\theta_i > \frac{P_1'-c_0}{g(x_1')-g(x_0)}$. These consumers buy the high-quality product both before and after the quality improvement. They can be divided into two sub-types. The first sub-type consists of consumers with values of θ_i such that $\frac{P_1-c_0}{g(x_1)-g(x_0)} < \theta_i < \tilde{\theta}_c$ (where $\tilde{\theta}_c$ satisfies $\tilde{\theta}_c g(x_1) - P_1 = \tilde{\theta}_c g(x_1') - P_1'$ and represents the consumer who values quality just enough to be equally well off before and after the quality increase), i.e., $\theta_i \in (\frac{P_1'-c_0}{g(x_1')-g(x_0)}, \frac{P_1'-P_1}{g(x_1')-g(x_1)})$. These consumers suffer a reduction in their utility as they value the quality improvement less than they dislike the increase in the product price. The second sub-type consists of individuals for whom $\theta_i > \tilde{\theta}_c$, i.e., $\theta_i \in (\frac{P_1'-P_1}{g(x_1')-g(x_1)}, \theta_{MAX})$. These consumers are made better off by the quality improvement.[11]

The effect of the quality increase on total consumer surplus is ambiguous: the gains to consumers with the highest values of θ_i may be larger or smaller than the losses to those with lower values.[12]

---

[10]There are no consumers who don't buy at all before the quality increase and buy from the fringe after, or the reverse, because the increase in the quality of the dominant firm's product does not affect the utility of either of these choices. There are also no consumers who don't buy at all before and buy from the dominant firm after, or who buy from the fringe before and from the dominant firm after, because our main result shows that the quality increase causes quantity to (weakly) decrease. Finally, there are no consumers who buy from the dominant firm before and don't buy at all after, because under Assumption 1 those who purchase from the dominant firm prefer the fringe's product to not buying at all.

[11]Our assumptions, including the assumption that $c_1 \leq c_0$, ensure that the thresholds are ranked as in Figure 2.

[12]For example, assume θ has a density of 1 on $[0,u]$ where $u \leq 1$ and there is a mass point with mass 1-u at $\theta = u$. One can show that in this case the effects on consumer surplus can be of either sign depending on u. For u low enough, no consumer is better off from the quality increase. Note that if $u = 1$ the distribution becomes U[0,1], in which case total consumer surplus unambiguously increases.

The effect of the quality increase on total consumer surplus will depend on whether $F(\cdot)$ has a fat or a thin tail above $\tilde{\theta}_c$. Note that the set of consumers for whom $\theta_i > \tilde{\theta}_c$ may be empty. This is because P_1' does not depend on the support or the shape of $F(\cdot)$ above $\frac{P_1'-c_0}{g(x_1')-g(x_0)}$, and so θ_{MAX} could be bigger or smaller than $\tilde{\theta}_c$. If $\tilde{\theta}_c > \theta_{MAX}$, then all consumers are weakly worse off by the quality increase.

To summarize, the only consumers who can possibly be affected by the increase in the dominant firm's quality are those who were willing to buy the dominant firm's initial (lower quality) product. Among those, the consumers with relatively low willingness-to-pay for quality are worse off (the increase in quality is not sufficient to compensate for the increase in price) while the consumers with sufficiently high willingness-to-pay for quality are better off. But if the distribution $F(\cdot)$ is relatively tight, there may not be any consumers in that latter set, or even if there are, aggregate consumer surplus will decline.

The quality increase will have no effect on fringe producers, as they earn zero profits both before and after. It will increase the profits of the dominant firm, since it could have kept the price unchanged and earned the pre-improvement profit, but chose to raise it instead. The effect on total welfare is ambiguous.

At this point, it is instructive to briefly contrast these results with the effects of an exogenous and costless improvement in the quality of the fringe firms. Such a quality increase would increase consumer surplus of every consumer, because (i) it would increase the value of the fringe product without increasing its price, since that price is equal to marginal cost, and (ii) it would decrease the price of the dominant firm's product, since its residual demand would shift in and also become flatter, without decreasing its value.[13] Thus, both options that consumers face result in higher surplus to them.[14]

[13]The quality improvement of the fringe's product from x_0 to x_0' increases the willingness of consumers to pay for it by $\theta_i[g(x_0') - g(x_0)]$. The residual demand of the dominant firm is shifted downwards by this same amount, and thus becomes flatter since high θ consumers are those with the highest willingness to pay for the dominant firm's product.

[14]These observations are not sufficient to sign the welfare effects of a fringe quality improvement. Consumers who switch from the dominant firm's product to the fringe's product gain in consumer surplus, but their switch to the fringe reduces the dominant's firms profits (recall that fringe firm's have zero profit since they price at marginal cost). Indeed, the marginal consumer who switches to the fringe reduces welfare that consumer's welfare will increase by a very small amount, but the fringe firm's profits from selling to that consumer falls by a discrete amount.

4 Discussion and Extensions

The most natural extension to our model would be to allow all firms to be strategic, rather than assuming a non-strategic competitive fringe. We did not pursue this extension because a small amount of strategic interaction (supported, perhaps, by a small amount of differentiation among the fringe firms) will not materially affect our results. In what follows, we take up other more meaningful extensions.

4.1 Markets without the Competitive Fringe: The Monopoly Case

Suppose the fringe was completely absent and the dominant firm was a pure monopolist. Further suppose that $V = 0$, as in standard vertical differentiation models. Would a similar result obtain? In this case, the pivot point of the demand curve and the rotation point of the marginal revenue curve will both have a height of zero. Clearly Proposition 1 will not hold, as the height of the rotation point is below c_1. Under standard vertical differentiation models, then, the competitive fringe is essential for obtaining our results.

But if we return to our more general preference specification, then we can obtain a result like the one in Proposition 1 without a fringe, but only under more restrictive conditions. A sufficient condition is that $V_{MIN} > c_1$. To most easily demonstrate this, suppose that $V_i = V > c_1$ for all i. In this case, the value of θ for the consumer who is indifferent between purchasing and not purchasing the good is given by

$$\theta_c = \frac{P_1 - V}{g(x_1)}. \tag{7}$$

The monopolist's profit function is $\pi = (P_1 - c_1)[1 - F(\frac{P_1 - V}{g(x_1)})]$ and the first order condition of profit maximization with respect to P_1 is given by

$$\frac{\partial \pi}{\partial P_1} = \left[1 - F\left(\frac{P_1 - V}{g(x_1)}\right)\right] - (P_1 - c_1) f\left(\frac{P_1 - V}{g(x_1)}\right)\frac{1}{g(x_1)} = 0, \tag{8}$$

which in turn can be re-written as

$$[1 - F(\theta_c^*)] - \frac{P_1^* - c_1}{P_1^* - V} f(\theta_c^*)\theta_c^* = 0. \tag{9}$$

In this environment, increases in product quality lead to decreases in quantity, as Proposition 2 states.

Proposition 2 *Suppose there is no competitive fringe. Then an increase in x_1, holding c_1 constant, leads to a reduction in the monopolist's sales if and only if $V > c_1$.*

Proof. Since by assumption $V > c_1$, the ratio $(P_1^* - c_1)/(P_1^* - V)$ is decreasing in P_1^*. Consider an increase in x_1 accompanied by an increase in P_1 such that θ_c remains unchanged. Then, the left hand side of equation (9) would be positive. A positive value of the left hand side of (9) implies that the firm's profit would increase if it further raised its price. Thus, an increase in P_1 that leads to no change in the monopolist's sales is smaller than the profit maximizing increase. Therefore, the profit maximizing price increase would reduce the firm's sales. The proof of the converse follows by reversing the signs in the above exposition.[15] $\square$

This example fits the Johnson and Myatt (2006) framework of demand rotation, albeit as a boundary case, with the rotation at the edge of the support of consumer willingness-to-pay. However, the model with the competitive fringe, which is outside the Johnson and Myatt framework, "works" without the need to assume that the product has a high value even in the complete absence of the salient attribute.

We now show that the model with the competitive fringe and the monopoly model can fit in a unified framework. Since the competitive fringe is non-strategic, the dominant firm can be thought of as a monopolist, albeit one facing the residual demand rather than market demand. The net willingness-to-pay for the dominant firm's product for consumers whose best alternative is purchasing from the fringe is

$$\begin{aligned} W_i &= U_i^{dominant} - U_i^{fringe} \\ &= (V_i + \theta_i g(x_1) - P_1) - (V_i + \theta_i g(x_0) - c_0) \\ &= c_0 + \theta_i (g(x_1) - g(x_0)) - P_1. \end{aligned} \tag{10}$$

Since $g(\cdot)$ can be any increasing function, we can re-parametrize it as $\gamma(x_1) = g(x_1) - g(x_0)$ and write

$$W_i = c_0 + \theta_i \gamma(x_1) - P_1. \tag{11}$$

Note that equation (11) is of the same form as equation (1) but with V_i being replaced by the marginal cost of the fringe. The price of the outside option c_0 takes the place of the value of the good in the absence of the attribute V_i. Since we have shown in the preceding section that the only factor that determines whether the monopoly output will decline with an increase in quality is the relationship

[15] This proposition could have been proven using steps analogous to those in Proposition 1, but where the rotation point of the MR curve is at a height of V rather than c_0. But we find it useful to show how both results can be proved in a more "brute force" approach based on profit maximization with respect to price.

between V and the marginal cost of the monopolist, and not the shape of the quality function $g(\cdot)$, it follows that in the presence of the competitive fringe the only relevant factor is the comparison between the marginal cost of the fringe and that of the dominant firm.

4.2 Multi-product Firms and Cost Changes

Our stylized model makes two assumptions regarding the environment following the introduction of the new high-quality product. The first is that the old high-quality product is discontinued upon introduction of the new one. The analysis in Itoh (1983) is directly relevant to what happens if this is not the case.[16] If the dominant firm retains both products, then following Itoh's Proposition 1, the optimal price of the original high-quality product remains unchanged, and so the market share of the dominant firm also remains unchanged. Consumer surplus goes up, as consumers either consume the product they used to and pay the same price, or they consume a better product at a higher price, which by revealed preference makes them better off. Welfare also goes up, since both consumer surplus and profits go up as long as all products have positive market share, as ensured by Assumption 1.

It is worthwhile noting that in many cases the introduction of a new product (e.g., the iPad or other electronics) is accompanied by the discontinuation of the older product, as we assume in the main body of the paper. An explanation for this is the presence of substantial fixed costs at the product level. The presence of such costs would make it unprofitable to manufacture, market, and distribute multiple versions of the same product, making the single-product case the salient case. Evans and Salinger (2005, 2008) present empirical evidence of the importance of fixed costs at the product level and develop a theoretical model of the relevance of such fixed costs in evaluating tying and bundling conduct. Moreover, in the non-temporal interpretation of our model, comparing a world with a dominant firm's sole product of a particular quality versus a world where quality is even higher, it is not meaningful to consider the co-existence of both products. Our second assumption is that costs are the same for both versions of the high-quality product. If instead the higher-quality version has higher costs, then our results become stronger: the price of the new product is increasing in the production cost; hence, the dominant firm's market share, consumer surplus and total welfare will all decrease.

[16] The competitive fringe in our model is equivalent to the outside option in Itoh, since no consumer is indifferent between purchasing from the dominant firm and not purchasing at all, and since the dominant firm does not have an 100% market share for consumers with any value of V_i.

A larger departure from our simple framework involves a simultaneous change in quality of both the dominant firm and the fringe. For example, following the introduction of the new product by the dominant firm, the old product could become generic and be produced by the fringe at its old marginal cost. The effects of this depend on the relative magnitudes of the differences $g(x_1) - g(x_0)$ and $g(x_1') - g(x_1)$. If these two differences are the same, then there is no change in the dominant firm's demand (see equation (3)), and hence in its price and market share. This is not surprising since the dominant firm has a better product, but not better relative to the new product of the fringe. Consumer surplus goes up, however, since consumers will purchase uniformly better products at the old prices. If the second difference is larger than the first, then our "unconventional" results continue to hold with regard to quantities, but not with regard to consumer surplus, since products will be uniformly weakly better for consumers (even after allowing for higher prices). If the second difference is smaller than the first, then our results do not hold even for quantities. However, a seeming paradox will remain: even though the quality gap between the dominant firm and the fringe gets smaller, the dominant firm's market share nevertheless goes up.

4.3 The Limits of this Framework

The results outlined so far depend upon the standard (and reasonable) assumption in vertical differentiation models that willingness-to-pay for the product is a linear function of a monotonic transformation $g(\cdot)$ of the product attribute. We now consider a modification of the model that departs from this linear assumption by allowing utility to be quadratic in the attribute

$$U_{ij} = V_i + \theta_i x_j + \delta x_j^2 - P_j, \tag{12}$$

where $g(\cdot)$ has been replaced by the identity function and δ is a parameter common to all consumers. A possible example in which willingness-to-pay could be quadratic is one where product quality takes the form of travel time to reach the firm's location. A higher value of x_j means lower travel time, possibly because of lower traffic on the road leading to the firm (and θ_i and δ would both be negative). Now $\theta_i x_j$ would represent the monetary loss from forgoing paid work to travel to the firm, or the monetary cost of the fuel required to drive there. Since wages and fuel efficiency vary across potential consumers and their vehicles, the parameter θ_i could differ across individuals. The term δx_j^2 would represent the pure disutility of travel and would be convex in the required travel time. Though this could also differ across consumers, for simplicity it is assumed not to.

Under these preferences, the value of θ_i for the consumer who is indifferent between purchasing from the dominant firm and purchasing from the fringe is given by

$$\begin{aligned}
\theta_c x_0 + \delta x_0^2 - c_0 &= \theta_c x_1 + \delta x_1^2 - P_1 \Rightarrow \\
\theta_c (x_1 - x_0) &= P_1 - c_0 - \delta(x_1^2 - x_0^2) \Rightarrow \\
\theta_c &= \frac{P_1 - c_0}{x_1 - x_0} - \delta(x_1 + x_0). \qquad (13)
\end{aligned}$$

For simplicity from now on we assume that θ_i is distributed U$[0,1]$. The residual demand for the dominant firm's product is

$$Q = 1 - \frac{P_1 - c_0}{x_1 - x_0} + \delta(x_1 + x_0), \qquad (14)$$

which solving for price yields

$$P_1 = c_0 + (x_1 - x_0) + \delta(x_1^2 - x_0^2) - (x_1 - x_0)Q. \qquad (15)$$

Notice the pivoting component, which pushes the y-axis intercept upwards and decreases the slope one-for-one for an increase in the value of x_1, and an additional parallel shift component that pushes out the demand at a rate of $2\delta x_1$ as x_1 increases. The optimal output level of the dominant firm is obtained by equating the MR with marginal cost

$$Q^* = \frac{1}{2} + \frac{1}{2}\frac{c_0 - c_1}{x_1 - x_0} + \frac{\delta}{2}(x_1 + x_0). \qquad (16)$$

This preference structure not only leads the demand curve to pivot in response to a quality increase, but also causes it to translate outward. Notice that $\delta = 0$ yields our original model and Q^* is invariant to quality when $c_1 = c_0$ and decreasing in x_1 when $c_1 < c_0$. But if $\delta > 0$ then there is an additional, strictly positive term that does not depend on the relationship between the two marginal costs. This means that Q^* unambiguously increases when $c_1 = c_0$, and can increase or decrease when $c_1 < c_0$ depending on which effect dominates.

5 Concluding Remarks

In this paper we identify a mechanism through which quality improvements that lead to increased willingness-to-pay for a product lead to reduced equilibrium output of that product. The result does not come from any trivial source such as reduced long run sales of a product whose durability increases. Rather, it comes from the fact that higher product quality of a dominant firm that is facing a competitive fringe greatly reduces the elasticity of the dominant firm's residual demand curve when its marginal cost is lower than that of the fringe. This elasticity effect leads to such a large increase in the profit maximizing price that the dominant firm sells fewer units. We also show that the effect of a quality increase on consumer surplus (and on total surplus) is ambiguous, but that it is possible for *all* consumers to be made weakly worse off, with some being strictly worse off.

A number of markets can (to a first approximation) be described as consisting of a dominant firm competing against a number of much smaller and less efficient rivals, and the standard vertical differentiation model on which we rely is a reasonable approximation of consumer preferences for products that are differentiated by quality, so our model is likely to have reasonably broad applicability. And even in situations where other quantity-increasing effects dominate the quantity-reducing effect analyzed here, its presence will tend to make the quantity increase smaller than it otherwise would be. At the very least we have shown that a quality improvement in the product of a dominant firm facing a competitive fringe has an effect of indeterminate sign on that firm's output, and that in an important special case, it is guaranteed to have a negative effect.

References

Beath, J. and Y. Katsoulacos, *The Economic Theory of Product Differentiation*, Cambridge University Press, Cambridge UK, 1991.

Choi, C. J. and H.S. Shin [1992] "A Comment on a Model of Vertical Product Differentiation," *Journal of Industrial Economics*, vol. 40, pages 229-231.

Deltas, G., D. Harrington, and M. Khanna [2012] "Oligopolies with (Somewhat) Environmentally Conscious Consumers: Market Equilibrium and Regulatory Intervention," *Journal of Economics and Management Strategy*, forthcoming.

Evans, D. S., and M. A. Salinger [2005] "Why Do Firms Bundle and Tie? Evidence from Competitive Markets and Implications for Tying Law, *Yale Journal on Regulation*, vol. 22, pages 37-89.

Evans, D. S., and M. A. Salinger [2008] "The Role of Cost in Determining When Firms Offer Bundles and Ties, *Journal of Industrial Economics*, vol. 56, pages 143-168.

Frascatore, M. R. [1999] "Vertical Product Differentiation when Quality is Scarce: the Case of $n > 2$ Firms," *Australian Economic Papers*, vol. 38, pages 120-130.

Gabszewics, J. J., A. Shaked, and J. Sutton [1986] "Segmenting the Market: The Monopolist's Optimal Product Mix," *Journal of Economic Theory*, vol. 39, pages 273-289.

Gabszewicz, J. J. and J.-F. Thisse [1979] "Price Competition, Quality and Income Disparities," *Journal of Economic Theory*, vol. 20, pages 340-359.

Itoh, M. [1983] "Monopoly, Product Differentiation and Economic Welfare," *Journal of Economic Theory*, vol. 31, pages 88-104.

Johnson, J.P. and D. P. Myatt [2006] "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economics Review*, vol. 96, pages 756-784.

Lehmann-Grube U. [1997] "Strategic Choice of Quality when Quality is Costly: the Persistence of the High-Quality Advantage," *Rand Journal of Economics*, vol. 28, pages 372-384.

Motta, M. [1993] "Endogenous Quality Choice: Price vs. Quantity Competition," *Journal of Industrial Economics*, vol. 41, pages 113-131.

Noh, Y. H. and G. Moschini [2006] "Vertical Product Differentiation, Entry-Deterrence Strategies, and Entry Qualities," *Review of Industrial Organization*, vol. 29, pages 227-252.

Shaked, A. and J. Sutton [1983] "Natural Oligopolies," *Econometrica*, vol. 51, pages 1469-1483.

Figure 1: Demand Pivots and Marginal Revenue Rotates

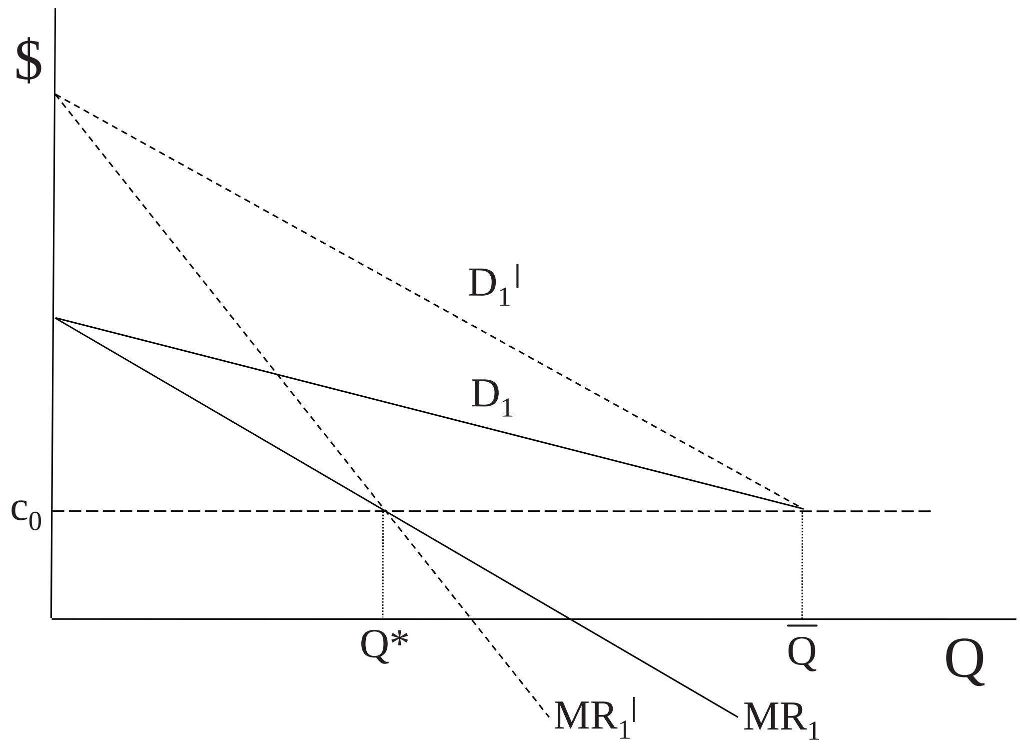

The solid lines represent the dominant firm's demand and marginal revenue at its original quality, and the dashed lines at the higher quality. Note that demand pivots and that marginal revenue rotates about points of height c_0.

Figure 2: Consumer Surplus Effects of an Increase in the Dominant Firm's Quality