On the Simple Economics of Advertising, Marketing, and Product Design

By Justin P. Johnson and David P. Myatt*

We propose a framework for analyzing transformations of demand. Such transformations frequently stem from changes in the dispersion of consumers' valuations, which lead to rotations of the demand curve. In many settings, profits are a U-shaped function of dispersion. High dispersion is complemented by niche production, and low dispersion is complemented by mass-market supply. We investigate numerous applications, including product design; advertising, marketing and sales advice; and the construction of quality-differentiated product lines. We also suggest a new taxonomy of advertising, distinguishing between hype, which shifts demand, and real information, which rotates demand. (*JEL* D8, L1, M3).

We propose a framework for analyzing consumer demand that is broadly applicable and yet easily understood in terms of simple economic concepts. It unifies analysis of seemingly disparate economic phenomena, such as the effects of advertising choices, product design, and income inequality, by revealing that they are fundamentally identical in their impact on the shape of demand. A consequence is that our analysis leads to numerous predictions concerning the changing shape of demand and firms' behavior and profitability.

The foundation of our framework is the observation that many forces influence the dispersion of consumers' valuations, leading to a rotation, as opposed to a shift, of the demand curve. It follows that understanding the simple economics of demand rotations helps to explain many phenomena. Surprisingly, however, while demand rotation is an elementary concept, it has received remarkably little formal study. We therefore provide some general results concerning changes in demand dispersion (equivalently, demand rotation). We show that firms have preferences for extremes, favoring either very high or very low levels of dispersion. Thus, a firm's overall mix of activities, including advertising and product design, will be chosen either to maximize or minimize dispersion. These choices are, in turn, tied to whether a firm pursues a niche (low-volume) as opposed to a mass-market (high-volume) position. Before explaining further, we motivate our study with some details of specific applications.

First, consider an alteration to a product's design. If consumers unanimously prefer the new version, then the simple consequence is that the demand curve shifts outward, so that for any given quantity, revenue must increase. Other situations are less straightforward, as a design change may appeal to some consumers while displeasing others. Often this will change the dispersion of demand (that is, the dispersion of the distribution of consumers' willingness to pay), leading not to a shift but rather to a rotation of the demand curve.

As further motivation, consider advertising and marketing activities. If product promotion is unambiguously persuasive, or informs consumers of a product's existence, then it will shift the demand curve outward. As with product design, however, the effects may be more complex. This will be so when advertising provides presales information that enables consumers to ascertain better their true underlying

* Johnson: Johnson Graduate School of Management, Cornell University, Ithaca, NY 14853 (e-mail: jpj25@cornell.edu); Myatt: Department of Economics, University of Oxford, Oxford OX1 3UQ, United Kingdom (e-mail: david.myatt@economics.ox.ac.uk). We thank colleagues, especially Ian Jewitt, Paul Klemperer, Meg Meyer, Vithala Rao, Chris Wallace, and seminar participants at the University of Bristol, Duke University, LSE, MIT, Northwestern University, Oxford University, University of Southampton, UCL, Wake Forest University, the UBC 2004 Summer IO Conference, and the 2005 International Industrial Organization Conference, as well as two referees, for helpful comments.

idiosyncratic preferences for the product: it may discourage some customers from purchasing while encouraging others. The consequence is a change in the dispersion of valuations, and a rotation of the demand curve.

Of course, some activities may involve both a shift in demand and a change in dispersion. For instance, a marketing campaign may hype a product's existence while giving details of the product's style and function. Our approach encompasses such situations, since we show that the combination of effects frequently reduces to a rotation of the demand curve.

While the analysis of demand rotations is somewhat more subtle than that involving mere shifts in demand, the results, nonetheless, turn out to be fairly straightforward. Our core results are easiest to see from the perspective of a monopolist. When consumers' valuations for a product are relatively homogeneous, a firm typically will choose to serve a large fraction, or "mass market," of potential consumers. Heuristically, the marginal consumer is "below average" in the distribution. Following an increase in dispersion, the demand curve rotates clockwise. This will push the willingness to pay of the marginal consumer down, and profits will fall. The firm will wish to minimize dispersion; furthermore, any reduction in dispersion will tend to enhance the desire to retain a mass-market posture. On the other hand, when consumers are heterogeneous, the firm will restrict sales to a relatively small "niche" of potential consumers. The marginal consumer will then be "above average" and response positively to increases in dispersion; profits increase as the demand curve rotates clockwise. Such an increase in dispersion will tend to reinforce the desire to retain a niche posture. Building upon this intuition, we show that, in a wide variety of circumstances, profits are "U-shaped" (that is, quasi-convex) in the dispersion of demand. Consequently, a firm will wish to use any tools at its disposal to pursue either maximal or minimal dispersion. We identify, therefore, two distinct marketing mixes that a firm may wish to deploy.

First, when a firm stakes out a mass-market position, its goal is to minimize the dispersion of demand. In pursuit of this goal, the product will be designed to have universal appeal, with controversial design features being eschewed; it will offer something for everyone. Any advertising will highlight the existence of the product, but will not allow consumers to learn precisely their true match with the product's characteristics. Such a marketing mix will ensure that consumers share similar valuations for the product, so that we will see a "plain-vanilla" product design promoted by an advertising campaign consisting of "pure hype."

Second, when adopting a niche posture, a firm aims to maximize the dispersion of demand. The product design will have extreme characteristics that appeal to specialized tastes. Many consumers will strongly dislike the product, but those who like it will love it. This product design will be accompanied by advertising and detailed sales advice containing "real information" that allows consumers to learn of their true match with the product's "love-it-or-hate-it" attributes.

A number of applications flow from our basic observation that many factors influence the shape of demand. For instance, building upon the intuition above, we introduce a new taxonomy of advertising, distinguishing between hype and real information. Promotional hype corresponds to the traditional notions of informative and persuasive advertising. It highlights the existence of the product, promotes any feature that is unambiguously valuable, or otherwise increases the willingness to pay of all consumers; it shifts the demand curve outward. Real information, on the other hand, allows a consumer to learn of his personal match with the product's characteristics; as we show, it rotates the demand curve. Importantly, supplying real information may sometimes lower, rather than raise, sales (and profits). Our taxonomy incorporates this possibility, distinguishing it from nearly all previous work on advertising.

Our results are not artifacts of an assumption that a firm sells but a single product. We extend our analysis to multiproduct firms offering product lines of vertically differentiated goods. In this setting, a consumer's type corresponds to a preference for increased quality, and we determine when profits are U-shaped in the dispersion of such types. We also relate the length and mix of a product line to consumer-type dispersion. When types are more dispersed, a firm frequently serves a smaller overall share of the market, but does so with a longer product line.

While much of our exposition focuses on monopoly, and our key intuitions are most easily conveyed in that case, our framework and results can be extended to handle competition. Under perfect competition, our results continue to hold. In the case of Cournot oligopoly, important strategic interactions are engendered, and the analysis is more complicated. Nonetheless, our results on the relationship between consumer-type dispersion and the structure of product lines offered by multiproduct firms continue to hold. Furthermore, when additional structure is placed upon the shape of demand, an oligopolist's profits continue to be U-shaped in dispersion. We also study the effects of increased competition on the relationship between demand dispersion and firm profitability.

Our work is related to several fields of economic inquiry. We believe that fully conveying such relationships is best accomplished by deferring thorough review to the individual subsequent sections. A few notes, however, are in order here. In Section I, we investigate the response of profits and output to the dispersion of demand. Our notion of dispersion builds upon the classic work of Michael Rothschild and Joseph E. Stiglitz (1970, 1971) and the single-crossing property of distribution functions studied by Peter A. Diamond and Stiglitz (1974) and John S. Hammond (1974). In Section II, our application to product design is based upon Lancastrian characteristics (Kevin J. Lancaster, 1966, 1971). In Section III, we turn to advertising, sales advice, and other marketing activities. Our study of the incentives to equip consumers with private information exploits the insights of Tracy R. Lewis and David E. M. Sappington (1991, 1994), and complements Marco Ottaviani and Andrea Prat's (2001) work on public information supply. Section IV's consideration of vertically differentiated goods builds upon the classic model of Michael Mussa and Sherwin Rosen (1978), as well as work by Johnson and Myatt (2003, 2006).

I. Demand Dispersion and a Preference for Extremes

Here we provide a framework upon which we build in Sections II to IV. We consider the monopoly supply of a single product and investigate the relationship between the shape of demand and profits, quantity, and price. Other market structures are postponed until Section V.

A. *Consumer Demand*

There is a unit mass of consumers. A consumer is willing to pay up to θ for a single unit of a particular product, where θ is drawn from the distribution $F_s(\theta)$ with support on some interval $(\underline{\theta}_s, \bar{\theta}_s)$.[1] The parameter $s \in S = [s_L, s_H]$ indexes a family of distributions. For $\theta \in (\underline{\theta}_s, \bar{\theta}_s)$, we assume that $F_s(\theta)$ is twice continuously differentiable in θ and s, and write $f_s(\theta)$ for the strictly positive density.[2] Thus, given a price p, a fraction $z = 1 - F_s(p)$ of consumers will choose to buy the product. We will often work with the inverse demand curve for the product. Hence, if z units are to be sold, the market clearing price must satisfy $p = P_s(z) \equiv F_s^{-1}(1 - z)$.[3] Thus, $s \in S$ also indexes a family of inverse demand curves.

To ascertain the impact of changing demand, we impose structure on the family of distributions. One possibility would be to assume that $F_s(\theta)$ is decreasing in s for all θ, so that the family is ordered by first-order stochastic dominance (Erich L. Lehmann, 1955). The economic interpretation is that an increase in s results in an upward shift of the inverse demand curve.

Simple shifts in demand, however, are not our focus. We consider the *shape* of demand in particular situations in which an increase in s results in a distribution that is "more disperse" or "more heterogeneous."[4] We consider concrete examples of such scenarios in Sections II and III. For now, however, we require a

[1] In Section IV, we study a model of price discrimination over quality, which is formally equivalent to price discrimination over quantities when consumers have demand for multiple units.

[2] We also require the lower bound $\underline{\theta}_s$ and upper bound $\bar{\theta}_s$ of the support to be continuous functions of s.

[3] The inverse is well defined for $z \in (0, 1)$. We complete the specification by setting $P_s(0) = \bar{\theta}_s$ and $P_s(1) = \underline{\theta}_s$.

[4] Our distinction is motivated by Anthony B. Atkinson's (1970, p. 245) suggestion to "separate shifts in the distribution from changes in its shape and confine the term inequality to the latter aspect."

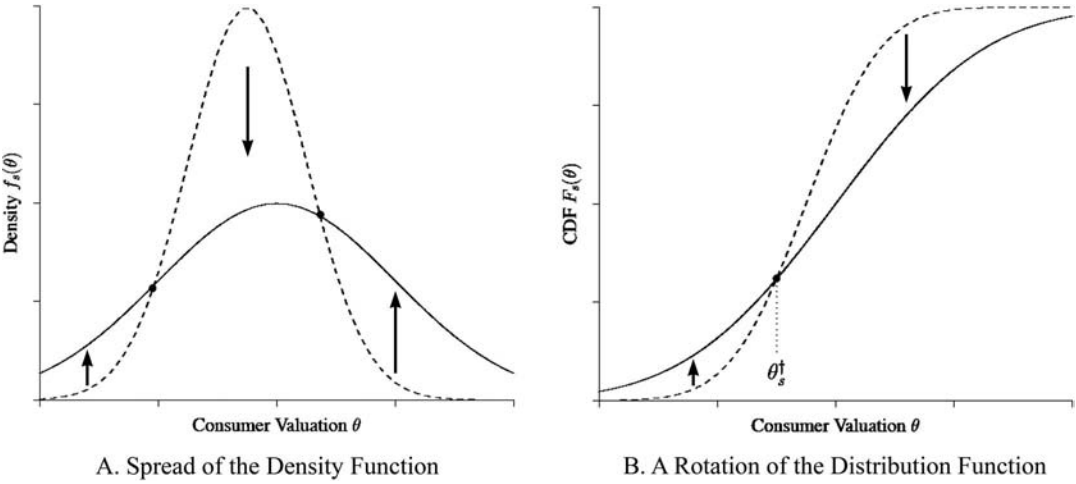

FIGURE 1. THE RESPONSE OF $f_s(\theta)$ AND $F_s(\theta)$ TO AN INCREASE IN DISPERSION

measure of "riskiness" to rank different distributions.[5]

An immediate candidate for such a measure is second-order stochastic dominance. Following the familiar Rothschild and Stiglitz (1970) definition, an increase in s results in increased risk if the expectation of any concave function decreases with s. Furthermore, a riskier distribution may be obtained via the addition of a sequence of mean-preserving spreads.

Although we do not use second-order dominance directly, we do build upon the concept of a (not necessarily mean-preserving) spread. Specifically, an increase in s results in a spread if it moves density away from the center of $f_s(\theta)$ and toward the upper and lower tails, as illustrated in Figure 1A.[6] Such a spread results in a clockwise rotation of the distribution function: there is some *rotation point* $\theta_s^\dagger$ such that $F_s(\theta)$ is increasing in s for $\theta < \theta_s^\dagger$, and decreasing in s for $\theta > \theta_s^\dagger$. In fact, a rotation is a weaker measure of increased dispersion than a spread.[7] It is central to our analysis, and hence we state it as a formal definition.

DEFINITION 1: *A local change in s leads to a* rotation *of $F_s(\theta)$ if, for some $\theta_s^\dagger$ and each $\theta \in (\underline{\theta}_s, \overline{\theta}_s)$,*

$$\theta \gtrless \theta_s^\dagger \Leftrightarrow \frac{\partial F_s(\theta)}{\partial s} \lesseqgtr 0.$$

If this holds for all s, then $\{F_s(\theta)\}$ is ordered by a sequence of rotations. *Equivalently,*

$$z \gtrless z_s^\dagger \Leftrightarrow \frac{\partial P_s(z)}{\partial s} \lesseqgtr 0$$

$$\text{where } z_s^\dagger = 1 - F_s(\theta_s^\dagger),$$

so that the family of inverse demand curves $\{P_s(z)\}$ is ordered by a sequence of rotations.

Definition 1 stipulates that two cumulative distribution functions or inverse demand functions, differing by a (clockwise) rotation, must cross only once.[8] Employing this notion, Dia-

[5] There is no uncertainty over the demand curve faced by a firm; our work is distinct from that of Hayne E. Leland (1972) and Donald V. Coes (1977), who considered a risk-averse firm faced by a stochastic demand curve.

[6] The notion of a spread does not require preservation of the mean. For instance, Diamond and Stiglitz (1974) considered preservation of $E[u(\theta)]$, and Ian Jewitt (1989) defined a notion of location-independent risk.

[7] For a rotation, the distribution functions cross once. Their densities, however, may cross an arbitrary number of times. If they differ by a spread, then their densities will cross only twice. Although we restrict our analysis to rotations throughout the main paper, our results (particularly Proposition 1) extend to measures of increased dispersion that involve multiple rotation points. We return to this issue later in the paper, and explain further in Appendix A.

[8] Definition 1 could accommodate first-order shifts in $F_s(\theta)$ by permitting either $\theta_s^\dagger = \underline{\theta}_s$ or $\theta_s^\dagger = \overline{\theta}_s$. It also allows for

mond and Stiglitz (1974) described the difference between distributions as satisfying a *single-crossing property,* whereas Hammond (1974) defined the corresponding random variables to be *simply intertwined.* Such a relationship corresponds to an increase in riskiness in the following sense: if, for some increasing function $u(\theta)$, $E[u(\theta)]$ is decreasing in s, and $v(\theta)$ is more risk-averse than $u(\theta)$ in the sense of Kenneth J. Arrow (1970) and John W. Pratt (1964), then $E[v(\theta)]$ is also decreasing in s.[9] We note that Definition 1 places no general restrictions on the response of the slope of demand to an increase in s: a clockwise rotation steepens the inverse demand curve when evaluated at $z_s^\dagger$; away from this quantity, however, this curve may well flatten.

Henceforth in this section, we order the family of distributions by a sequence of rotations.[10, 11] Equivalently, an increase in s results in a clockwise rotation of the inverse-demand curve around the quantity-price pair of $z = z_s^\dagger$ and $p = \theta_s^\dagger$. This pair may, however, move with s.

B. *Variance-Ordered and Elasticity-Ordered Distributions*

Here we describe two families of distributions that satisfy Definition 1. First, we consider distributions that share the same basic shape, but differ by mean and variance.

DEFINITION 2: *The family of distributions is* ordered by increasing variance *if*

$$F_s(\theta) = F\left(\frac{\theta - \mu(s)}{\sigma(s)}\right),$$

where $F(\cdot)$ is a continuous distribution with zero mean, unit variance, and strictly positive density, $\mu(s)$ and $\sigma(s)$ are continuously differentiable, $\sigma(s) > 0$, and $\sigma'(s) > 0$. The corresponding inverse-demand curve satisfies $P_s(z) = \mu(s) + \sigma(s)P(z)$ where $P(z) = F^{-1}(1 - z)$.

$F(\theta)$ determines the basic shape of each member of a variance-ordered family; $\mu(s)$ and $\sigma(s)$ are the mean and standard deviation. To verify the final part of the definition, notice that selling z units at price $P_s(z)$ requires $z = 1 - F_s(P_s(z))$, then solve for $P_s(z)$.[12]

An increase in the dispersion parameter s increases the standard deviation $\sigma(s)$. We do not, however, restrict to mean-preserving increases in risk; an increase in s may influence the mean as well as the standard deviation, so that the overall change involves both a pure (mean-preserving) rotation and a pure shift in the mean (that is, a shift in demand). Nevertheless, the net effect is a clockwise rotation of the inverse demand curve, which falls into our broad definition of a rotation. In fact, setting $\sigma(s) = s$ without loss of generality, inverse-demand becomes $P_s(z) = \mu(s) + sP(z)$ and the rotation quantity satisfies $z_s^\dagger = 1 - F(-\mu'(s))$.[13] Thus, the quantity-price pair of $z_s^\dagger$ and $\theta_s^\dagger$ around which the demand curve rotates is determined by the degree to which an increase in dispersion goes hand in hand with a shift in the mean.

For our second specification, we relate Definition 1 to the textbook notion of elasticity. Evaluated at a quantity-price pair of $z_s^\dagger$ and $\theta_s^\dagger$, a clockwise rotation reduces elasticity. If we order a family of inverse demand functions by

movement of the support $(\underline{\theta}_s, \bar{\theta}_s)$. When $\theta_s^\dagger \in (\underline{\theta}_s, \bar{\theta}_s)$, the support must (at least weakly) expand with s.

9 Formally, $v(\theta)$ is a concave transformation of $u(\theta)$ and hence displays a higher coefficient of absolute risk-aversion. Then, if an agent dislikes an increase in s, an agent who is more risk-averse will also dislike such an increase. This result was the focus of Hammond's (1974) paper, appeared as a lemma in the work of Sanford J. Grossman and Oliver D. Hart (1983, p. 151), and formed a basis for the contributions of Jewitt (1987, 1989).

10 We do not require demand curves separated by discrete differences in s to be ordered by rotation; Definition 1 refers only to local changes in s. Such discretely separated demand curves will, however, be ranked in the upper and lower tails. For instance, for $r < s$ and $z < \inf_{s' \in [r,s]} z_{s'}^\dagger$, we can be sure that $P_s(z) > P_r(z)$. If the rotation quantity $z_s^\dagger$ is constant then any two members of the family will indeed differ by a rotation.

11 When a family of distributions share a common mean, a rotation ordering ensures that any two members are ordered by second-order stochastic dominance. A partial converse holds. Suppose that, for a finite set $\tilde{S}$, distributions are ordered by second-order stochastic dominance. Following Rothschild and Stiglitz (1970) and Mark J. Machina and Pratt (1997), we may construct a sequence of mean-preserving spreads (and hence rotations) linking each pair, yielding an (expanded) family of distributions ordered by a sequence of rotations.

12 Write $F_s(P_s(z)) = 1 - z \Leftrightarrow F((P_s(z) - \mu(s))/\sigma(s)) = 1 - z \Leftrightarrow P_s(z) = \mu(s) + \sigma(s)F^{-1}(1 - z) = \mu(s) + \sigma(s)P(z)$.

13 Equivalently, $\theta_s^\dagger = \mu(s) - s\mu'(s)$. Observe that $z_s^\dagger$ is increasing if and only if $\mu''(s) \geq 0$.

decreasing elasticity, they must also be ordered by a sequence of rotations.[14] Of course, the rotation ordering is a relatively weak implementation of decreasing elasticity. For a stronger implementation, we might require an increase in s to reduce the elasticity everywhere.[15] To satisfy this more stringent criterion, we specify a family of inverse demand curves where each member has constant elasticity, and where that elasticity decreases with s.

DEFINITION 3: *A family of inverse demand functions is ordered by decreasing elasticity if* $\log P_s(z) = \mu(s) - s \log z$, *where* $\mu(s)$ *is continuously differentiable and* $s_H < 1$.

Under this specification, the elasticity of demand is $1/s$. The restriction $s_H < 1$ (so that demand is never inelastic) ensures the existence of a monopoly solution.[16] When $\mu'(s) < 0$, the family is ordered by a sequence of rotations, with a rotation quantity satisfying $z_s^\dagger = \exp(\mu'(s))$.[17] Although Definition 3 is couched in terms of the inverse demand functions, it is equivalent to specifying a family of distributions $F_s(\theta) = 1 - \exp(\mu(s)/s)\theta^{-1/s}$, where the support of θ has a lower bound $\underline{\theta}_s = \exp(\mu(s))$. Henceforth, when we refer to a family as "elasticity ordered" we mean that Definition 3 is satisfied. (Similarly, a "variance ordered" family will satisfy Definition 2.)

C. *A Monopolist's Preference for Extremes*

Here we study the response of a monopolist to changing dispersion. The monopolist faces costs of $C(z)$ from the production of z units. We write z_s^* for her optimized quantity, and $\pi(s)$ for her profits.

To facilitate our analysis, we distinguish between two separate cases. When $z_s^* > z_s^\dagger$, the monopolist acts as a "mass market" supplier to a relatively large fraction of consumers. As dispersion increases, the demand curve rotates clockwise around $z_s^\dagger$, and the willingness to pay of the marginal consumer is pushed down. This lowers profits: a mass-market monopolist dislikes increased dispersion. On the other hand, when $z_s^* < z_s^\dagger$, we say that the monopolist acts as a "niche" supplier by restricting supply to a relatively small fraction of the consumer base.[18] Following an increase in dispersion, the willingness to pay of the marginal consumer is pushed up. This raises profits: a niche monopolist likes increased dispersion. Summarizing:

$$z_s^* \gtrless z_s^\dagger \Leftrightarrow \frac{\partial \pi(s)}{\partial s} \lessgtr 0.$$

This observation does not, by itself, allow us to characterize the response of profits across the entire range of dispersion parameters, since z_s^* and $z_s^\dagger$ might switch ranking a number times. When $z_s^\dagger$ increases with s, however, the range of niche operations expands as dispersion increases. The result is that, once dispersion is high enough for the monopolist to switch from a mass-market to a niche posture, then for larger values of dispersion she will never wish to switch back. Figure 2 illustrates this argument, which we state formally as a proposition. A formal proof of this result follows from Lemma 1 below.[19]

PROPOSITION 1: *Suppose that the family of distributions is ordered by a sequence of rotations. If the rotation quantity* $z_s^\dagger$ *is increasing in* s (*equivalently,* $\theta_s^\dagger$ *is decreasing*), *then monopoly profits are quasi-convex in* s, *and hence maximized at an extreme* $s \in \{s_L, s_H\}$.

Thus, as s increases, profits first fall and then rise: they are "U-shaped," so that profits are high when consumers are either homogeneous or highly idiosyncratic. Of course, quasi-convexity also allows for profits that are monotonically increasing or decreasing in s.

Monotonicity of the rotation quantity is easy to check. To see this, fix z and suppose that

[14] If they were not, then we would be able to find a quantity-price pair around which the demand curve would rotate counterclockwise in response to an increase in s. This would lead to an increase in elasticity.

[15] Note this does not imply that the slope of the demand curve increases everywhere.

[16] Recall that a monopolist always chooses to operate on an elastic portion of her demand curve.

[17] When $\mu'(s) > 0$, an increase in s shifts $P_s(z)$ upward.

[18] Our terminology is related to that of Colin J. Aislabie and Clement A. Tisdell (1988). They referred to flat and steep demand curves as "bandwagon type" and "snob type," respectively.

[19] Any formal proofs omitted from the main text are contained in Appendix B.

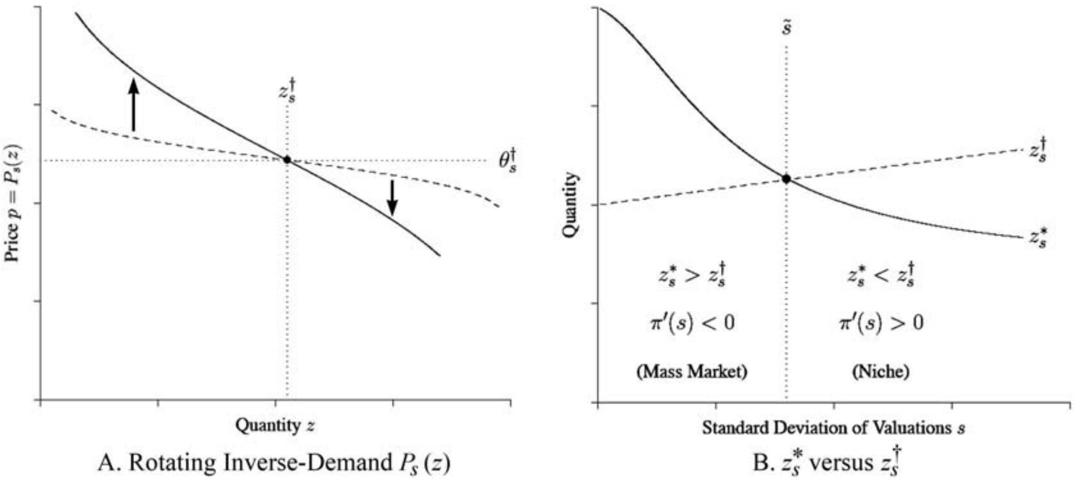

FIGURE 2. INVERSE-DEMAND ROTATION AND NICHE VERSUS MASS MARKET

$P_s(z)$ is locally increasing in s. By the definition of a rotation, this means that $z < z_s^\dagger$. If $z_s^\dagger$ is increasing, this means that for larger values of s, $P_s(z)$ must continue to increase. In other words, the inverse-demand function is a quasi-convex function of the dispersion parameter.

LEMMA 1: *When the family of distributions is ordered by a sequence of rotations, the rotation quantity $z_s^\dagger$ is increasing in s if and only if $P_s(z)$ is quasi-convex in s for all z. If the family is variance ordered, these statements hold if and only if $\mu'(s)/\sigma'(s)$ is weakly increasing in s. If it is elasticity ordered, these statements hold if and only if $\mu'(s)$ is weakly increasing in s.*

A few comments are in order. First, Lemma 1 may be stated in terms of the family of distribution functions $\{F_s(\theta)\}$. In fact, $z_s^\dagger$ is increasing in s if and only if $\theta_s^\dagger$ is decreasing in s, which in turn holds if and only if $F_s(\theta)$ is quasi-concave in s for each θ. Second, the criterion for variance-ordered families to generate quasi-convex profits is that the mean is a convex function of the standard deviation.[20] As we show subsequently, this criterion is satisfied for many applications. Third, Lemma 1 provides a straightforward proof of Proposition 1: (a) given monotonicity of $z_s^\dagger$, $P_s(z)$ is quasi-convex in s; and (b) $\pi(s) = \max_{z\in[0,1]}\{zP_s(z) - C(z)\}$ is the maximum of quasi-convex functions, and so is itself quasi-convex.

Appropriate variants of Lemma 1 and Proposition 1 hold even when a family of demand curves is not ordered by a sequence of rotations. To move beyond Definition 1, imagine a situation where a local increase in s raises and lowers $P_s(z)$ in multiple regions, separated by multiple rotation quantities. Such quantities form an alternating sequence of clockwise and counterclockwise rotation quantities. For quasi-convexity of profits, it is sufficient that each clockwise rotation quantity is increasing and that each counterclockwise rotation quantity is decreasing. In Appendix A we investigate this issue more fully.

D. *The Response of Monopoly Output*

In the previous section, we categorized a monopolist's marketing posture as either being a niche or a mass-market one, depending on whether she is serving fewer or more than $z_s^\dagger$ consumers, respectively. Here we turn attention to the response of monopoly output to changes in the dispersion parameter.

We begin by writing $\mathrm{MR}_s(z) \equiv P_s(z) + zP_s'(z)$ for the marginal revenue associated with the dispersion parameter s, and assume that it decreases with z. We further assume that $C(z)$ is increasing, convex, and continuously differentiable. Under these standard conditions, the mo-

[20] For $\sigma(s) = s$, it reduces to $\mu''(s) \geq 0$. Otherwise, and if $\mu(s)$ is increasing, it is equivalent to $-\mu''(s)/\mu'(s) \leq -\sigma''(s)/\sigma'(s)$, so that $\mu(s)$ is less concave than $\sigma(s)$, in the Arrow (1970) and Pratt (1964) sense.

nopolist's solution is characterized by the familiar condition $\mathrm{MR}_s(z_s^*) = C'(z_s^*)$. It is easy to see that the monopoly quantity increases in s if and only if marginal revenue evaluated at z_s^* increases with s. To characterize the behavior of z_s^* across the whole range of dispersion parameters, we can exploit the notion of rotation ordering introduced in Definition 1. Specifically, we say that an increase in s rotates marginal revenue if, for some $z_s^\ddagger$,

$$z \gtreqless z_s^\ddagger \Leftrightarrow \frac{\partial\ \mathrm{MR}_s(z)}{\partial s} \lesseqgtr 0.$$

We observe that since the demand curve rotates, and hence steepens, around $z_s^\dagger$, marginal revenue must fall when evaluated at $z_s^\dagger$. This implies, in turn, that $z_s^\ddagger < z_s^\dagger$.[21]

For many specifications, including the variance-ordered and elasticity-ordered cases of Definitions 2 and 3, increased dispersion results in a clockwise rotation of marginal revenue.

LEMMA 2: *If a family of inverse demand functions is either variance-ordered or elasticity-ordered, the corresponding family of marginal revenue curves is rotation-ordered.*

When $z_s^* > z_s^\ddagger$, an increase in s prompts a reduction in monopoly supply. Similarly, when $z_s^* < z_s^\ddagger$, it prompts an increase in monopoly supply. In fact, so long as $z_s^\ddagger$ is increasing in s, we can ensure that the monopolist's output is a "U-shaped" function of demand dispersion.

PROPOSITION 2: *If a family of marginal revenue curves is rotation-ordered and $z_s^\ddagger$ is increasing in s, then the monopoly quantity z_s^* is quasi-convex in s. (For a variance-ordered family with $\sigma(s) = s$, when $\mu(s)$ is convex these statements hold and profits are convex in s.)*

Bringing our results together, we identify three categories of response as s increases. When $z_s^\ddagger < z_s^\dagger < z_s^*$, output and profits fall with s: a contracting mass market. For intermediate values of s, where $z_s^\ddagger < z_s^* < z_s^\dagger$, profits rise with s: a contracting niche market. Finally, for larger values of dispersion, we have $z_s^* < z_s^\ddagger < z_s^\dagger$: an expanding niche market.

E. *Closing Comments*

We close this section by noting that whether a monopolist prefers high or low levels of dispersion may depend on her production costs. Intuitively, when costs are higher she is more likely to prefer to be a niche player, and hence to prefer extreme heterogeneity of demand. While this intuition need not always hold, it does in many cases. For instance, it holds whenever a family of distributions is variance-ordered with a constant mean and constant marginal costs.[22]

In many situations, a firm will have little control over the dispersion of consumer demand. For instance, an increase in dispersion might correspond to an exogenous increase in income inequality. In other situations, however, a firm may use advertising, marketing, and product design decisions to influence the heterogeneity of consumers' valuations. We now turn our attention to such decisions, beginning with the issue of product design.

II. Product Design

Here we consider the issue of product design faced by a monopolist. We show that, under natural conditions, product design decisions reduce to choosing the dispersion of demand. Product design decisions, therefore, provide a microfoundation for our analysis of Section I.

A. *Combining Characteristics*

A monopolist must assemble her product from a convex function of different characteristics.[23] For instance, we might imagine that a

[21] Precisely, at $z = z_s^\dagger$, we have $(\partial\ \mathrm{MR}_s(z)/\partial s) = (\partial P_s(z)/\partial s) + z(\partial P_s'(z)/\partial s) = z(\partial P_s'(z)/\partial s) < 0$, since demand becomes steeper at this point. Hence, $z_s^\dagger$ must lie to the right of marginal revenue's rotation point $z_s^\ddagger$.

[22] We omit a formal proof but describe part of the argument. It can be shown that $\pi(s_L) = \pi(s_H)$ implies $z_{s_L}^* > z_{s_H}^*$. Then, an envelope theorem argument shows that an increase in c lowers $\pi(s_L)$ more than $\pi(s_H)$, so if $s = s_H$ is optimal for some marginal cost c, it is also optimal for higher c.

[23] This "bundling" of characteristics ensures that our analysis is related to the contributions of George J. Stigler (1968), William J. Adams and Janet L. Yellen (1976), Richard Schmalensee (1982, 1984), and R. Preston McAfee et al. (1989). These authors analyzed the incentive of a multiproduct monopolist to sell goods separately, or to

restaurateur is deciding upon the combination of different ingredients that go into a meal.[24] Thus we are taking Lancaster's (1971) "characteristics" approach, in which (Lancaster, 1975, p. 567) the consumer "is assumed to derive ... satisfaction from characteristics which cannot in general be purchased directly, but are incorporated in goods." Other examples include automobiles and computers, which may be blends of performance, practicality, size, and weight.

We impose two simplifications. First, the expected valuation for the final product is invariant to the exact combination of characteristics. Second, we adopt the normal distribution.

APPLICATION 1: *The monopolist's product consists of n weighted characteristics, indexed by i, with typical weight* $\alpha_i \in [0, 1]$ *where* $\sum_{i=1}^{n} \alpha_i = 1$. *A consumer's valuation for the product satisfies* $\theta = \mu + \sum_{i=1}^{n} \alpha_i \eta_i$, *where* $\boldsymbol{\eta}$ *is an* $n \times 1$ *multivariate normal* $\boldsymbol{\eta} \sim N(0, \boldsymbol{\Sigma})$ *and* $|\boldsymbol{\Sigma}| > 0$.

Consumer valuations for the product are normally distributed, and satisfy

$$\theta \sim N(\mu, s^2)$$

$$\text{where } s^2 = \boldsymbol{\alpha}'\boldsymbol{\Sigma}\boldsymbol{\alpha} = \sum_{i=1}^{n} \alpha_i^2 \sigma_i^2 + 2 \sum_{i<j} \alpha_i \alpha_j \rho_{ij} \sigma_i \sigma_j,$$

where σ_i^2 is the variance of characteristic i, and ρ_{ij} is the correlation coefficient between η_i and η_j. From Section I, inverse demand is given by $P_s(z) = \mu + sP(z)$, where $P(z) = \Phi^{-1}(1 - z)$;[25] the characteristics mix matters only insofar as it influences the variance s^2 of valuations. Increases in s correspond to rotations and, since the rotation point is fixed (as μ is constant), Proposition 1 applies: a monopolist chooses either to minimize or maximize s. In fact, Proposition 2 also applies, so that profits are convex in s, not only quasi-convex.

To understand the implications for the characteristics mix, observe that s^2 is convex in the weights $\{\alpha_i\}$. To maximize s, the monopolist will wish to place all weight on the characteristic with the highest variance: she will "pander to the highest extreme." Many consumers will strongly dislike this product, but those who like it will love it. In contrast, to minimize s, the monopolist will often choose an interior solution (that is, $\alpha_i \in (0, 1)$ for each i) that offers "something for everyone." This product will not arouse strong disagreement, so that the distribution of willingness to pay is clumped around the mean μ.

PROPOSITION 3: *For Application 1, a monopolist will wish either to* (*a*) *set* $\alpha_i = 1$ *where* $\sigma_i^2 \in \max_j\{\sigma_j^2\}$, *or* (*b*) *choose the unique vector of convex weights* $\boldsymbol{\alpha}^*$ *to minimize* $\boldsymbol{\alpha}'\boldsymbol{\Sigma}\boldsymbol{\alpha}$.

While there is no uncertainty here, a design with something for everyone is equivalent to choosing a variance-minimizing stock portfolio with arbitrary covariance matrix, but common expected returns. When $n = 2$, $s^2 = \text{var}[\theta] = \alpha_1^2\sigma_1^2 + \alpha_2^2\sigma_2^2 + 2\alpha_1\alpha_2\rho_{12}\sigma_1\sigma_2$.[26] This is minimized when

$$\frac{\alpha_1}{\alpha_2} = \frac{\sigma_2^2 - \rho_{12}\sigma_1\sigma_2}{\sigma_1^2 - \rho_{12}\sigma_1\sigma_2},$$

so long as this is an interior solution (i.e., $\rho_{12} < \min\{\sigma_1/\sigma_2, \sigma_2/\sigma_1\}$). This is the optimal Lancastrian bundle for a mass-market posture. From a portfolio choice perspective, if the attributes are uncorrelated, their relative variances determine the mix, with the higher variance attribute receiving less weight. Similarly, an

employ pure or mixed bundles. Here, increasing the contribution of one characteristic to the mix entails the reduction of another. Nevertheless, our results just below illustrate how combinations of negatively correlated characteristics may reduce the risk in the distribution of valuations, as in Figure I of McAfee et al. (1989).

[24] This is an example of Lancaster (1966, p. 133), who noted that "[a] meal (treated as a single good) possesses nutritional characteristics but it also possesses aesthetic characteristics, and different meals will possess these characteristics in different relative proportions." As a literal illustration, Andrea Prat suggested to us the example of a radio station that must divide its airtime between different genres of music.

[25] Following common notation, $\Phi(\cdot)$ is the cumulative distribution function of the standard normal.

[26] This case is closely related to Schmalensee's (1982) pricing of product bundles in which reservation prices are drawn from a bivariate normal. He noted that "[pure] bundling is shown to operate by reducing buyer diversity" and that "changes in [dispersion] affect both the level and the elasticity of demand."

increase in ρ_{12} increases the weight on the first attribute if and only if it has a lower variance than the second.

In Application 1, the mean willingness to pay across the population is identical for each design. However, this is not required for design decisions to reduce to the study of demand rotation. One natural possibility is that the product design decision actually corresponds to deciding how many distinct features to add to a product. Adding more features may raise (or even lower) the average willingness to pay, yet also increase the variance of willingness to pay. The result of adding a new feature, however, will often be a rotation of the demand curve, although the rotation point may move as more features are added. For example, let us suppose that a consumer's willingness to pay is the sum of his valuations for the original product and for a new feature, drawn independently from two distributions. As long as the distributions are strongly unimodal (that is, their densities are log-concave), the convolution resulting from the product's development will reduce to a rotation.[27] Hence, even with changing average willingness to pay, product design or development decisions reduce to choosing a demand curve from a set that is ordered by rotations.[28]

B. *Design-Dependent Production Costs*

Most of our formal analysis in this paper assumes that a firm bears no additional costs depending on the level of demand dispersion. While this is completely appropriate when changes in dispersion are brought about by exogenous forces, endogenous choices of s may (but need not) be costly. Here, we discuss how incorporating this possibility influences our results, in the context of a product-design decision. Much of our discussion is relevant for other applications. We suppose the choice of s influences marginal costs of production. Fixed costs of dispersion generate qualitatively similar conclusions.

The primary issue is how costs of dispersion influence the quasi-convexity of payoffs and the preference for extremes. Let $c(s)$ denote the (constant) marginal production cost.[29] Depending on the circumstances, this cost might be either increasing or decreasing in s. That is, it might be more expensive to build a product that is a blend of many characteristics, or instead be more costly to achieve extremes along a single dimension. To ascertain sufficient conditions for extremes to be preferred, however, the curvature of $c(s)$ is critical, not the sign of its slope. For brevity, we consider increasing $c(s)$, and discuss how its curvature influences our results; related interpretations emerge from consideration of decreasing $c(s)$.

First suppose that $c(s)$ is concave, as would be the case when the greatest marginal cost increases are borne as the firm first begins moving toward a very specialized product. Define $\mu(s) = \mu - c(s)$, and note that profits given s and z can be written as $z(P_s(z) - c(s)) = z(\mu + sP(z) - c(s)) = z(\mu(s) + sP(z))$. That is, the dependence of costs on s can be subsumed into the profit margin simply by defining $\mu(s)$ appropriately. Since $\mu(s)$ is convex when $c(s)$ is concave, Lemma 1 and Proposition 1 indicate that overall profits are quasi-convex—a firm prefers extremes of product design. If $c(s)$ is instead convex, so that pushing closer to a niche ideal hastens marginal cost's climb, then $\mu(s)$ is concave, and there is no guarantee that extremes are optimal. The reason is that high levels of s may involve rapid increases in costs that overwhelm the rate at which the market price climbs, for given niche output levels.

Our conclusion is that production costs that are influenced by dispersion may provide a natural explanation as to why we do not observe only very low or very high levels of demand dispersion induced by product design. We note, however, that introducing such costs does not always overturn our results; convex costs of dispersion are necessary, though not sufficient, for this to occur.[30] Moreover, while the analysis

[27] Following earlier work by Samuel Karlin and Frank Proschan (1960) and Karlin (1957, 1968), Jewitt (1987) showed that such a convolution crosses the original distribution at most once.

[28] Furthermore, profits will still be quasi-convex, as long as the conditions identified in Lemma 1 are satisfied.

[29] Since different choices of characteristic weights $\{\alpha_i\}$ may give rise to the same value of s, we suppose that $c(s)$ is that product design having lowest marginal cost among all product designs that induce dispersion s.

[30] For example, if $P_s(z)$ is more convex than $c(s)$ in s, then it can be shown that profits remain convex in s.

is more complex when dispersion influences costs, understanding the demand side of the problem remains as simple as before, since increases in s continue to correspond to demand rotations.

III. Advertising and Information Provision

Here, we provide a different microfoundation for our core theory of demand rotations developed in Section I. We ask how advertising, sales, and marketing activities might shape demand. We identify two functions of such activities. First, they may involve promotional *hype,* which highlights the product's existence, shifting the demand curve outward. Second, they may involve the provision of *real information,* which often rotates the demand curve.

While we focus primarily on advertising, our analysis applies to many situations in which the availability of real information may be influenced. For example, the sponsorship of product reviews and demonstrations, the employment of knowledgeable and honest sales staff, return policies, and the provision of "test drives" may all increase the precision of real information.

A. *Hype versus Real Information in Advertising*

Traditionally, studies of advertising and their textbook counterparts have defined two main categories of advertising: advertisements that are *persuasive* and those that are *informative.*[31, 32] According to this taxonomy, advertising is persuasive if it increases each given consumer's willingness to pay for a product, while it is informative if it allows previously ignorant consumers to learn of a product's existence. While conceptually distinct, from a monopolist's perspective there is little qualitative difference between informative and persuasive advertisements: both serve to shift demand outward, and so can only increase sales at each possible price.[33]

The assumption that advertising will *always* increase sales is restrictive.[34] More generally, advertising, sales advice, and marketing may allow consumers privately to learn of their personal match with a product, and hence their true valuation for it. This need not always increase demand, as some consumers will learn that the product is not suited to their tastes even as others realize that it is. For instance, when an automobile manufacturer advertises the sporty nature of her product, this may dissuade consumers who seek a comfortable ride.

In response to these observations, we suggest a different taxonomy: an advertisement consists of both *hype* and *real information.* The hype corresponds to basic publicity for the product; a consumer might learn of the product's existence, price, availability, and any objective quality.[35] Absent other issues, hype will always increase demand. In contrast, real information allows consumers to evaluate their subjective preferences for a product, as when it emphasizes the sporty nature of an automobile.[36, 37] We show that real information increases the dispersion of consumers' valuations, and hence rotates the product's demand curve.

[31] Kyle Bagwell's (forthcoming) literature survey adds the "complementary" view: following Stigler and Gary S. Becker (1977) and Becker and Kevin M. Murphy (1993), supplies of advertising are complementary to the good being sold.

[32] We omit any strategic or dynamic issues. For instance, Phillip Nelson (1974), Carl Shapiro (1982), and Paul R. Milgrom and John Roberts (1986) viewed advertising as a signal of quality. We also neglect the interplay between advertising and market structure, a topic that attracted many studies following William S. Comanor and Thomas A. Wilson (1967, 1971).

[33] If, following Avinash Dixit and Victor Norman (1978), Shapiro (1980), and Gene M. Grossman and Shapiro (1984), we were to analyze welfare, then the distinction would return. This is, however, beyond the scope of this paper.

[34] The idea that advertising may make some people disinclined to buy is implicit in some of the existing literature, for example Schmalensee (1978). He considered a behavioral model in which advertising contains no direct information, but consumers are more likely to try highly advertised products. Equilibria may exist in which low-quality firms advertise more than high-quality firms. In such an equilibrium, if a consumer were instead rational, he would take higher levels of advertising as a negative indication of quality.

[35] An "objective quality" is a product feature that is valued by every consumer. For instance, an automobile manufacturer may advertise unambiguously valuable characteristics such as reliability and fuel economy.

[36] This real information may well stem from prior experience with a firm's product line (Johnson, 2005).

[37] Our emphasis on real information reflects concerns of the marketing literature. Alan J. Resnik and Bruce L. Stern (1977, p. 50) suggested that "for a commercial to be considered informative, it must permit a typical viewer to make a more intelligent buying decision after seeing the commercial than before seeing it."

This taxonomy suggests a two-step decision problem for a monopolist as she formulates her advertising campaign. For the first step, she must decide on the size of the campaign. An increase in the campaign size might involve the purchase of additional advertising space in newspapers, or a rise in the frequency of radio and television commercials. This will entail increased expenditure: hype is costly. For the second step, the monopolist must also choose the real-information content of her advertising. In many scenarios, whatever decisions are made at this stage will have few significant cost implications, and henceforth we omit any considerations of cost in the increased provision of real information.

Under our classification, any advertisement necessarily contains elements of hype, but can contain varying levels of real information. When an advertisement tells (or reminds) consumers only of the existence of a product, it is *pure hype,* and consumers learn little of their personal match with the product's characteristics.[38] Relating this to Robert Dorfman and Peter O. Steiner's (1954, p. 826) description of advertising as an expenditure that "influences the shape or position of a firm's demand curve," we see that the traditional notions correspond operationally to what we call pure hype, and change the *position* of a firm's demand curve, whereas what we call real information rotates demand, thereby changing its *shape.*

Our view is that the existing advertising literature, stemming from the classic contributions of Stigler (1961) and Gerard R. Butters (1977), helps us to understand the phenomenon of hype. In contrast, the idea that advertising may enable consumers to learn about their taste for a product has received surprisingly little attention in the literature. A notable exception is the contribution of Lewis and Sappington (1994).[39] They note that "suppliers often have considerable control over what is known about their products ... a supplier can help inform buyers about the uses and potential profitability of purchasing her goods and services." In contrast, Ottaviani and Prat (2001) expanded upon the work of Milgrom and Robert Weber (1982) by considering the incentive for a monopolist to commit to the public revelation of information, while authors such as Nicola Persico (2000) have considered the incentives for the responding decision maker (in this case, a consumer) to acquire additional information.[40]

This taxonomy is also connected to the ideas of vertical and horizontal product differentiation.

[38] A number of authors have assessed empirically the real information content of advertisements. Resnik and Stern (1977) and Stern et al. (1977) introduced a range of 14 "evaluative criteria" that reflect useful information in a television commercial. These criteria include the communication of quality, product components, taste, and packaging. These authors and their followers (Grahame R. Dowling, 1980; William Renforth and Sion Raveed, 1983; Preben Senstrup, 1985; Marc G. Weinberger and Harlan E. Spotts, 1989; and others) viewed a commercial as "informative" if it met one or more of the 14 criteria—a low hurdle. Nevertheless, these studies found that an extremely large fraction of advertisements were uninformative. For instance, Stern and Resnik (1991) found that, among a sample of 462 U.S. television commercials broadcast in Oregon in 1986, only 51.2 percent were informative. The proportion of completely uninformative advertisement drops dramatically in related studies of print advertisements (Gene R. Laczniak, 1979; David B. Taylor, 1983; Charles S. Madden et al., 1986). Nevertheless, we conclude that "pure hype" advertising is an established phenomenon.

[39] There are, of course, other related contributions that highlight the implications of giving an agent access to improved private information. For instance, Joel Sobel (1993) considered a principal's preferences when an agent is either informed or uninformed (intermediate situations were not considered) in a contracting problem. Jacques Crémer et al. (1998) modified David Baron and Roger Myeron's (1982) regulation model so that the agent must pay a cost to acquire his private information (for related work, see Chifeng Dai and Lewis, 2005, and the references therein). They investigated the principal's incentive to influence the agent's information-acquisition decision. Simon P. Anderson and Régis Renault (2005) considered consumer search for products with random match value. Although consumers always learn their true match before buying, the existence of search costs and the possibility of hold-up by a firm causes advertising that provides either match-specific information or price-specific information to be optimal in different cases. Kenneth L. Judd and Michael R. Riordan (1994) investigated a firm's incentives to influence the precision of information about the quality of its product. This precision affects profits by influencing the price distortions associated with signalling quality to consumers.

[40] In the closing sentence of their paper, Ottaviani and Prat (2001) stated that "[in] contrast to the case of public information, no general principle has yet emerged on the value of private information in monopoly." We suggest that a robust economic principle emerges from the analysis of this paper. Specifically, an increase in private information results in increased dispersion of consumer valuations, which corresponds to a rotation of the demand curve. When such rotations satisfy the conditions of Proposition 1 (as many do), the monopolist prefers the extremes, reaffirming the "all or nothing" insight of Lewis and Sappington (1994).

The hype component of advertising increases the perceived vertical position of the product, while real information influences the perceived horizontal position. Similarly, a firm influences the vertical position when it provides, for example, information regarding objective quality, while the horizontal position is affected by information about product attributes that some consumers dislike.

The remainder of this section is organized as follows. Our first model of consumer learning is based upon a setting of Lewis and Sappington (1994). We show how real information yields a rotation of the inverse demand curve. We then consider a second, richer, specification which incorporates issues of risk aversion and the idiosyncrasy of product design.

B. *Truth or Noise*

To formalize the notion that advertising may help consumers to learn, we separate a consumer's valuation (or willingness to pay) from his true payoff. Formally, his true taste is determined by some (unknown) parameter ω. Prior to making any purchase, the consumer may observe some signal x. We will refer to x as an advertisement. However, x may represent the outcome from any other sales or marketing activity, such as the consumer's inspection of a product sample. Lewis and Sappington's (1994) specification is obtained when the advertisement x represents "truth or noise" about the true payoff ω.

APPLICATION 2: *A risk-neutral consumer's true (monetary) payoff ω is drawn from $G(\omega)$, forming his prior distribution. He observes an advertisement x, but not ω. With probability $s \in [s_L, s_H] \subseteq [0, 1]$, $x = \omega$. With probability $1 - s$, x is an independent draw from $G(x)$.*

Under this specification, the advertisement x perfectly reveals the consumer's true preference with probability s, but is otherwise noise. The parameter s represents the accuracy of the information source. If $s = 1$, then x is perfectly revealing, whereas if $s = 0$, then it is pure noise.[41] It is straightforward to observe that $G(x)$ represents the marginal distribution of the advertisement x, as well as the marginal distribution (and hence prior) of the consumer's preference ω. Upon receipt of the advertisement x, a consumer is unable to distinguish between truth or noise. Bayesian updating, he obtains the posterior expectation

$$\theta(x) = \mathrm{E}[\omega | x] = sx + (1 - s)\mathrm{E}[\omega].$$

If $s = 0$, so that the advertisement is always noise, the consumer retains his prior expected valuation of $\mathrm{E}[\omega]$. If $s > 0$, then his posterior expectation is strictly increasing in x. If the monopolist sells z units at a common price, then she will sell to all consumers receiving an advertisement greater than $x = G^{-1}(1 - z)$. This requires her to set a price $P_s(z)$ satisfying

$$P_s(z) = sG^{-1}(1 - z) + (1 - s)\mathrm{E}[\omega].$$

This is linear, and hence convex, in s. Applying Lemma 1, the conditions of Proposition 1 are satisfied. In fact, $\pi(s) = \max_z\{zP_s(z) - C(z)\}$ is the maximum of convex functions, and is itself convex: this is even stronger than the quasi-convexity of Proposition 1.

PROPOSITION 4: *For Application 2, profits are convex in s and maximized by $s \in \{s_L, s_H\}$.*

In particular, if $[s_L, s_H] = [0, 1]$, then the monopolist would prefer either full information or complete ignorance. Notice that when $s = 0$, all consumers share a common willingness to pay of $\mathrm{E}[\omega]$. With a constant marginal cost of c, the monopolist will supply $z_0^* = 1$ units at a price $p_0^* = \mathrm{E}[\omega]$, so long as $\mathrm{E}[\omega] > c$. When $s = 1$, θ is drawn from $G(\cdot)$. Hence,

$$\pi(1) > \pi(0) \Leftrightarrow \max_z z[G^{-1}(1 - z) - c]$$

$$> \max\{\mathrm{E}[\omega] - c, 0\}.$$

This analysis, in essence, replicates Propositions 1 and 2 from Lewis and Sappington (1994).[42] Our approach, however, differs from

[41] These extreme cases are related to Yeon-Koo Che's (1996) study of return policies, which allow consumers to perfectly learn their true valuations ($s = 1$). Without such a policy, consumers are wholly ignorant ($s = 0$).

[42] Lewis and Sappington (1994) considered a price-discriminating monopolist. As we show in Section IV, our

theirs, and illustrates the simple economics of the result: an increase in the supply of real information increases the dispersion of the distribution of posterior expectations, and hence rotates the demand curved faced by a monopolist.

Which extreme ($s = s_L$ or $s = s_H$) is preferred by a monopolist depends upon the values taken by s_L and s_H. For instance, as s_L increases (perhaps via word-of-mouth communication or independent product reviews), the monopolist may be prompted to switch her preference from s_L to s_H.[43] Alternatively, the monopolist faces an incentive to lower s_L, perhaps by destroying any available real information.[44]

Application 2 has limitations. First, the truth or noise specification is extreme, although its tractability helps to illustrate key ideas in a clean fashion. Second, the demand rotations generated by changes in s are special in that $z_s^\dagger = 1 - G(\mathrm{E}[\omega])$ is constant, so that the inverse demand curve $P_s(z)$ rotates around a price point equal to the mean $\mathrm{E}[\omega]$.[45] Third, increases in s do not change the mean of the valuation distribution. This follows from the assumption of risk neutrality. If consumers are risk averse, then an increase in s will reduce the risk premium paid by a consumer who is uncertain of her preference for it. A reduction in risk premia will increase any willingness to pay. Hence, any rotation of demand will be accompanied by a shift; we might expect the rotation quantity $z_s^\dagger$ to move.

C. *Advertising and Risk Aversion*

In addressing these limitations, our next application allows for both consumer risk aversion and idiosyncrasy in product design.

APPLICATION 3: *The prior distribution of a consumer's true (monetary) utility satisfies* $\omega \sim N(\mu, \kappa^2)$. *He observes an advertisement* x, *but not* ω. *Conditional on* ω, $x \sim N(\omega, \xi^2)$. *His preferences exhibit constant absolute risk-aversion (CARA) with coefficient* λ.

For this example, the variance κ^2 indexes the dispersion of true consumer payoffs. Our interpretation is that this represents *idiosyncrasy* in the product's design.[46] When κ^2 is small, all consumers value the product in a similar way—a "plain-vanilla" design. In contrast, when κ^2 is large, true payoffs are more variable—a "love-it-or-hate-it" design. The second key parameter is ξ^2, indexing the noise in the advertising signal; an alternative representation is to write $\psi = 1/\xi^2$ for the *precision* of any real information provided to consumers.

The normal-CARA combination is widely used in many fields of economics, and leads to tractable results.[47] Given the receipt of an advertisement x, a consumer Bayesian updates his beliefs to obtain posterior beliefs over ω. Standard calculations confirm that his willingness to pay for the product will be the certainty equivalent $\theta(x)$ satisfying

$$\theta(x) = \frac{1}{1 + \psi\kappa^2}\left[\mu - \frac{\lambda\kappa^2}{2}\right] + \frac{\psi\kappa^2}{1 + \psi\kappa^2}\,x.$$

approach extends straightforwardly to a supplier of multiple quality-differentiated products.

[43] In empirical work, Daniel A. Ackerberg (2003, p. 1011) recognized that "if consumers obtain idiosyncratic information from consumption, we might expect prior experience and the resulting accumulation of information to generate relatively higher variance (across consumers) in experienced consumers' behaviors (for example, some consumers find out they like the brand, some find out they do not)." Ackerberg (2001) examined whether advertising provides information, or achieves its effects via other avenues such as prestige.

[44] This idea underpinned Patrick DeGraba's (1995) model of buying frenzies. By selling fewer units, the monopolist creates excess demand. Thus, consumers who delay purchasing, hence acquiring real information, will find no units available. Pascal Courty's (2003) analysis of ticket pricing under uncertainty was based on a related model, while in a paper by Andrew R. Biehl (2001), a similar mechanism prompts a monopolist to sell rather than lease.

[45] $\partial P_s(z_s^\dagger)/\partial s = G^{-1}(1 - z_s^\dagger) - \mathrm{E}[\omega] = 0$ solves to yield $z_s^\dagger = 1 - G(\mathrm{E}[\omega])$ or equivalently $\theta_s^\dagger = \mathrm{E}[\omega]$.

[46] A further microfoundation may be obtained via the Lancastrian specification of Application 1.

[47] For instance, Ackerberg's (2003) empirical examination of advertising, learning, and consumer choice followed Tülin Erdem and Michael P. Keane (1996) by using a normal specification for consumer learning. The normal-CARA specification is also central to the "career concerns" literature sparked by Bengt R. Holmström (1982). An implication of our results (including Propositions 4 and 5) is that the monopolist prefers extremes when consumers have access to real information. Interestingly, this conclusion also emerges from analyses of learning within organizations. For instance, Margaret A. Meyer (1994) considered task assignment in a team-production setting, and its implications for learning about agents' abilities. She found (pp. 1171–75) that the principal prefers tasks to be either completely shared or completely specialized; there is no interior solution.

This is a weighted average of the ex ante certainty equivalent and the ex post advertisement realization, where the weights depend upon the relative precision of the prior and the signal.

To characterize the demand curve, we consider the distribution of θ. Realized advertisements follow the distribution $x \sim N(\mu, \kappa^2 + \xi^2)$. Consumer valuations are linear in x, and satisfy

$$\theta \sim N\left(\mu - \frac{\lambda \kappa^2}{2(1 + \psi\kappa^2)}, \frac{\psi\kappa^4}{1 + \psi\kappa^2}\right).$$

For any choice of κ^2 and ψ, the distribution remains within the normal family. Thus, changes in either parameter yield a variance-ordered family with a changing mean:

$$P_{\kappa^2,\psi}(z) = \mu - \frac{\lambda\kappa^2}{2(1 + \psi\kappa^2)} + P(z)\sqrt{\frac{\psi\kappa^4}{1+\psi\kappa^2}}$$

$$\text{where } P(z) = \Phi^{-1}(1 - z),$$

where the inverse-demand curve is now indexed by both κ^2 and ψ, rather than a single dispersion parameter s. Notice that the standard deviation $\sqrt{\psi\kappa^4/(1 + \psi\kappa^2)}$ is increasing in κ^2 and ψ: the valuation distribution is riskier when the product design is more idiosyncratic (an increase in κ^2) or when advertising is more informative (an increase in ψ).

In contrast to Application 2, and so long as $\lambda > 0$, the mean valuation is *not* invariant to κ^2 and ψ. Fixing ψ, $E[\theta]$ is decreasing in κ^2; for a more idiosyncratic product, a purchase represents more of a gamble. This is reflected by a higher risk premium, and the increase in variance is accompanied by an inward shift of the inverse-demand curve. Fixing κ^2, an increase in ψ increases $E[\theta]$; more informative advertising reduces the risk premium, and hence the rotation is accompanied by an outward demand shift. Thus increases in κ^2 and ψ both rotate the demand curve clockwise, but shift the mean in opposite directions.

These properties ensure that the quantity around which the demand curve rotates will not be constant. Nevertheless, Proposition 1 admits such possibilities, and a monopolist will wish to choose extreme values for her product-design idiosyncrasy and the precision of information. We write $\pi(\kappa^2, \psi)$ for profits as a function of $\kappa^2 \in [\kappa_L^2, \kappa_H^2]$ and $\psi \in [\psi_L, \psi_H]$.

PROPOSITION 5: *For Application 3, $\pi(\kappa^2, \psi)$ is quasi-convex in κ^2 and in ψ. Furthermore,*

$$\frac{\partial \pi}{\partial \kappa^2} > 0 \Rightarrow \frac{\partial \pi}{\partial \psi} > 0,$$

$$\text{and } \frac{\partial \pi}{\partial \psi} < 0 \Rightarrow \frac{\partial \pi}{\partial \kappa^2} < 0.$$

If $\lambda = 0$ then these implications also hold with the reverse inequalities: with risk-neutral consumers, profits are increasing in κ^2 if and only if they are increasing in ψ.

Proposition 5 implies that the monopolist will choose, if she is able to do so, $\kappa^2 \in \{\kappa_L^2, \kappa_H^2\}$ and $\psi \in \{\psi_L, \psi_H\}$: an "all-or-nothing" approach to design idiosyncrasy and real-information provision. Furthermore, she will never choose $\kappa^2 = \kappa_H^2$ together with $\psi = \psi_L$. To see why, notice that an increase in κ^2 increases the variance of the valuation distribution, while reducing the mean. Increasing ψ can achieve a similar increase in variance, while increasing the mean. Hence increases in κ^2 will always be accompanied by increases in ψ: more idiosyncratic products are complemented by detailed advertising and marketing activities.[48]

When consumers are risk neutral, the inverse demand curve reduces to $P_{\kappa^2,\psi}(z) = \mu + sP(z)$, where $s = \sqrt{\psi\kappa^4/(1 + \psi\kappa^2)}$. The monopolist's desire to choose extreme values for s manifests itself as the pairing of κ_L^2 with ψ_L or of κ_H^2 with ψ_H: a monopolist will never engage in highly informative advertising of a plain-vanilla product.

When consumers are risk averse, the situation is subtly different. It is possible that the monop-

[48] Similarly, a plain-vanilla design complements a promotional campaign of pure hype. This does not imply that a mass-market monopolist will refrain from changing a product's design. Replacing an old product may well destroy any real information about it, hence lowering ψ_L. Thus pure hype may be accompanied by regular design revisions.

olist may pair detailed advertising (ψ_H) with a vanilla product (κ_L^2). To see why, recall that increases in κ^2 and ψ shape demand in different ways; the demand curve rotates around different points. In a slight abuse of notation, we write $z_{\kappa^2}^{\dagger}$ for the quantity satisfying $\partial P_{\kappa^2,\psi}(z_{\kappa^2}^{\dagger})/\partial\kappa^2 = 0$, and similarly for $z_{\psi}^{\dagger}$. Based on earlier analysis, $z_{\kappa^2}^{\dagger} < ½ < z_{\psi}^{\dagger}$.[49] If the monopoly supply satisfies $z_{\kappa^2}^{\dagger} < z^* < z_{\psi}^{\dagger}$, then profits will be locally increasing in ψ, and yet decreasing in κ^2: more detailed advertising of a less idiosyncratic product is desirable. Since $z_{\psi}^{\dagger}$ is increasing with λ while $z_{\kappa^2}^{\dagger}$ is decreasing (see footnote 49), this outcome is more likely for large λ: the optimal policy then reduces consumers' risk as much as possible.[50]

This discussion, driven by a possible tension between design idiosyncrasy and real information, reveals that families of distributions may be rotation-ordered and yet violate the monotonicity criterion of Proposition 1. For such cases, it may be (but need not be) that profits are not quasi-convex. In particular, suppose that initial increases in a dispersion parameter s primarily serve to increase ψ, but further increases primarily increase κ^2. Then, since the point at which the demand rotates is moving from a higher quantity to a lower quantity, in some region of prices it may be that the demand curve is first pushed out (from an increase in ψ) and then pulled back (from an increase in κ^2). If a monopolist were pricing in that region (as she may be) profits might first increase, but then decrease.

It follows that there is no fully general result that all rotation-ordered families of distributions must induce quasi-convex profits.[51] Nonetheless, inasmuch as a monopolist can *independently* control both product idiosyncrasy (κ^2) and real information (ψ), she always prefers extreme values for each, although which extreme value is preferred (low or high) may (occasionally, and not when the consumer is risk neutral) differ.

IV. Product Lines

So far, we have restricted attention to the sale of a single product. Here, we extend our analysis by studying a multiproduct monopolist selling a range of quality-differentiated goods. We explore the relationship between her product line and the dispersion of demand. Constructing a standard model of quality-based price discrimination, following Mussa and Rosen (1978), we use Johnson and Myatt's (2003, 2006) "upgrades approach" for our analysis. This approach may be used for other market structures, which are considered in Section V.

A. *The Upgrades Approach to Product-Line Design*

A monopolist is able to offer n distinct product qualities, where the quality of product i is q_i, and $0 < q_1 < \cdots < q_n$. A consumer's type θ is his willingness to pay for a single unit of quality. Thus, if a type θ consumes quality q at price p, he receives a net payoff of $\theta q - p$. Faced with a set of prices $\{p_i\}$, a consumer purchases a single unit of the product that maximizes $\theta q_i - p_i$, unless doing so yields a negative payoff, in which case he purchases nothing.[52]

The monopolist sets quantities for each of her n products. As in Section I, we write $P_s(z)$ for the type θ with a mass z of others above him. This corresponds to the inverse demand curve for a single product of quality $q = 1$, and is the basis for the entire system of inverse demand.

[49] Recall that $P_{\kappa^2,\psi}(1/2) = \mathrm{E}[\theta]$. This is the certainty equivalent of a purchase following receipt of a "neutral" advertisement $x = \mu$. $\mathrm{E}[\theta]$ is increasing in ψ, but decreasing in κ^2, yielding the stated inequalities. In fact, straightforward algebraic manipulation confirms that

$$z_{\psi}^{\dagger} = \Phi(\lambda\sqrt{\psi\kappa^4/(1+\psi\kappa^2)})$$

$$\text{and}\quad z_{\kappa^2}^{\dagger} = \Phi(-\lambda/(2+\psi\kappa^2)\sqrt{\psi(1+\psi\kappa^2)}).$$

By inspection, these rotation quantities are increasing in both κ^2 and ψ^2.

[50] As $\lambda \to 0$, these two different rotation quantities converge to ½ , and this possibility is eliminated.

[51] Similarly, there is no general result relating Blackwell-ordered (David Blackwell, 1953) advertising-signal distributions to quasi-convexity of payoffs. Lewis and Sappington (1994) provided a counterexample, which, utilizing the insights of our approach, can be seen to involve a decreasing rotation quantity $z_s^{\dagger}$.

[52] This framework can incorporate quantity-based second-order price discrimination by a monopolist supplying multiple-unit bundles. Product i then corresponds to a particular multi-unit bundle size.

We write $\{z_i\}$ for the supplies of the n qualities. For the market to clear, the individual with $\sum_{j=1}^{n} z_j$ others above him must be just indifferent between purchasing quality q_1 and not purchasing at all. We define the cumulative variables $\{Z_i\}$ to satisfy $Z_i = \sum_{j=i}^{n} z_j$. With this notation, $p_1 = q_1 P_s(Z_1)$. Thus, the consumer of type $\theta = P_s(Z_1)$ is just willing to buy product 1. Similarly, a consumer with Z_i others above him must be just indifferent between products i and $i-1$, so that $p_i - p_{i-1} = P_s(Z_i)(q_i - q_{i-1})$. Writing $p_0 = q_0 = 0$,

$$\Delta p_i = P_s(Z_i)\Delta q_i \qquad \text{for all } i \in \{1, \ldots, n\}.$$

We interpret Δp_i as the price of an "upgrade" from quality q_{i-1} to the next quality q_i. Thus, we can view the monopolist as supplying Z_1 units of a "baseline" product of quality q_1, at a price of p_1. She then supplies successive upgrades to this baseline product in order to achieve qualities above this. So, a product of quality q_2 consists of a baseline product at a price of p_1, bundled together with an upgrade Δq_2 priced at Δp_2. Similarly, quality q_3 consists of quality q_2 bundled together with an upgrade Δq_3 priced at Δp_3.

On the cost side, the monopolist manufactures quality q_i at a constant marginal cost of c_i. The profits of the monopolist can be determined, given that she has chosen a profile of upgrade supplies $\{Z_i\}$, where $Z_i \geq Z_{i+1}$ is required for this profile to be feasible (since $z_i = Z_i - Z_{i-1} \geq 0$). In particular, profits are $\pi = \sum_{i=1}^{n} Z_i(\Delta p_i - \Delta c_i)$, or equivalently

$$\pi = \sum_{i=1}^{n} \pi_i$$

$$\text{where } \pi_i = \Delta q_i \times Z_i\left[P_s(Z_i) - \frac{\Delta c_i}{\Delta q_i}\right].$$

Hence a monopolist's profits are the sum of her profits in each of the upgrade markets. We impose two simplifying conditions.[53] First, quality-adjusted marginal revenue $\mathrm{MR}_s(z) = P_s(z) + zP_s'(z)$ is decreasing in z for all s. Second, there are decreasing returns to quality:[54]

$$\frac{\Delta c_n}{\Delta q_n} > \frac{\Delta c_{n-1}}{\Delta q_{n-1}} > \cdots > \frac{\Delta c_2}{\Delta q_2} > \frac{\Delta c_1}{\Delta q_1} > 0.$$

This assumption implies that a monopolist will not offer a product line that exhibits "gaps." That is, if products i and $k > i$ are optimally in positive supply, then so is any product j with $i \leq j \leq k$. We can see this by examining the unconstrained supply Z_{is}^* that maximizes profits in upgrade market i for any given dispersion parameter s. Since marginal revenue is decreasing, we can employ the usual first-order condition, equating marginal cost to marginal revenue, where possible, to characterize Z_{is}^*. In fact,

$$\mathrm{MR}_s(Z_{is}^*) = P_s(Z_{is}^*) + Z_{is}^* P_s'(Z_{is}^*) = \frac{\Delta c_i}{\Delta q_i},$$

so long as the quality-adjusted marginal revenue crosses $\Delta c_i/\Delta q_i$; otherwise the optimal (unconstrained) upgrade output in market i is either zero or one. Now, given this, and the fact that the right-hand side is strictly increasing in i, Z_i^* must be decreasing. It is strictly decreasing whenever we have an interior solution: $Z_i^* \in (0, 1)$ implies $Z_i^* > Z_{i+1}^*$. Since the monotonicity constraint $Z_i \geq Z_{i+1}$ on upgrade supplies is satisfied, the set of upgrade supplies that maximize profits in each upgrade market independently also solves the monopolist's multiproduct maximization problem, incorporating the appropriate monotonicity constraints. Moreover, as claimed, there clearly can be no gaps in the product line of the firm.

Note that, when an interior solution maximizes π_i, the equalization of marginal revenue and (quality-adjusted) marginal cost in the upgrade market reveals the simple economics of second-degree price discrimination: it corresponds to the usual monopoly solution in the space of upgrade supplies, rather than the supplies of the products themselves.

[53] Neither of these conditions is required for Proposition 6 to hold.

[54] If we were to interpret quality q_i as a multi-unit bundle, then the decreasing returns to quality assumption holds if (a) the monopolist faces a constant marginal cost of producing each unit, and (b) consumers experience decreasing marginal utility as a bundle's size grows.

B. *Product Lines and Distribution Families*

We now consider a family of rotation-ordered type distributions $\{F_s(\theta)\}$, with associated inverse-demand functions $\{P_s(z)\}$. From our work in Section I we know that the profits of a monopolist selling a single product are either quasi-convex or convex in s. Similar results hold for a multiproduct monopolist.[55]

PROPOSITION 6: *If $P_s(z)$ is a convex function of s for each $z \in [0, 1]$, then profits $\pi(s)$ are convex in s, and maximized by $s \in \{s_L, s_H\}$. Hence, if the family is either variance-ordered with $\sigma(s) = s$, or elasticity-ordered, and $\mu(s)$ is convex, then $\pi(s)$ is convex.*

The intuition behind Proposition 6 is somewhat more complicated than when there is but a single product. To better understand it, note that we may divide the monopolist's upgrades into two subsets for any fixed s. We recall that $z_s^\dagger$ is the quantity around which the (now quality-normalized) inverse demand rotates from a change in s. Suppose that, for some i,

$$\frac{\Delta c_i}{\Delta q_i} > \mathrm{MR}_s(z_s^\dagger) > \frac{\Delta c_{i-1}}{\Delta q_{i-1}} \Rightarrow Z_{is}^* < z_s^\dagger < Z_{(i-1)s}^*.$$

Hence, for all upgrades $j \geq i$, optimal supply is below the rotation quantity $z_s^\dagger$, and for all upgrades $j < i$ supply is above $z_s^\dagger$. Profits are the sum of two components:

$$\pi(s) = \underbrace{\sum_{j<i} \Delta q_j \times Z_{js}^* \left[P_s(Z_{js}^*) - \frac{\Delta c_j}{\Delta q_j} \right]}_{\text{Mass-market upgrades}}$$

$$+ \underbrace{\sum_{j \geq i} \Delta q_j \times Z_{js}^* \left[P_s(Z_{js}^*) - \frac{\Delta c_j}{\Delta q_j} \right]}_{\text{Niche upgrades}}.$$

The "mass-market upgrades," which include the baseline product q_1, have optimal supplies serving the mass market. A local increase in s will reduce the profitability of such upgrades. In contrast, the supplies of "niche upgrades," which include supply of the upgrade to the maximum-feasible quality q_n, are restricted to a niche market. The profitability of them is increasing in s. Thus, the two sets of upgrades present a tension for the monopolist: a mass-market posture is optimal for some, while a niche posture is optimal for the remainder. That is, increasing s raises the profits in some upgrade markets but lowers it in others.[56] Despite this tension, overall profits are convex in the dispersion parameter s.[57]

To investigate this tension further, suppose that there are N potential qualities, but that the monopolist is initially restricted to qualities q_n and below, where $n < N$. If $\Delta c_n/\Delta q_n < \mathrm{MR}_s(z_s^\dagger)$, then a monopolist's product line consists entirely of mass-market upgrades, and her incentive is to lower s. If, however, n increases sufficiently, perhaps due to innovation, and if $\Delta c_N/\Delta q_N > \mathrm{MR}_s(z_s^\dagger)$, then $\Delta c_n/\Delta q_n > \mathrm{MR}_s(z_s^\dagger)$. The monopolist now has more niche upgrades in her portfolio. Therefore, as n increases, her profits are increasingly biased toward a niche operation, in the sense that for a given s, the local change in profits from an increase in s is increasing. Moreover, under some mild technical conditions, increases in n make it more likely that the monopolist will wish to choose s_H over s_L.

There are numerous further implications that follow immediately from Proposition 6. For ex-

[55] As stated previously, Proposition 6 does not require either the assumption of decreasing marginal revenue or that of decreasing returns to quality assumed earlier in this section. These two assumptions ensure that the constraints $Z_{is} \geq Z_{(i+1)s}$ may be neglected. Johnson and Myatt (2003) exploited failures of these two assumptions to explain the phenomena of "fighting brands" and "product-line pruning."

[56] This suggests that a firm would like to set different levels of s in different upgrade markets. Each upgrade market is influenced, however, by the same underlying distribution, and so whatever moves the demand curve in one upgrade market moves it similarly in another.

[57] It is now clear why, unlike in the case of a single product, mere quasi-convexity of $P_s(z)$ in s is insufficient to guarantee quasi-convexity of profits here, and yet convexity of $P_s(z)$ ensures that profits are convex. The technical reason is that, while quasi-convexity of $P_s(z)$ does ensure that profits are quasi-convex in each upgrade market, there is no guarantee that the sum of the profits is quasi-convex, since the sum of quasi-convex functions need not be quasi-convex. Convexity of $P_s(z)$ resolves this issue, since sums of convex functions are convex, and moreover leads to the stronger conclusion that profits are convex in s. We note that convexity of $P_s(z)$ in s holds for Applications 1–3.

ample, much of the discussion in Sections I to III for single-product firms is readily applicable to the multiproduct setting, and we do not recapitulate here. On the other hand, the question of how the entire product line is influenced by changes in s is necessarily not considered in the single-product setting, and so we turn now to that issue.

C. *Product-Line Transformations*

Now we consider how either exogenous or endogenous changes in the dispersion of consumer demand for quality influence a monopolist's product line. To sharpen our analysis, we consider a family of distributions for which the corresponding quality-normalized marginal revenue curves are rotation-ordered. This applies, for instance, when the family is either variance-ordered or elasticity-ordered, as confirmed by Lemma 2. Recall that for such a family, and for each s, there is a quantity $z_s^{\ddagger}$ such that

$$z \gtrless z_s^{\ddagger} \Leftrightarrow \frac{\partial \ \mathrm{MR}_s(z)}{\partial s} \lessgtr 0.$$

To ease notational burden, we write $\mathrm{MR}_s^{\ddagger} \equiv \mathrm{MR}_s(z_s^{\ddagger})$. For an upgrade i satisfying $\Delta c_i/\Delta q_i > \mathrm{MR}_s^{\ddagger}$, it must be that $Z_{is}^* < z_s^{\ddagger}$. The marginal revenue from an extra unit of this upgrade is locally increasing in s, and hence so is Z_{is}^*. On the other hand, the opposite inequality $\Delta c_i/\Delta q_i > \mathrm{MR}_s^{\ddagger}$ implies that Z_{is}^* is locally decreasing in s. Summarizing,

$$\frac{\Delta c_i}{\Delta q_i} \gtrless \mathrm{MR}_s^{\ddagger} \Leftrightarrow \frac{\partial Z_{is}^*}{\partial s} \gtrless 0.$$

Thus, the supply of higher upgrades (with correspondingly higher quality-adjusted marginal cost $\Delta c_i/\Delta q_i$) increases with s, while the supply of lower upgrades decreases with s. Recall that Z_{is}^* is the total supply of all qualities q_i and greater. Hence, setting $Z_0 = 1$ for convenience for all s, $1 - Z_{is}^*$ is the supply of qualities strictly below i. Thus, $\{1 - Z_{is}^*\}$ may be used to characterize the distribution of qualities offered by the monopolist. The argument above demonstrates that this distribution undergoes a clockwise rotation as s increases.

PROPOSITION 7: *Suppose that for a family of distributions, the associated quality-normalized marginal revenue curves are ordered by a sequence of rotations.* (*This holds for families that are either variance-ordered or elasticity-ordered.*) *Then, the associated distributions of qualities* $\{1 - Z_{js}^*\}$ *offered by a monopolist are ordered by a sequence of rotations. Furthermore, an increase in dispersion results in an expansion of the product line: if a product is offered in positive supply for some* s, *then it will continue to be offered for higher* s.

When $\Delta c_1/\Delta q_1 < \mathrm{MR}_s^{\ddagger} < \Delta c_n/\Delta q_n$, an increase in s will result in an increase in the dispersion of a monopolist's product line. In other situations, the notion of a rotation is broad enough to encompass shifts in the distribution of qualities. When $\mathrm{MR}_s^{\ddagger} < \Delta c_1/\Delta q_1$, an increase in s will result in an expansion of Z_{is}^* for all i, and hence the product line shifts upward. Similarly, when $\mathrm{MR}_s^{\ddagger} < \Delta c_n/\Delta q_n$, the entire product line moves down.[58]

The final statement of Proposition 7 allows us to predict changes in the boundaries of a product line. In particular, the firm will expand the set of qualities offered (so that the product line lengthens) as dispersion increases; furthermore, the extension may occur at both the low- and high-quality ends. One concrete application is the evolution of a product line over time. If demand becomes more disperse over time for whatever reasons, for example a growth in information or increased demographic dispersion, the optimal strategy for a monopolist will be to begin with a focused product line, but to then expand it in both directions.

Of course, this does not imply that increased dispersion will increase the total quantity supplied, since Z_{1s}^* (the supply of the baseline product) is decreasing in s when $\mathrm{MR}_s^{\ddagger} > \Delta c_1/\Delta q_1$. Thus, as the dispersion of types increases, the monopolist may serve less of the total market. She will do so, however, using more products.

[58] The three possible responses of the product line to s fall within Hammond's (1974) notion of "simply related" distributions. Thus, a change in s results in two distributions of θ that are simply related, and the corresponding distributions of qualities supplied are also simply related.

V. Competition

Our analysis so far has considered the case of monopoly. Our key results concerning the influence of dispersion on a firm's activities and profits carry over straightforwardly to a perfectly competitive industry.[59] We focus here instead on both single-product and multiproduct Cournot oligopolies, which involve strategic effects.

A. *Single-Product Cournot Oligopoly*

We now consider m symmetric firms competing in quantities. We begin with the single-product case and assume, for each s, that marginal revenue is decreasing, and that each firm faces a constant marginal cost of c.

Under these conditions, there is a unique Cournot Nash equilibrium (see, for instance, Xavier Vives, 1999) in which the industry produces a total output z_s^* satisfying the usual condition

$$\mathrm{MR}_s(z_s^*) = c$$

where

$$\mathrm{MR}_s(z) \equiv P_s(z) + \frac{zP_s'(z)}{m} = P_s(z)\left[1 - \frac{\eta_s(z)}{m}\right],$$

and where $\eta_s(z)$ is the inverse of the elasticity of demand. By inspection, $\mathrm{MR}_s(z)$ is the marginal revenue faced by a firm when it produces an equal share of total output z.

When a family of demand curves is rotation-ordered, price rises or falls with s depending on whether $z < z_s^\dagger$ or $z > z_s^\dagger$. In the absence of competition, these inequalities also determine whether a monopolist's profits rise or fall. Under quantity competition, however, there is a further, strategic, effect: the change in s influences opponents' output, and therefore an individual firm's profits. Writing $\pi(s)$ for a firm's profits in equilibrium,

$$\frac{d\pi(s)}{ds} = \frac{z_s^*}{m}\left[\frac{\partial P_s(z^*)}{\partial s} + \frac{(m-1)P_s'(z_s^*)}{m}\frac{dz_s^*}{ds}\right].$$

The first bracketed term is the direct effect on market price from an increase in s. The second bracketed term is the strategic effect on an oligopolist's profits, and from inspection it depends on whether an increase in dispersion raises or lowers industry output. Industry output will, in turn, depend upon the change in the slope of the inverse demand curve:

$$\frac{dz_s^*}{ds} = -\frac{\partial\,\mathrm{MR}_s(z_s^*)/\partial s}{\partial\,\mathrm{MR}_s(z_s^*)/\partial z}$$

$$= -\frac{1}{\partial\,\mathrm{MR}_s(z_s^*)/\partial z}\left[\frac{\partial P_s(z_s^*)}{\partial s} + \frac{z_s^*}{m}\frac{\partial^2 P_s(z_s^*)}{\partial s\partial z}\right].$$

The rotation ordering of Definition 1 places no general restriction on the effect of s on the slope of the inverse demand curve.[60] Thus, even if a family of distributions (and hence demand curves) is rotation ordered, we cannot guarantee that profits are quasi-convex in s, simply because of the strategic effect. It follows that

[59] For each given s, a competitive equilibrium is determined by the intersection of the demand curve and the industry supply curve. Whenever the rotation quantity $z_s^\dagger$ lies to the left of the industry supply curve, a clockwise rotation leads to a fall in both output and price, as the demand curve traces the supply curve down. Once the (increasing) rotation quantity passes through to the right of the supply curve, further clockwise demand rotations trace the supply curve up, leading to an increase in output and price. In other words, industry output and price are both quasi-convex in s. Now, recalling that the output and profits of a perfectly competitive firm rise and fall with the price, we see that an individual firm's output and profits are quasi-convex in s. Turning to a competitive multiproduct industry, output of each upgrade is determined by the intersection of a rotating demand curve and the quality-adjusted upgrade costs. Following the spirit of the proof of Proposition 7, in which a rotating marginal revenue curve was critical, demand rotation ensures that the distribution of qualities supplied by the industry is ordered by a sequence of rotations, and that the industry product line expands as demand rotates.

[60] Under a rotation ordering, inverse demand will become steeper with s when evaluated at $z_s^\dagger$. Thus,

$$z_s^* = z_s^\dagger \Rightarrow \frac{\partial P_s(z_s^*)}{\partial s} = 0 \Rightarrow$$

$$\frac{d\pi(s)}{ds} = -\frac{(m-1)(z_s^*)^2}{m^3}\frac{P_s'(z_s^*)}{\partial\,\mathrm{MR}_s(z_s^*)/\partial z}\frac{\partial^2 P_s(z_s^*)}{\partial s\partial z} > 0,$$

so that oligopoly profits increase with s. When $z_s^* = z_s^\dagger$, an increase in s has no first-order effect on the price $P_s(z^*)$. Thus, the only effect remaining is the strategic effect. Since inverse demand steepens, a firm's competitors reduce their output, which increases profits.

Proposition 1 cannot carry over to an oligopolistic environment unless further restrictions are placed on the slope of demand and its response to s.

We now consider a family of demand curves that are elasticity-ordered (Definition 3), so that the inverse elasticity of demand satisfies $\eta_s(z) = s$ for all z. Under this specification,

$$\mathrm{MR}_s(z) = P_s(z)\left[1 - \frac{s}{m}\right]$$

$$\Rightarrow \frac{dz_s^*}{ds} = \frac{z_s^*}{s}\left[\frac{\partial \log P_s(z_s^*)}{\partial s} - \frac{1}{m-s}\right],$$

so that the strategic effect is predictable. When the demand curve exhibits constant elasticity, output and profits are easy to calculate, as confirmed by the following result.

PROPOSITION 8: *If the family of demand curves is elasticity-ordered, then the corresponding family of marginal revenue curves is rotation-ordered. Cournot output and profits satisfy*

$$z_s^* = \left[\frac{(m-s)e^{\mu(s)}}{mc}\right]^{1/s}$$

$$\text{and } \pi(s) = \frac{sc}{m(m-s)}\left[\frac{(m-s)e^{\mu(s)}}{mc}\right]^{1/s}.$$

When $\mu''(s) \geq 0$, profits are convex in s. Output is quasi-convex in s if $\mu''(s) \geq (m-s)^{-2}$.

The criterion $\mu''(s) \geq 0$ in this proposition holds if and only if the rotation quantity $z_s^\dagger$ is increasing. Thus, when the monotonicity criterion of Proposition 1 holds for an elasticity-ordered family, oligopoly profits (rather than just monopoly profits) are a convex function of dispersion, and are maximized at an extreme $s \in \{s_L, s_H\}$.

B. *Multiproduct Cournot Oligopoly*

The stronger conclusion of convexity allows us to extend the work of Section IV. Specifically, we may consider a set of m firms that must simultaneously choose supplies of n different quality-differentiated products. Adopting the upgrades approach, and given decreasing returns to quality, a symmetric equilibrium of the multiproduct Cournot oligopoly game corresponds to a profile of industry upgrade outputs $\{Z_{is}^*\}$ satisfying $\mathrm{MR}_s(Z_{is}^*) = \Delta c_i/\Delta q_i$, where marginal revenue is defined for an oligopolist as described just above.[61] Put simply, a multiproduct equilibrium reduces to a collection of single-product equilibria in each upgrade market.[62] From Proposition 8, profits are convex in s for each upgrade market, and hence total profits are also convex; the conclusion of Proposition 6 continues to hold under the more general setting of quantity competition.

Proposition 8 also states that the m-firm marginal revenue curve rotates clockwise following an increase in s. This property is all we need to ensure that the distribution of qualities offered by the industry rotates, just as is the case under monopoly (Proposition 7).

PROPOSITION 9: *For elasticity-ordered demand curves, the associated distributions of qualities $\{1 - Z_{js}^*\}$ offered by a Cournot industry are ordered by a sequence of rotations. An increase in dispersion results in an expansion of the product line: if a product is offered in positive supply for some s, it will continue to be offered for higher s. Furthermore, if the rotation quantity $z_s^\dagger$ is increasing, multiproduct oligopoly profits are convex in s and maximized by $s \in \{s_L, s_H\}$.*

From this analysis, we conclude that while our results cannot hold in full generality in oligopoly settings, there are natural specifications under which they do in fact persist. Our results on changing product lines are particularly robust. For instance, when a family of distributions is variance-ordered and exhibits decreasing marginal revenue (Definition 2), the associated m-firm marginal revenue curves are ordered by a sequence of rotations. It follows

[61] In fact, the analysis of Johnson and Myatt (forthcoming) reveals that such an equilibrium is unique.

[62] Precisely, this is true because of the maintained assumptions that marginal revenue is decreasing and $\Delta c_i/\Delta q_i$ is increasing. See Johnson and Myatt (2006) for a comprehensive analysis of multiproduct Cournot oligopoly, and Johnson and Myatt (2003) for the consequences of a failure of decreasing marginal revenue.

that the extension of Proposition 7 continues to hold in this setting, and in any others for which marginal revenue curves are rotation-ordered.

C. *Preferences for Dispersion and Industry Concentration*

We now ask how an increase in competition changes a firm's preference for dispersion. Recall that we describe a monopolist as a niche player whenever $z_s^* < z_s^\dagger$, or equivalently when an increase in dispersion increases her profits; similarly she is a mass-market supplier when increased dispersion hurts her profits. Naturally, we can extend the same terminology to an oligopoly. We impose the conditions of Proposition 9 so that profits are U-shaped. Given that this is so, there will exist some dispersion parameter $\tilde{s} \equiv \arg\min \pi(s)$ that minimizes a firm's profits: this yields a mass-market industry when $s < \tilde{s}$ and a niche industry when $s > \tilde{s}$.

PROPOSITION 10: *Suppose that the family of demand curves is elasticity-ordered, and the rotation quantity $z_s^\dagger$ is increasing. Then $\tilde{s} \equiv \arg\min_{s \in [s_L, s_H]} \pi(s)$ is increasing in m.*

This says that if a firm dislikes any local increase in dispersion, that firm will continue to dislike increased dispersion when the number of competitors rises. Heuristically, the reason is that heightened competition expands total industry output, which makes it more likely that the industry's marginal consumer is "below average." This suggests, for instance, that firms in more competitive industries are more likely to be hurt by consumer learning (whether generated exogenously or endogenously) or some other increase in dispersion, such as wider income inequality or more idiosyncratic consumer tastes.

VI. Concluding Remarks

We have proposed a framework for analyzing both exogenous and endogenous transformations of the demand facing firms. Our approach is based on the observation that changes in demand frequently correspond to changes in the dispersion of the underlying willingness to pay of consumers, which lead to a rotation of the demand curve. We investigated numerous applications of our framework, including product design decisions, advertising and marketing activities, and product-line choices of multiproduct firms. The optimal advertising and product design depend, in turn, on the desire to adopt either a mass-market or niche posture. A niche position is complemented by high levels of dispersion, and a mass-market position by low dispersion. We also suggested a new taxonomy of advertising, distinguishing between hype, which shifts demand, and real information, which rotates demand. While our framework is broadly applicable, it is also straightforward. The simple economics of demand rotations can help us to understand many phenomena from diverse areas of application.

Appendix A: Incorporating Multiple Rotation Points

Quasi-Convexity of Profits. Section IC claimed that Definition 1 may be extended to cope with multiple rotation points. To do this, consider a distribution function $F_s(\theta)$ indexed by s. We index J (where J is odd) rotation points by j, and order them so that $\theta_s^{\dagger(j)} < \theta_s^{\dagger(j+1)}$. For convenience of notation, we define $\theta_s^{\dagger(0)} = -\infty$ and $\theta_s^{\dagger(J+1)} = +\infty$. An increase in s will correspond to an increase in dispersion if, for $1 \leq j \leq J + 1$,

$$\theta_s^{\dagger(j-1)} < \theta < \theta_s^{\dagger(j)} \Rightarrow \frac{\partial F_s(\theta)}{\partial \theta} \begin{cases} >0 & j \text{ odd} \\ <0 & j \text{ even} \end{cases}.$$

Furthermore, $\partial F_s(\theta_s^{\dagger(j)})/\partial s = 0$ for $1 \leq j \leq J$. Increasing s rotates $F_s(\theta)$ clockwise around odd-numbered rotation points, and counterclockwise around even-numbered rotation points.[63]

We now explain the extensions to Lemma 1 and Proposition 1. Suppose that each (odd-numbered) clockwise rotation point is decreasing in s and each (even-numbered) counterclockwise rotation point is increasing. Take any

[63] When J is odd, $F_s(\theta)$ moves up and down in the extreme lower and upper tails, respectively. This is required, for example, if the family of distributions is to be ordered by second-order stochastic dominance. This maintains the spirit of our earlier analysis. Note that $J = 1$ yields the specification of Definition 1.

fixed θ. If $F_s(\theta)$ is not quasi-concave in s, then there exist $s' < s''$ such that $F_{s'}(\theta) = F_{s''}(\theta) > F_{\tilde{s}}(\theta)$ for some $\tilde{s} \in (s', s'')$. This implies that $F_s(\theta)$ is strictly locally decreasing in s for some $\hat{s} \in (s', \tilde{s})$, which, in turn, means that for some even j, $\theta_{\hat{s}}^{\dagger(j-1)} < \theta < \theta_{\hat{s}}^{\dagger(j)}$. But by hypothesis, $\theta_s^{\dagger(j-1)}$ is decreasing and $\theta_s^{\dagger(j)}$ is increasing, meaning that $F_s(\theta)$ is strictly decreasing over $(\hat{s}, s'')$, and hence over $(\tilde{s}, s'')$. This contradicts $F_{s''}(\theta) > F_{\tilde{s}}(\theta)$, and we conclude $F_s(\theta)$ is quasi-concave in s for each θ. By the argument following Lemma 1 in the text, quasi-concavity of $F_s(\theta)$ implies quasi-convexity of profits.

Multiple Submarkets. Multiple rotation points can arise when demand is drawn from different submarkets. Specifically, suppose that valuations in a submarket $i \in \{1, \ldots, n\}$ with mass α_i follow a distribution $F^i_{s_i}(\theta)$, with a dispersion parameter s_i. Overall demand is determined by the mixture $F(\theta) = \sum_{i=1}^n \alpha_i F^i_{s_i}(\theta)$. If s_i orders $F^i_{s_i}(\theta)$ by a sequence of rotations, then it also orders $F(\theta)$ in the same way, and if $F^i_{s_i}(\theta)$ is quasi-concave in s_i for each θ, then a monopolist's profits will be quasi-convex in s_i. The monopolist will prefer extreme choices for s_i in each of the n submarkets. However, the different dispersion parameters rotate the mixture distribution around different rotation points. Furthermore, which extreme is preferred may differ by submarket.[64]

While advertising tactics may vary by submarket, it is, perhaps, more difficult for product designs to differ. Ideally, a monopolist would like to offer a plain-vanilla design to submarket i and an idiosyncratic design to submarket j. If forced to offer the same design, however, then things are more complex. Effectively, the monopolist is constrained to set $s_i = s_j = s$. Thus, an increase in s (more idiosyncratic design, for instance) may simultaneously rotate $F(\theta)$ (and hence the corresponding demand curve) around multiple rotation points.

We illustrate with a specific example. Suppose that $G(\cdot)$ is a symmetric unimodal distribution with full support, density $g(\cdot)$, zero mean, and unit variance. With two submarkets, we build a mixture $F_s(\theta) = \alpha G((\theta - \mu_1)/s) + (1 - \alpha)G((\theta - \mu_2)/s)$, so that the expected consumer valuation is $\alpha\mu_1 + (1 - \alpha)\mu_2$. If $\theta_s^\dagger$ is a rotation point, then

$$\frac{\partial F_s(\theta_s^\dagger)}{\partial s} = \frac{1}{s^2}\left[\alpha(\mu_1 - \theta_s^\dagger)g\left(\frac{\theta_s^\dagger - \mu_1}{s}\right) + (1 - \alpha)(\mu_2 - \theta_s^\dagger)g\left(\frac{\theta_s^\dagger - \mu_2}{s}\right)\right] = 0.$$

This can hold only when $\mu_1 > \theta_s^\dagger > \mu_2$, and hence any rotation point must lie between the two submarket means. When this first-order condition is satisfied, the second derivative is

$$\frac{\partial^2 F_s(\theta_s^\dagger)}{\partial s^2} = \frac{1}{s^4}\left[\alpha(\mu_1 - \theta_s^\dagger)^2 g'\left(\frac{\theta_s^\dagger - \mu_1}{s}\right) + (1 - \alpha)(\mu_2 - \theta_s^\dagger)^2 g'\left(\frac{\theta_s^\dagger - \mu_2}{s}\right)\right].$$

Since $g(\cdot)$ is unimodal and $\mu_1 > \theta > \mu_2$, the first $g'(\cdot)$ term is positive and the second negative.[65] Hence, this second derivative can take either sign, and $F_s(\theta)$ may be quasi-convex rather than quasi-concave. We can obtain sharper results by specifying $G(\cdot)$ to be a standard normal distribution, so that $g(x) = (2\pi)^{-1}e^{-(x^2/2)}$. The standard normal density has the property $g'(x) = -xg(x)$. This enables us to combine the second-order condition on $F_s(\theta)$ with the first-order condition, and find that $\partial^2 F_s(\theta_s^\dagger)/\partial s^2 \leq 0$ if and only if $\theta_s^\dagger \geq (\mu_1 + \mu_2)/2$. Thus, whether $\theta_s^\dagger$ is increasing or decreasing in s depends upon whether it lies above or below the halfway point between the two submarket means. Even if $\theta_s^\dagger$ is increasing in s, profits do not necessarily reach a local maximum at an interior value of s. In fact, $\theta_s^\dagger$ must rise more quickly than the monopoly price if the quasi-convexity of profits is to fail.

[64] Suppose that it is optimal to set a price p satisfying $\theta_{s_i}^\dagger > p > \theta_{s_j}^\dagger$ for two submarkets i and j. Heuristically, the marginal consumer is "below average" in submarket i and "above average" in submarket j. Thus, the monopolist may wish to use pure promotional hype in i, while supplying detailed real information to submarket j.

[65] Since $g(\cdot)$ is unimodal, $g(\theta - \mu_i)$ is increasing if and only if $\theta < \mu_i$.

Appendix B: Omitted Proofs

Propositions 1, 3, 4, 6, and 9 follow from arguments presented in the main text.

PROOF OF LEMMA 1:

The argument in the text demonstrates that if $z_s^\dagger$ is weakly increasing in s, then $P_s(z)$ is quasi-convex in s. Now suppose that $z_s^\dagger$ is not weakly increasing. Given this, we may find $s' < s''$ and a z such that $z_{s'}^\dagger > z > z_{s''}^\dagger$, and

$$\frac{\partial P_{s'}(z)}{\partial s} > 0 > \frac{\partial P_{s''}(z)}{\partial s},$$

so that $P_s(z)$ is first increasing and then decreasing in s: quasi-convexity fails. Similar logic applies to $F_s(\theta)$. For a variance-ordered family, $P_s(z) = \mu(s) + \sigma(s)P(z)$. Fixing s, let us choose z so that $\partial P_s(z)/\partial s = \mu'(s) + \sigma'(s)P(z) = 0$. Then for quasi-convexity we need

$$\frac{\partial^2 P_s(z)}{\partial s^2} = \mu''(s) + \sigma''(s)P(z)$$

$$= \mu''(s) - \frac{\mu'(s)\sigma''(s)}{\sigma'(s)} \geq 0 \Leftrightarrow \frac{d}{ds}\left[\frac{\mu'(s)}{\sigma'(s)}\right] \geq 0.$$

For the elasticity-ordered case, $z_s^\dagger = \exp(\mu'(s))$ is increasing if and only if $\mu''(s) \geq 0$.

PROOF OF LEMMA 2:

Setting $\sigma(s) = s$ without loss, $P_s(z) = \mu(s) + sP(z)$, and hence marginal revenue satisfies $\mathrm{MR}_s(z) = \mu(s) + s[P(z) + zP'(z)] = \mu(s) + s\,\mathrm{MR}(z)$, where $\mathrm{MR}(z)$ is the marginal revenue associated with an inverse-demand curve $P(z)$. Observe that

$$\frac{\partial\,\mathrm{MR}_s(z)}{\partial s} = \mu'(s) + \mathrm{MR}(z)$$

$$\Rightarrow \frac{\partial^2 \mathrm{MR}_s(z)}{\partial z \partial s} = \frac{\partial\,\mathrm{MR}(z)}{\partial z} = \frac{1}{s}\frac{\partial\,\mathrm{MR}_s(z)}{\partial z} < 0,$$

since marginal revenue is decreasing in z. Thus, if $\mathrm{MR}_s(z)$ is decreasing in s, then it is decreasing for all larger z. Hence, for some $z_s^\ddagger \in [0, 1]$, marginal revenue is increasing in s for $z < z_s^\ddagger$ and decreasing in s for $z > z_s^\ddagger$: the family of marginal revenue curves is ordered by a sequence of rotations. For the elasticity-ordered case, see Proposition 8.

PROOF OF PROPOSITION 2:

The main proposition follows from the argument in the text. For the variance-ordered case, the logic used in the proof of Lemma 1 ensures that $z_s^\ddagger$ is increasing in s. $P_s(z)$ is convex in s when $\mu(s)$ is convex, and hence $\pi(s) = \max_{z\in[0,1]}\{zP_s(z) - C(z)\}$ is the maximum of convex functions and so is convex.

PROOF OF PROPOSITION 5:

A consumer updates his prior $\omega \sim N(\mu, \kappa^2)$ following the observation of the signal $x|\omega \sim N(\omega, \xi^2)$. Standard calculations confirm that his posterior satisfies

$$\omega|x \sim N\left(\frac{\kappa^2 x + \xi^2\mu}{\kappa^2 + \xi^2}, \frac{\kappa^2\xi^2}{\kappa^2 + \xi^2}\right)$$

$$\text{or equivalently } N\left(\frac{\mu + \psi\kappa^2 x}{1 + \psi\kappa^2}, \frac{\kappa^2}{1 + \psi\kappa^2}\right).$$

Given CARA preferences, the consumer's willingness to pay will be the certainty equivalent $\mathrm{E}[\omega|x] - \lambda\,\mathrm{var}[\omega|x]/2$. Substituting in for the mean and variance yields the expression for $\theta(x)$ given in the text. We may now turn to the proposition itself. In a slight abuse of notation, let us write σ^2 for the variance of θ. Following some algebraic manipulation,

$$P_{\kappa^2,\psi}(z) = \mu - \frac{\lambda\kappa^2}{2} + \frac{\lambda s^2}{2} + \sigma P(z)$$

$$\text{where } s = \sqrt{\frac{\psi\kappa^4}{1 + \psi\kappa^2}}.$$

Let us fix κ^2. By inspection, $P_{\kappa^2,\psi}(z)$ is convex in s, and s in turn is increasing in ψ. Thus, $P_{\kappa^2,\psi}(z)$ is quasi-convex in ψ. Hence profits are quasi-convex in ψ. Furthermore,

$$\frac{\partial P_{\kappa^2,\psi}(z_s^\dagger)}{\partial \psi} = 0 \Leftrightarrow \lambda s = -P(z_\psi^\dagger) \Leftrightarrow z_\psi^\dagger = \Phi(\lambda s),$$

which, upon substitution of s, yields $z_\psi^\dagger$ in footnote 49. Next, let us fix ψ. Differentiating,

$$\frac{\partial P_{\kappa^2,\psi}(z)}{\partial \kappa^2} = -\frac{\lambda}{2} + [\lambda s + P(z)]\frac{\partial s}{\partial \kappa^2} = 0$$

$$\Leftrightarrow z = z^{\dagger}_{\kappa^2} = \Phi\left(-\frac{\lambda}{2(\partial s/\partial \kappa^2)} + \lambda s\right).$$

To assess the quasi-convexity of profits in κ^2, we apply Proposition 1. Examining $z^{\dagger}_{\kappa^2}$,

$$\frac{\partial z^{\dagger}_{\kappa^2}}{\partial \kappa^2} \geq 0 \Leftrightarrow \frac{\lambda}{2(\partial s/\partial \kappa^2)^2}\frac{\partial^2 s}{\partial(\kappa^2)^2} + \lambda\frac{\partial s}{\partial \kappa^2}$$

$$\geq 0 \Leftrightarrow 2\left[\frac{\partial s}{\partial \kappa^2}\right]^3 + \frac{\partial^2 s}{\partial(\kappa^2)^2} \geq 0.$$

To verify this last inequality, differentiate s with respect to κ^2 to obtain

$$\frac{\partial s}{\partial \kappa^2} = \frac{(2 + \psi\kappa^2)\sqrt{\psi}}{2(1 + \psi\kappa^2)^{3/2}} \quad \text{and}$$

$$\frac{\partial^2 s}{\partial(\kappa^2)^2} = -\frac{(4 + \psi\kappa^2)\psi^{3/2}}{4(1 + \psi\kappa^2)^{5/2}}.$$

The desired inequality becomes

$$\frac{(2 + \psi\kappa^2)^3\psi^{3/2}}{4(1 + \psi\kappa^2)^{9/2}} \geq \frac{(4 + \psi\kappa^2)\psi^{3/2}}{4(1 + \psi\kappa^2)^{5/2}} \Leftrightarrow$$

$$(2 + \psi\kappa^2)^3 \geq (4 + \psi\kappa^2)(1 + \psi\kappa^2)^2 \Leftrightarrow 4 \geq 0,$$

which establishes the claimed quasi-convexity of profits. Next, we observe that

$$\frac{\partial P_{\kappa^2,\psi}(z)}{\partial \kappa^2} = -\frac{\lambda}{2} + \frac{\partial P_{\kappa^2,\psi}(z)}{\partial \psi} \times \frac{\partial s/\partial \kappa^2}{\partial s/\partial \psi}.$$

Inspection of this expression yields the remaining claims of the proposition.

PROOF OF PROPOSITION 7:

The first claim follows from the argument in the text. For the final claim, suppose $Z^*_{is} > Z^*_{(i+1)s}$, so that the complete product i is in positive supply. If for some higher s', product i is not being supplied, then either $Z^*_{is'} = Z^*_{(i+1)s'} = 1$, or $Z^*_{is'} = Z^*_{(i+1)s'} = 0$; we can rule out interior possibilities because $\Delta c_i/\Delta q_i$ is strictly increasing and marginal revenue is decreasing. Now, it cannot be that $Z^*_{is'} = Z^*_{(i+1)s'} = 1$, since $\text{MR}_s(1)$ is decreasing in s, so that the fact $Z^*_{(i+1)s} < 1$ implies $Z^*_{(i+1)s'} < 1$. Also, it can't be that $Z^*_{is'} = Z^*_{(i+1)s'} = 0$, since $\text{MR}_s(0) = P_s(0)$ is increasing in s, so that $Z^*_{is} > 0$ implies that $Z^*_{is'} > 0$.

PROOF OF PROPOSITION 8:

As noted in the text, for demand with constant elasticity $1/s$,

$$\text{MR}_s(z) = P_s(z)\left(1 - \frac{s}{m}\right)$$

$$\Rightarrow \frac{\partial\, \text{MR}_s(z)}{\partial s} = \frac{m - s}{m}\frac{\partial P_s(z)}{\partial s} - \frac{P_s(z)}{m} \gtrless 0$$

$$\Leftrightarrow \frac{1}{m - s} \lessgtr \frac{\partial \log P_s(z)}{\partial s} = \mu'(s) - \log z$$

$$\Leftrightarrow z \lessgtr \exp\left(\mu'(s) - \frac{1}{m - s}\right) \equiv z^{\ddagger}_s,$$

so that the marginal revenue faced by a firm rotates clockwise around $z^{\ddagger}_s$. Now,

$$\text{MR}_s(z) = c \Leftrightarrow P_s(z) = e^{\mu(s)}z^{-s} = \frac{mc}{m - s}$$

$$\Leftrightarrow z = z^*_s \equiv \left[\frac{(m - s)e^{\mu(s)}}{mc}\right]^{1/s}.$$

Substitution leads to the expression for $\pi(s)$. Evaluating the derivative of profits,

$$\frac{d\pi(s)}{ds} = \frac{z^*_s}{m}\left[\frac{\partial P_s(z^*)}{\partial s}\right.$$

$$\left. + \frac{(m - 1)P'(z^*_s)}{m}\frac{z^*_s}{s}\left[\frac{\partial \log P_s(z^*_s)}{\partial s} - \frac{1}{m - s}\right]\right]$$

$$= \frac{z^*_s P_s(z^*_s)}{m^2}\left[\frac{\partial \log P_s(z^*)}{\partial s} + \frac{m - 1}{m - s}\right]$$

$$= \frac{e^{\mu(s)}[z^*_s]^{1-s}}{m^2}\left[\mu'(s) - \log z^*_s + \frac{m - 1}{m - s}\right].$$

Hence, taking the second derivative to evaluate convexity, we obtain

$$\frac{d^2\pi(s)}{ds^2} = \frac{e^{\mu(s)}[z_s^*]^{1-s}}{m^2}\left[\mu''(s) - \frac{d\log z_s^*}{ds} + \frac{m-1}{(m-s)^2}\right] + \frac{e^{\mu(s)}[z_s^*]^{1-s}}{m^2}\left[\mu'(s) - \log z_s^* + \frac{m-1}{m-s}\right] \times \left[\mu'(s) - \log z_s^* + (1-s)\frac{d\log z_s^*}{ds}\right].$$

For the desired result, we need this to be positive when $\mu''(s) \geq 0$. By inspection, it is sufficient to show that the desired inequality holds when $\mu''(s) \geq 0$. We seek to evaluate:

$$\left[\frac{m-1}{(m-s)^2} - \frac{d\log z_s^*}{ds}\right] + \left[\mu'(s) - \log z_s^* + \frac{m-1}{m-s}\right] \times \left[\mu'(s) - \log z_s^* + (1-s)\frac{d\log z_s^*}{ds}\right] \geq 0.$$

It is straightforward to confirm that

$$\frac{d\log z_s^*}{ds} = \frac{1}{s}\left[\mu'(s) - \log z_s^* - \frac{1}{m-s}\right],$$

and hence on substitution we obtain the criterion

$$[\mu'(s) - \log z_s^*]^2 - \frac{2(1-s)}{m-s}[\mu'(s) - \log z_s^*] + \frac{(m-s) + (2s-1)(m-1)}{(m-s)^2} \geq 0.$$

This is most difficult to satisfy when $\mu'(s)$ minimizes the left-hand side, which happens when

$$\mu'(s) - \log z_s^* = \frac{1-s}{m-s}.$$

Evaluating at this point, we obtain the inequality

$$\frac{(m-s) + (2s-1)(m-1) - (1-s)^2}{(m-s)^2} \geq 0$$

$$\Leftrightarrow s \leq 2m - 1.$$

This final inequality holds since (following Definition 3) for an elasticity ordered family we restrict to elastic demand satisfying $s < 1$. Finally, we consider quasi-convexity of Cournot output, or equivalently quasi-convexity of $\log z_s^*$. Observe that

$$\frac{d^2\log z_s^*}{ds^2} = \frac{1}{s}\left[\mu''(s) - 2\frac{d\log z_s^*}{ds} - \frac{1}{(m-s)^2}\right].$$

By inspection, when $d\log z_s^*/ds = 0$, this expression is positive if and only if $\mu''(s) > (m - s)^{-2}$, as required. This completes the proof.

PROOF OF PROPOSITION 10:

For ease of exposition we consider only interior values of $\tilde{s} \in (s_L, s_H)$. Building upon expressions obtained in the proof of Proposition 8, observe that:

$$\frac{d\pi(s)}{ds} \gtrless 0 \Leftrightarrow \mu'(s) \gtrless \log z_s^* - \frac{m-1}{m-s}.$$

Differentiating the right-hand side of this expression with respect to m,

$$\frac{d}{dm}\left\{\log z_s^* - \frac{m-1}{m-s}\right\} = \frac{s(m-1)}{m(m-s)^2} > 0.$$

Evaluated at $s = \tilde{s}$, we know that $d\pi(s)/ds = 0$. An increase in m results in $d\pi(s)/ds < 0$, and hence an increase in $\tilde{s}$, the desired result.

REFERENCES

Ackerberg, Daniel A. "Empirically Distinguishing Informative and Prestige Effects of Advertising." *RAND Journal of Economics*, 2001, *32*(2), pp. 316–33.

Ackerberg, Daniel A. "Advertising, Learning, and Consumer Choice in Experience Good

Markets: An Empirical Examination." *International Economic Review*, 2003, *44*(3), pp. 1007–40.

Adams, William J. and Yellen, Janet L. "Commodity Bundling and the Burden of Monopoly." *Quarterly Journal of Economics*, 1976, *90*(3), pp. 475–98.

Aislabie, Colin J. and Tisdell, Clement A. "Profit Maximization and Marketing Strategies: Demand Rotation and Social Influences [a Note on Price-Advertising Reaction Function]." *Managerial and Decision Economics*, 1988, *9*(1), pp. 77–82.

Anderson, Simon P. and Renault, Régis. "Advertising Content." *American Economic Review*, 2006, *96*(1), pp. 93–113.

Arrow, Kenneth J. "The Theory of Risk Aversion," in Kenneth J. Arrow, ed., *Essays in the theory of risk bearing*. Amsterdam: North-Holland, 1970, pp. 90–120.

Atkinson, Anthony B. "On the Measurement of Inequality." *Journal of Economic Theory*, 1970, *2*(3), pp. 244–63.

Bagwell, Kyle. "The Economic Analysis of Advertising," in C. Mark Armstrong and Robert H. Porter, eds., *Handbook of industrial organization*. Vol. 3. Amsterdam: Elsevier North-Holland (forthcoming).

Baron, David P. and Myerson, Roger B. "Regulating a Monopolist with Unknown Costs." *Econometrica*, 1982, *50*(4), pp. 911–30.

Becker, Gary S. and Murphy, Kevin M. "A Simple Theory of Advertising as a Good or Bad." *Quarterly Journal of Economics*, 1993, *108*(4), pp. 941–64.

Biehl, Andrew R. "Durable-Goods Monopoly with Stochastic Values." *RAND Journal of Economics*, 2001, *32*(3), pp. 565–77.

Blackwell, David. "Equivalent Comparisons of Experiments." *The Annals of Mathematical Statistics*, 1953, *24*(2), pp. 265–72.

Butters, Gerard R. "Equilibrium Distributions of Sales and Advertising Prices." *Review of Economic Studies*, 1977, *44*(3), pp. 465–91.

Che, Yeon-Koo. "Customer Return Policies for Experience Goods." *Journal of Industrial Economics*, 1996, *44*(1), pp. 17–24.

Coes, Donald V. "Firm Output and Changes in Uncertainty." *American Economic Review*, 1977 (*Papers and Proceedings*), *67*(2), pp. 249–51.

Comanor, William S. and Wilson, Thomas A. "Advertising Market Structure and Performance." *Review of Economics and Statistics*, 1967, *49*(4), pp. 423–40.

Comanor, William S. and Wilson, Thomas A. "On Advertising and Profitability." *Review of Economics and Statistics*, 1971, *53*(4), pp. 408–10.

Courty, Pascal. "Ticket Pricing under Demand Uncertainty." *Journal of Law and Economics*, 2003, *46*(2), pp. 627–52.

Crémer, Jacques; Khalil, Fahad and Rochet, Jean-Charles. "Contracts and Productive Information Gathering." *Games and Economic Behavior*, 1998, *25*(2), pp. 174–93.

Dai, Chifeng; Lewis, Tracy R. and Lopomo, Giuseppe. "Delegating Management to Experts." *RAND Journal of Economics* (forthcoming).

DeGraba, Patrick. "Buying Frenzies and Seller-Induced Excess Demand." *RAND Journal of Economics*, 1995, *26*(2), pp. 331–42.

Diamond, Peter A. and Stiglitz, Joseph E. "Increases in Risk and in Risk Aversion." *Journal of Economic Theory*, 1974, *8*(3), pp. 337–60.

Dixit, Avinash and Norman, Victor. "Advertising and Welfare." *Bell Journal of Economics*, 1978, *9*(1), pp. 1–17.

Dorfman, Robert and Steiner, Peter O. "Optimal Advertising and Optimal Quality." *American Economic Review*, 1954, *44*(5), pp. 826–36.

Dowling, Graham R. "Information Content in US and Australian Television Advertising." *Journal of Marketing*, 1980, *44*(4), pp. 34–37.

Erdem, Tülin and Keane, Michael P. "Decision-Making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets." *Marketing Science*, 1996, *15*(1), pp. 1–20.

Grossman, Gene M. and Shapiro, Carl. "Informative Advertising with Differentiated Products." *Review of Economic Studies*, 1984, *51*(1), pp. 63–81.

Grossman, Sanford J. and Hart, Oliver D. "Implicit Contracts under Asymmetric Information." *Quarterly Journal of Economics*, 1983, *98*(3), pp. 123–56.

Hammond, John S. "Simplifying the Choice between Uncertain Prospects where Preference Is Non-Linear." *Management Science*, 1974, *20*(7), pp. 1047–72.

Holmström, Bengt R. "Managerial Incentive Schemes—a Dynamic Perspective," in Lars Wahlbeck, ed., *Essays in economics and management in honour of Lars Wahlbeck*. Helsinki: Svenska Handelshogskolan, 1982, pp. 209–30.

Jewitt, Ian. "Risk Aversion and the Choice between Risky Prospects: The Preservation of Comparative Statics Results." *Review of Economic Studies*, 1987, *54*(1), pp. 73–85.

Jewitt, Ian. "Choosing between Risky Prospects: The Characterization of Comparative Statics Results, and Location Independent Risk." *Management Science*, 1989, *35*(1), pp. 60–70.

Johnson, Justin P. "Entry-Level Products with Consumer Learning." *Contributions to Economic Analysis & Policy*, 2005, *4*(1), art. 9.

Johnson, Justin P. and Myatt, David P. "Multiproduct Quality Competition: Fighting Brands and Product Line Pruning." *American Economic Review*, 2003, *93*(3), pp. 748–74.

Johnson, Justin P. and Myatt, David P. "Multiproduct Cournot Oligopoly." *RAND Journal of Economics*, 2006.

Judd, Kenneth L. and Riordan, Michael H. "Price and Quality in New Product Monopoly." *Review of Economic Studies*, 1994, *61*(4), pp. 773–89.

Karlin, Samuel. "Pólya Type Distributions, II." *Annals of Mathematical Statistics*, 1957, *28*(2), pp. 281–308.

Karlin, Samuel. *Total positivity*. Stanford: Stanford University Press, 1968.

Karlin, Samuel and Proschan, Frank. "Pólya Type Distributions of Convolutions." *Annals of Mathematical Statistics*, 1960, *31*(3), pp. 721–36.

Laczniak, Gene R. "Information Content in Print Advertising." *Journalism Quarterly*, 1979, *56*(2), pp. 324–27.

Lancaster, Kelvin J. "A New Approach to Consumer Theory." *Journal of Political Economy*, 1966, *74*(2), pp. 132–57.

Lancaster, Kelvin J. *Consumer demand: A new approach*. New York: Columbia University Press, 1971.

Lancaster, Kelvin J. "Socially Optimal Product Differentiation." *American Economic Review*, 1975, *65*(4), pp. 567–85.

Lehmann, Erich L. "Ordered Families of Distributions." *Annals of Mathematical Statistics*, 1955, *26*(3), pp. 399–419.

Leland, Hayne E. "Theory of the Firm Facing Uncertain Demand." *American Economic Review*, 1972, *62*(3), pp. 278–91.

Lewis, Tracy R. and Sappington, David E. M. "All-or-Nothing Information Control." *Economics Letters*, 1991, *37*(2), pp. 111–13.

Lewis, Tracy R. and Sappington, David E. M. "Supplying Information to Facilitate Price Discrimination." *International Economic Review*, 1994, *35*(2), pp. 309–27.

Machina, Mark J. and Pratt, John W. "Increasing Risk: Some Direct Constructions." *Journal of Risk and Uncertainty*, 1997, *14*(2), pp. 103–27.

Madden, Charles S.; Caballero, Majorie J. and Matsukubo, Shinya. "Analysis of Information Content in US and Japanese Magazine Advertising." *Journal of Advertising*, 1986, *15*(3), pp. 38–45.

McAfee, R. Preston; McMillan, John and Whinston, Michael D. "Multiproduct Monopoly, Commodity Bundling, and Correlation of Values." *Quarterly Journal of Economics*, 1989, *104*(2), pp. 371–83.

Meyer, Margaret A. "The Dynamics of Learning with Team Production: Implications for Task Assignment." *Quarterly Journal of Economics*, 1994, *109*(4), pp. 1157–84.

Milgrom, Paul R. and Roberts, John. "Price and Advertising Signals of Product Quality." *Journal of Political Economy*, 1986, *94*(4), pp. 796–821.

Milgrom, Paul R. and Weber, Robert J. "A Theory of Auctions and Competitive Bidding." *Econometrica*, 1982, *50*(5), pp. 1089–1122.

Mussa, Michael and Rosen, Sherwin. "Monopoly and Product Quality." *Journal of Economic Theory*, 1978, *18*(2), pp. 301–17.

Nelson, Philip. "Advertising as Information." *Journal of Political Economy*, 1974, *82*(4), pp. 729–54.

Ottaviani, Marco and Prat, Andrea. "The Value of Public Information in Monopoly." *Econometrica*, 2001, *69*(6), pp. 1673–83.

Persico, Nicola. "Information Acquisition in Auctions." *Econometrica*, 2000, *68*(1), pp. 135–48.

Pratt, John W. "Risk Aversion in the Small and the Large." *Econometrica*, 1964, *32*(1), pp. 122–36.

Renforth, William and Raveed, Sion. "Consumer Information Cues in Television Advertising:

A Cross Country Analysis." *Journal of the Academy of Marketing Science*, 1983, *11*(3), pp. 216–25.

Resnik, Alan J. and Stern, Bruce L. "An Analysis of Information Content in Television Advertising." *Journal of Marketing*, 1977, *41*(1), pp. 50–53.

Rothschild, Michael and Stiglitz, Joseph E. "Increasing Risk I: A Definition." *Journal of Economic Theory*, 1970, *2*(3), pp. 225–43.

Rothschild, Michael and Stiglitz, Joseph E. "Increasing Risk II: Its Economic Consequences." *Journal of Economic Theory*, 1971, *3*(1), pp. 66–84.

Schmalensee, Richard. "A Model of Advertising and Product Quality." *Journal of Political Economy*, 1978, *86*(3), pp. 485–503.

Schmalensee, Richard. "Commodity Bundling by Single-Product Monopolies." *Journal of Law and Economics*, 1982, *25*(1), pp. 67–71.

Schmalensee, Richard. "Gaussian Demand and Commodity Bundling." *Journal of Business*, 1984, *57*(1), pp. S211–30.

Senstrup, Preben. "Information Content in TV Advertising: Consumer Policy Implications of the Growing Supply of TV Advertising in Europe." *Journal of Consumer Policy*, 1985, *8*(3), pp. 239–65.

Shapiro, Carl. "Advertising and Welfare: Comment." *Bell Journal of Economics*, 1980, *11*(2), pp. 749–52.

Shapiro, Carl. "Consumer Information, Product Quality, and Seller Reputation." *Bell Journal of Economics*, 1982, *13*(1), pp. 20–35.

Sobel, Joel. "Information Control in the Principal-Agent Problem." *International Economic Review*, 1993, *34*(2), pp. 259–69.

Stern, Bruce L. and Resnik, Alan J. "Information Content in Television Advertising: A Replication and Extension." *Journal of Advertising Research*, 1991, *31*(3), pp. 36–46.

Stern, Bruce L.; Resnik, Alan J. and Grubb, Edward L. "Information Content in Television Advertising: A Further Analysis," in Barnett A. Greenberg and Danny N. Bellenger, eds., *Contemporary marketing thought—1977: Educators' proceedings*. Chicago: American Marketing Association, 1977, pp. 358–61.

Stigler, George J. "The Economics of Information." *Journal of Political Economy*, 1961, *69*(3), pp. 213–25.

Stigler, George J. "A Note on Block Booking," in George J. Stigler, ed., *The organization of industry*. Homewood, IL: Richard D. Irwin, 1968, pp. 165–70.

Stigler, George J. and Becker, Gary S. "De Gustibus Non Est Disputandum." *American Economic Review*, 1977, *67*(2), pp. 76–90.

Taylor, David B. "The Information Content of Women's Magazine Advertising in the UK." *European Journal of Marketing*, 1983, *17*(5), pp. 28–32.

Vives, Xavier. *Oligopoly pricing: Old ideas and new tools*. Cambridge, MA: MIT Press, 1999.

Weinberger, Marc G. and Spotts, Harlan E. "A Situational View of Information Content in TV Advertising in the US and UK." *Journal of Marketing*, 1989, *53*(1), pp. 89–94.