

Search, Design, and Market Structure[†]

By HESKI BAR-ISAAC, GUILLERMO CARUANA, AND VICENTE CUÑAT*

The Internet has dramatically changed the nature of demand and competition. In many industries, new production and search technologies have altered the pattern of sales and market structure to the benefit of fringe goods in the long tail.¹ Another common view, however, informed by search models, is that cheaper access to information leads to very competitive markets with low price dispersion and a few, high-quality, superstar products.

While the long-tail and superstar phenomena are often portrayed as opposing each other, reality has proven more subtle, as the case of the book publishing industry suggests. Easier access to information on available titles has increased the number of specialized books that cater to ever more specific audiences. In recent years, while the overall size of the market has remained roughly constant, the composition of book sales has changed. From 2002 to 2007, the number of new titles and editions grew at the astonishing rate of approximately 10 percent each year; indeed, the number of new titles in 2007 alone surpassed the total published throughout the 1970s. At the same time, the startling and unprecedented success of books such as the *Harry Potter* series or *The Da Vinci Code* seems to suggest the presence of superstar effects.²

*Bar-Isaac: Department of Economics, Stern School of Business, New York University, 44 West Fourth Street 7-73, New York, NY 10012 (e-mail: heski@nyu.edu); Caruana: Casado del Alisal 5, 28014 Madrid, Spain (e-mail: caruana@cemfi.es); Cuñat: Room A450, Department of Finance, London School of Economics, Houghton Street, London WC2A 2AE, UK (e-mail: v.cunat@lse.ac.uk). Previous versions have circulated as “Costly Search and Design.” We thank for their helpful comments three anonymous referees, Simon Anderson, Michael Baye, Antonio Cabrales, Juanjo Ganuza, Avi Goldfarb, Maarten Janssen, Justin Johnson, George Mailath, Ram Rao, Eric Rasmussen, Michael Rauh, Régis Renault, Andrew Rhodes, Konstantinos Serfes, Silvia Sonderegger, Yaron Yehezkel, and excellent seminar participants at many seminars. Financial support from the NET Institute (<http://www.NETinst.org>) is gratefully acknowledged. Guillermo Caruana acknowledges the financial support of the Spanish Ministry of Science and Innovation through grant ECO2008-00801.

[†]To view additional materials, visit the article page at <http://dx.doi.org/10.1257/aer.102.2.1140>.

¹The term “long tail” refers to the well-documented and dramatic increase in the market share for goods in the tail of the sales distribution (that is, with relatively low sales). The phrase was coined in an article in *Wired* (Anderson 2004) and was later expanded and developed in Anderson (2006, 2009). Anderson focuses mostly on books, music, and film sales and rentals, but argues that the phenomenon is much broader. See Brynjolfsson, Hu, and Smith (2006) for references to theoretical and empirical academic work on long tails. In particular, Brynjolfsson, Hu, and Smith (2003) and Brynjolfsson, Hu, and Simester (2007) find evidence of long tails for online bookstores and an online retail women’s catalog, respectively.

²The facts in this paragraph are collected from a variety of sources. According to the US Census data, the overall size of the book market has remained relatively unchanged (revenues fell by five percent). The data on new book titles from 2002–2007 are available from Bowker (the exclusive US ISBN and SAN Agency), which also shows annual growth of six percent per year when excluding on-demand and short-run titles (Bowker Industry Statistics 2009). The data for the 1970s come from Greco (2004). Sales figures for yearly best-sellers, including the *Harry Potter* series and the *Da Vinci Code*, are found across different issues of *Publishers Weekly*.

The coexistence of the long-tail and superstar phenomena extends to other industries. Jeff Bewkes, the Chairman and CEO of Time Warner, points out that in media industries

*[a]udiences are at once fragmenting into niches and consolidating around blockbusters. Of course, media consumption has not risen much over the years, so something must be losing out. That something is the almost but not quite popular content that occupies the middle ground between blockbusters and niches. The stuff that people used to watch or listen to largely because there was little else on is increasingly being ignored.*³

In this paper, we allow for a richer choice of firm strategies than the search literature has typically considered. Specifically, firms choose the “design” of their products in addition to price. Our notion of design is broad and can accommodate not only physical design, but also marketing and information disclosure.⁴ We are, therefore, able to address how designs adapt as search costs fall and to consider the equilibrium effects on market structure, prices, and consumer surplus. In particular, our analysis leads naturally to long-tail and superstar effects arising simultaneously, and to prices and industry profits that are nonmonotonic in search costs.⁵

Formally, our notion of design choice builds on a recent and growing literature, notably Johnson and Myatt (2006), with an important antecedent in Lewis and Sappington (1994).⁶ While this literature has focused on design choices by monopolists, our paper extends this analysis to a competitive environment. To do so, we introduce product design, along the lines of Johnson and Myatt (2006), into a search model (Wolinsky 1986; Bakos 1997; or Anderson and Renault 1999). In particular, firms choose designs ranging from broad market designs that are inoffensive to all consumers to more niche or quirky designs that consumers either love or loathe. Meanwhile, each consumer searches among firms, paying a small cost to obtain a price quote from an additional firm and to learn about the extent to which that firm’s product suits his tastes.

The model allows us to address the impact of search engines, the Internet, and communication and information technologies more generally, by interpreting these as a fall in search costs. Our central results are on the equilibrium market outcomes. First, we show that more-advantaged firms choose broader designs, while disadvantaged firms prefer nichier ones. More importantly, allowing for an endog-

³ “A World of Hits.” 2009. *The Economist* 393(8659): 79.

⁴ Within the book publishing example, it would accommodate both publisher decisions to focus on particular topics or types of manuscripts, and marketing decisions such as making sample chapters available online.

⁵ There is a small related literature that considers firms that vary design in response to falling search costs. Larson (forthcoming) studies horizontal differentiation in a model of sequential search with a particular emphasis on welfare considerations in what can be viewed as a special case of our model. Kuksov (2004) presents a duopoly model in which consumers know the varieties available (but not their location) prior to search, and different designs come with different costs associated. Anderson and Renault (2009) also consider duopoly and, in a result similar to one in this paper, show that the low-quality firm has the greater incentives to release information on horizontal characteristics; Cachon, Terwiesch, and Xu (2008) and Watson (2009) focus specifically on multiproduct firms’ choices of product range. Our model allows for a wide range of designs and a much more general demand specification. It also focuses on different issues. For example, these papers do not consider sales distributions explicitly, and so do not address long-tail and superstar effects.

⁶ More recently, Bar-Isaac, Caruana, and Cufi a (2010, forthcoming) put more emphasis on consumers’ information-gathering decisions and highlight that these are co-determined with the firm’s pricing, design, and marketing strategies in equilibrium.

enous choice of product design reveals that lower search costs have an indirect effect on prices and profits through changes to the offered designs. Lower search costs induce more firms to choose niche designs, leading consumers to visit more firms. Consequently, prices and profits can be nonmonotonic in search costs. In particular, we argue that profits increase as search costs fall only when ex ante differences between firms are relatively small. Instead, if firms are sufficiently vertically differentiated, then reducing search costs intensifies price competition and leads to lower industry profits.

Reduced search costs and endogenous designs also have interesting effects on sales distributions. Lower search costs allow consumers to search longer and find “better” firms. This leads to a superstar effect, where better firms are even more successful.⁷ Furthermore, consumers are also more likely to buy better-suited products. For simplicity, in our model the overall size of the market stays constant. Consistent with Jeff Bewkes’s comments quoted above, however, the stars and the tail can both increase market share at the expense of middling firms.⁸ These firms, facing a more competitive environment, switch to niche designs with lower sales and higher markups, thereby releasing additional buyers. Some of these buyers will end up purchasing from other niche firms, which boost their sales and allow for a long-tail effect. Overall, our model shows how a reduction in search costs can simultaneously explain both superstar and long-tail effects.

Finally, by imposing functional forms, we obtain a simple and intuitive condition on the existence of long-tail effects. These should not arise if there is considerable vertical heterogeneity in the industry, while, at the same time, the range of possible designs is relatively narrow. This is consistent with the evidence of Goldmanis et al. (2010) on bookstores, new-auto dealers, and travel agencies, where different firms are often selling identical goods. Instead, if vertical heterogeneity is limited, and moving from a broad to a niche design has a relatively large effect, long-tail effects are likely to arise. This is consistent with the results of Brynjolfsson, Hu, and Smith (2003) who show an increase in product variety in the book-publishing industry, an industry in which the scope for niche targeting is very large.

I. Model

There is a continuum of risk-neutral firms and consumers of measure 1 and m , respectively. Each firm i produces a single product. Each consumer l has tastes described by a conditional utility function (net of any search costs) of the form

$$(1) \quad u_{li}(p_i) = v_i + \varepsilon_{li} - p_i$$

⁷Maris Goldmanis et al. (2010), who consider better firms to be those with lower costs, find the superstar effect both theoretically and empirically in their study of bookstores, travel agencies, and new-auto dealers.

⁸Elberse and Oberholzer-Gee (2007) present empirical evidence on simultaneous long-tail and superstar effects from the video rental industry. Similar evidence, albeit through the channel of recommender systems, appears in Oestreicher-Singer and Sundararajan (forthcoming) on books and Tucker and Zhang (2011) on wedding service vendors. Hervás-Drane (2010) provides further references and a model that contrasts two different channels (sequential search and ex ante recommendations) through which the Internet might generate superstar and long-tail effects.

if she buys product i at price p_i . The term v_i captures a natural advantage of firm i . A higher v_i can be thought of as a better innate vertical quality, but it can also be interpreted as a lower marginal production cost.⁹ Meanwhile, $\varepsilon_{li} \sim F_i$ is a match value between consumer l and product i . It captures idiosyncratic consumer preferences for certain products over others. We assume that realizations of ε_{li} are independent across firms and individuals.¹⁰

A consumer incurs a search cost c to learn the price p_i , the quality v_i , and the match value ε_{li} for the product offered by any particular firm i . Consumers search sequentially. The utility of a consumer l is given by

$$(2) \quad u_{lk}(p_k) - kc$$

if she buys product k at price p_k at the k th firm she visits. From now on, and for simplicity, we will omit the firm and consumer subscripts, unless there is ambiguity.

Firms cannot affect v , the exogenous quality of the good, which is distributed according to some continuously differentiable distribution $H(v)$ with support $[\underline{v}, \bar{v}]$. In Section V, we analyze the case where the distribution is degenerate, so that, ex ante, all firms are identical.

We introduce strategic design choice by assuming that the firm can affect the distribution of the match-specific component of consumer tastes F_s by picking a design $s \in S = [B, N]$. That is, designs range from most-broad (B) to most-niche (N). A design s leads to ε_{li} distributed according to $F_s(\cdot)$ with support on some bounded interval $(\underline{\theta}_s, \bar{\theta}_s)$, and with a log-concave density $f_s(\cdot)$ that is positive everywhere.¹¹ Regardless of design and intrinsic quality, the firm produces goods at a marginal cost of 0.¹²

The strategy for each firm, therefore, consists of a choice of price p and a product design $s \in S$. We suppose that there are no costs associated with choosing different designs s .¹³

We follow Johnson and Myatt (2006) in assuming that different product designs induce demand rotations. Formally, there is a family of rotation points θ_s^\dagger such that $\partial F_s(\theta)/\partial s < 0$ for $\theta > \theta_s^\dagger$ and $\partial F_s(\theta)/\partial s > 0$ for $\theta < \theta_s^\dagger$; further, θ_s^\dagger is decreasing

⁹In the latter interpretation, $v_i (< 0)$ represents the negative of marginal costs, and p_i should be understood as the absolute markup above the marginal cost. The price that a consumer faces would be $p_i - v_i$. Given that v_i enters the consumer's utility linearly, derivations and results would not change at all.

¹⁰Taking these realizations to be independent, while consistent with previous literature on search (Wolinsky 1986 and Anderson and Renault 1999), is not without loss of generality. It does not allow firms to target specific niches. That is, there is no spatial notion of differentiation or product positioning. Given that we assume a continuum of firms and no ability for consumers to determine location in advance, however, this assumption may be more reasonable.

¹¹See Bagnoli and Bergstrom (2005) for a broad discussion of log-concavity and functions that do and do not satisfy this condition. The assumption of log-concavity ensures that the failure rate $f_s(\theta)/(1 - F_s(\theta))$ is monotonic and, so, guarantees existence of a profit-maximizing monopolist price, which is continuous and increasing in constant marginal costs.

¹²Assuming constant marginal costs and no fixed costs simplifies the analysis considerably, though it can be relaxed in a similar fashion to Section IIB of Johnson and Myatt (2006). Within a framework of constant marginal costs, setting them to zero is without loss of generality. As already mentioned, differences in marginal costs play an identical role to differences in v .

¹³Relaxing this assumption would affect the results. It is not obvious, however, how costs should vary with design. In particular, when interpreted as information provision, assuming that different designs come at the same cost is reasonable. More generally, targeting a specific audience may be costly, but so is producing a fully compatible, wide-reaching product. Thus, we remain agnostic and focus fully on the demand-induced effect.

in s . The concept of a demand rotation is a formal approach to the notion that some designs lead to a wider spread in consumer valuations than others. In particular, a higher value of s should be interpreted as a “quirkier” product that appeals more to certain consumers and less to others; the bounds on s correspond to the most-broad (B) and the most-niche (N) designs. Alternatively, design can be interpreted as a marketing decision on how much information to release; here, a higher value of s corresponds to more information that shifts priors (up or down) and leads to more-dispersed valuations. This formalization of design choices is general enough to accommodate a wide range of concepts of product design. One of the contributions of Johnson and Myatt (2006) is to show that natural models of physical product design and information release provide microfoundations for such demand rotations.

The natural equilibrium concept to consider is weak perfect Bayesian Nash equilibrium (wPBNE). As is standard in the search literature, we impose that consumers keep the same (passive) beliefs about the distribution of future prices and designs on and off the equilibrium path.¹⁴ This restriction implies that a consumer’s search and purchase behavior can be described by a threshold rule U : she buys the current product if the utility she obtains, $u_{ii}(p_i)$, is more than or equal to U , and continues searching otherwise.¹⁵ One additional advantage to this notation is that U also represents the consumer surplus from participating in the market. To summarize, consumers choose a threshold U , and each v firm chooses a strategy (p, s) that consists of a price, p , and design, s .¹⁶ It can be shown that the set of wPBNE with passive beliefs corresponds with the set of Nash equilibria of a simplified static version of the model in which, simultaneously, consumers choose a threshold rule U as above, and firms choose their prices and designs. For simplicity, equilibrium in the rest of the paper refers to Nash equilibrium of this simplified game.

II. Equilibrium

A. Consumer Behavior

Suppose that a consumer expects each firm of type v to choose strategy (p_v, s_v) .¹⁷ Consider a consumer who can stop searching and obtain utility u . If the consumer, instead, samples an additional firm of type v , she will prefer the new product if $v + \varepsilon - p_v > u$. In this case, the additional utility obtained is $v + \varepsilon - (u + p_v)$, and so the expected incremental utility from searching at one more firm that is expected to have design s_v and price p_v and to be of quality v is

$$(3) \quad E_\varepsilon(\max\{v + \varepsilon - p_v - u, 0\}) = \int_{u+p_v}^{\infty} (v + \varepsilon - p_v - u) f_{s_v}(\varepsilon) d\varepsilon.$$

¹⁴This is a reasonable restriction. Given that there are infinite firms that do not coordinate their actions, a consumer cannot infer anything about what may be available in the future at other stores from today’s observed realization.

¹⁵The optimality of a threshold rule follows from the analysis of McCall (1970) of costly search in a stationary i.i.d. environment. Note, though, that different tie-breaking rules can be used when $u_{ii}(p_i) = U$. This is a source of multiplicity of equilibria. Since ties are zero-probability events, such multiplicity has no impact on any of the analysis in the paper.

¹⁶More broadly, we can allow firms to mix, so that each firm chooses an element $\sigma_v \in \Delta(\mathbb{R} \times [B, N])$.

¹⁷With a continuum of firm types and no atoms in the distribution, it is without loss of generality to assume that each type of firm chooses a pure strategy in design and price.

It is worth visiting one more firm if and only if the expected value of the visit is worth more than the cost, where the final expectation is taken over v (with an implicit firm strategy for both price and design); that is, as long as $E_v[E_\varepsilon(\max\{v + \varepsilon - p_v - u, 0\})] \geq c$, or, equivalently, if $u < U$ where U is implicitly defined by

$$(4) \quad \int_{-\infty}^{\infty} \left(\int_{U+p_v-v}^{\infty} (v + \varepsilon - p_v - U) f_{s_v}(\varepsilon) d\varepsilon \right) h(v) dv = c.$$

There is, at most, one solution to (4) since the left-hand side is strictly decreasing in U (as the integrand is decreasing in U and the lower limit of the inner integral is increasing in U). For large enough c , there is no feasible positive U that satisfies (4): no consumer would ever continue searching, and firms would have full monopoly power (as in Diamond 1971). In other words, the consumer initiates search if and only if $U \geq 0$.

B. Firm Profit Maximization

Consumers who visit a firm of type v buy as long as they receive a match ε such that $v + \varepsilon - p > U$ and, thus, purchase with probability $1 - F_s(p + U - v)$.

We define ρ as the expected probability that a consumer who visits a random firm buys from that firm; this is exogenous from the perspective of firm v . This definition allows us to calculate the demand for a given firm v . The expected number of consumers who visit it in the first round is m . A further $m(1 - \rho)$ consumers visit firm v in the second round after an unsuccessful visit to some other firm, a further $m(1 - \rho)^2$ visit in the third round, and so on. We can, therefore, write demand for firm v that chooses a design s and price p as

$$(5) \quad \frac{m}{\rho} (1 - F_s(p + U - v)),$$

and its profits as

$$(6) \quad \Pi = \frac{m}{\rho} p (1 - F_s(p + U - v)).$$

It is useful to define $p_{vs}(U)$ as firm v 's profit-maximizing price when the consumer's threshold is U and the design strategy is s . This price is implicitly determined by

$$(7) \quad p_{vs}(U) = \frac{1 - F_s(p_{vs}(U) + U - v)}{f_s(p_{vs}(U) + U - v)}.$$

Our first result, a consequence of the log-concavity assumption, ensures that p_{vs} is well-defined and behaves in a way that is intuitive: first, higher-quality firms charge higher prices; and second, firms charge lower prices if they face pickier consumers—that is, if customers have higher values of U . Note that all proofs in the paper appear in the appendices.

LEMMA 1: *The profit-maximizing price, $p_{vs}(U)$, associated with a design s is uniquely defined. It is continuously decreasing in the consumers' reservation threshold U ,*

and continuously increasing in the firm's quality v . Further, $p_{vs}(U) + U$ is continuously increasing in U , and $p_{vs}(U) - v$ is continuously decreasing in v .

Substituting for $p_{vs}(U)$, in (6), the firm's problem is to maximize the resulting expression with respect to its remaining strategic variable s . Note that neither the optimal price nor the optimal design choice depends on m or ρ , as these are just constant factors in profits.¹⁸

Johnson and Myatt (2006) show that, when designs are rotation-ordered and all designs cost the same, monopoly profits are quasi-convex in design. Thus, a monopolist would always choose an extremal design. In our competitive environment, since every firm has a local monopoly power, and the resulting residual demand is a truncation of the original one, it is still rotation-ordered, and, as a consequence, the same result applies.

PROPOSITION 1: *Firms choose extremal designs. That is, every firm chooses either the most-niche ($s = N$) or the most-broad ($s = B$) design.*

To gain some intuition for this result, first consider the case in which the optimal price at a given design s is below the point of rotation, so that the profit-maximizing quantity is greater than the quantity at the point of rotation $1 - F_s(\theta_s^\dagger)$. Then, decreasing s (and, thus, "flattening" out demand) will lead to a greater quantity being sold even if the price is kept fixed. Therefore, decreasing s must lead to higher profits. A similar argument applies when the optimal price is above the point of rotation.¹⁹

The result of Proposition 1 arises from the specific demand-rotation assumptions that we make on designs, following Johnson and Myatt (2006). The central intuition of the paper, however—that there are equilibrium effects through changing designs as search costs fall—applies much more generally. The demand-rotation framework simply provides an analytically tractable environment to explore such equilibrium effects.

In particular, Proposition 1 allows us to restrict attention to equilibrium strategies in which firm v chooses either a broad design (p_{vB}, B) or a niche one (p_{vN}, N), where p_{vB} and p_{vN} are defined by (7) for $s = B, N$, respectively.

Next, define $V(U)$ as the solution to

$$(8) \quad p_{vB}(U)(1 - F_B(p_{vB}(U) + U - V)) \\ = p_{vN}(U)(1 - F_N(p_{vN}(U) + U - V)).$$

If $V(U)$ lies in the feasible range $[\underline{v}, \bar{v}]$, then $V(U)$ captures the firm that is indifferent between choosing the broad or the niche strategy. If $V(U)$ falls outside this

¹⁸This highlights that search costs play a qualitatively different role from that of scale effects. This is a central point of Wolinsky (1986), and is discussed by Anderson and Renault (1999).

¹⁹For an alternative intuition, consider in the price-quantity space all feasible demand curves that arise from different designs. The rotation-ordering condition ensures that the upper envelope of these demands is traced out by the most-niche and most-broad designs.

range, with some abuse of notation, we redefine it to be the appropriate extreme of the range.²⁰ This definition allows us to characterize firm behavior.²¹

PROPOSITION 2: *Given a consumer search rule, U , there is a threshold type of firm $V(U)$, as characterized by (8), such that all firms with lower quality than this threshold, $v < V(U)$, choose a niche design, and all firms with $v > V(U)$ choose a broad one. Moreover, $V(U)$ is continuously increasing in U ; that is, as consumers search more intensively, more firms choose niche designs.*

Proposition 2 shows that the more severe the competition a firm faces (either because consumers are pickier and require more utility in order to purchase, or because the firm faces a disadvantage as compared to other firms), the more likely it is to choose a niche strategy. Loosely, the intuition here is that a firm in a disadvantageous position needs the consumer to “love” the good in order to buy it. The chance that this happens increases with a design that leads to dispersed valuations—a niche design. Instead, a high-value firm can appeal to many consumers by adopting the broad strategy and, thereby, can minimize the chance that a well-disposed consumer observes that a product is such a bad match that she would prefer not to purchase.

This result is economically rich and appealing. When interpreting v as relating to marginal costs, it states that low-cost firms try to attract a broad market, while high-cost firms, which must charge higher prices to be profitable, target niches. The result is also relevant when interpreting v as quality. For example, consider five-star hotels competing in a city. Although they are in the same category, they differ in an important dimension: location. Our model predicts that well-located hotels (center of the city, close to the airport or other facilities) are more likely to deliver standard services. Meanwhile, those with less-desirable locations are more likely to be specialized—for example, boutique hotels with distinctive styling or those catering to specific groups, such as customers with pets.

C. Equilibrium Summary

Given the analysis above, we can express an equilibrium as a pair (U, V) , where U summarizes the search and purchase behavior of consumers, and V determines which firms choose the broad or the niche strategy. These two values have to satisfy the following conditions. First, rewriting (4), consumers optimize their behavior when

$$(9) \quad c = \int_{-\infty}^V \left(\int_{p_{vN}(U)+U-v}^{\infty} (v + \varepsilon - p_{vN}(U) - U) f_N(\varepsilon) d\varepsilon \right) h(v) dv + \int_V^{\infty} \left(\int_{p_{vB}(U)+U-v}^{\infty} (v + \varepsilon - p_{vB}(U) - U) f_B(\varepsilon) d\varepsilon \right) h(v) dv.$$

²⁰Mathematically, we redefine V to be $\max\{v, \min\{\bar{v}, \cdot\}\}$ of the solution to (8).

²¹Note that firms that make no sales are indifferent about the design they choose. It is convenient, however, for the statement of results (while having no effect on equilibrium transactions) to assume that such firms respect the design choices implied by Proposition 2.

Second, as explained above, firms' maximizing behavior is summarized by the indifference of V , as in (8). Third, associated with broad and niche designs are profit-maximizing prices $p_{vB}(U)$ and $p_{vN}(U)$, as determined in (7). Finally, it must be worthwhile for a consumer to initiate search; that is, $U \geq 0$.

Given this characterization, we can compute the expected probability that a consumer buys when she visits a random firm. This is given by

$$(10) \quad \rho(U, V) \equiv \int_{-\infty}^V (1 - F_N(p_{vN}(U) + U - v)) h(v) dv \\ + \int_V^{\infty} (1 - F_B(p_{vB}(U) + U - v)) h(v) dv.$$

Note that there always exist equilibria where consumers prefer not to search, and (in order to sustain this) firms charge sufficiently high prices. When search costs are sufficiently high—specifically for $c > c_0$, where c_0 is defined in the proof of Proposition 3—this is the unique equilibrium. For lower search costs, we ignore this type of equilibria and fully focus on those that involve active search by consumers. Turning to the firms, these equilibria may involve either all firms choosing a niche (or broad) design, or some firms choosing a niche while others are choosing a broad design. The former case is straightforward to characterize. We summarize its results in the Proposition below. Meanwhile, the latter—where different firms choose different kinds of designs—is the interesting case, and the one to which we devote the rest of the paper.

PROPOSITION 3: *Let U_B denote the consumer stopping rule that makes the lowest-quality firm indifferent between choosing a niche and a broad design, and U_N , the consumer stopping rule that makes the highest-quality firm indifferent between choosing a niche and a broad design.*

There exist thresholds on search costs c_0 , c_N , and c_B such that there exists an equilibrium where all firms choose the broad design if and only if $U_B > 0$ and $c \in [c_B, c_0)$; and there exists an equilibrium where all firms choose the niche design if and only if $c < \underline{c}$, where

$$\underline{c} = \begin{cases} c_N & \text{if } U_N > 0 \\ c_0 & \text{if } U_N \leq 0 \end{cases}.$$

III. Comparative Statics: The Effect of Lower Search Costs

Despite having abstracted away from equilibria with no search, there is still scope for multiplicity in this model. Here, we argue that some equilibria are better behaved than others in terms of their stability properties. We can consider how the demand side, through changes to U , and the supply side, through changes to V , respond to perturbations from an equilibrium. Given the continuum of firms and consumers, the game has an infinite-dimensional space of strategy profiles, and so in the online Appendix B.1, we propose a simplified dynamic system that captures the interplay between V and U . The steady states of this system are the

Nash equilibria of our model. Thus, we refer to our Nash equilibria as stable if they are asymptotically stable within that dynamic system and, from now on, focus our attention on stable equilibria.

Multiple stable equilibria may still exist.²² Nevertheless, in the online Appendix B1, we show that an equilibrium is stable if and only if $U(\cdot)$ has a slope less than 1 at the equilibrium. This observation is key, as it allows us to obtain the same comparative statics at any stable equilibrium and, thus, understand how the industry changes in response to a reduction in the search costs:

PROPOSITION 4: *At any stable equilibrium, decreasing the search cost c raises both consumer surplus (higher U) and the fraction of niche firms (higher V).*

As highlighted at the start of the paper, there has been much recent discussion of the long tail of the Internet. Proposition 4 provides a theoretical result that speaks to the issue by demonstrating that, for stable equilibria, lower search costs bring more niche firms.

This does not necessarily imply that niche products sell more or are more profitable. There are competing effects that arise from consumers being more picky (that is, having higher U). There is a direct effect that leads firms to drop prices and sell less per consumer visit. There is also a countervailing effect: More consumers will visit any given firm (i.e., ρ is lower), not only because consumers are more picky, but also because more intermediate firms opt for a low-sales, high-markup strategy. The overall effect on sales and profits is, therefore, ambiguous. Indeed, in Sections V and VI, we show that either effect can dominate. First, we formally define our notions of long-tail and superstar effects.

IV. Long Tails and Superstars

DEFINITION 1: *We say that a superstar effect is present if the firm with the highest sales captures an increasing market share as search costs fall.*

DEFINITION 2: *We say that a long-tail effect is present if the firm with the lowest sales captures an increasing market share as search costs fall.*

Our definitions of long-tail and superstar effects may seem somewhat extreme in focusing only on one firm; but in this model, because of continuity, if the extreme firm behaves in a certain way, so do adjacent ones. Thus, our definitions imply a mass of firms at the head or tail of the sales distribution gaining market share.

Below, we study distributional changes when different designs coexist in equilibrium. We start, however, by stressing the importance of design heterogeneity for the emergence of long-tail effects. In Proposition B.1 in the online Appendix, we show that, for a wide class of distribution functions (including those with nondecreasing density), if all firms choose the same design, there are always superstar effects, but

²² Still, there are many cases, as exemplified in Sections V and VI, where a unique equilibrium exists.

never long-tail effects. This suggests that the documented long-tail effect cannot be solely a consequence of a fall in the cost of search. If firms continued delivering the same type of products, we would expect to see low-quality firms losing market share. It is through a change towards more-niche designs that the long tail arises.²³ Note, also, that holding design constant, firm profits decrease as search costs decrease,²⁴ which appears to be at odds with the rise of new firms on the Internet.

While it is plausible that the Internet has reduced fixed costs of entry, we demonstrate that when firms' designs are strategic choices, the long-tail effect arises naturally, and that as search costs fall, firm profits can increase, potentially leading to new-firm entry. We show these effects clearly by adding some further structure to the model. In Section V, we assume ex ante symmetry of all firms, and in Section VI, we allow for heterogeneous types and consider uniform distributions.

V. Homogeneous Firms

Here, we simplify our model by assuming that all firms are ex ante identical. The purpose is to show that the key ingredient to obtain our results is the endogenous design choice, and not firm heterogeneity per se, as in most papers in the long-tail literature.

Without loss of generality, we assume that $v = 0$ for all firms. To simplify notation, we drop the v subscripts throughout this section. Given that all firms are now alike, we need to consider the possibility of mixed-strategy equilibria. In particular, we denote λ as the proportion of firms that choose a niche rather than a broad design. Analogous to the characterization of Section IIC, equilibria can be summarized by (U, λ) , and conditions (8)–(10) can be written as

$$(11) \quad c = \lambda \int_{p_N(U)+U}^{\infty} (\varepsilon - p_N(U) - U) f_N(\varepsilon) d\varepsilon \\ + (1-\lambda) \int_{p_B(U)+U}^{\infty} (\varepsilon - p_B(U) - U) f_B(\varepsilon) d\varepsilon.$$

$$(12) \quad \lambda \in \arg \max \left\{ \lambda p_B(U) (1 - F_B(U) + U) \right. \\ \left. + (1 - \lambda) p_N(U) (1 - F_N(p_N(U) + U)) \right\}.$$

$$(13) \quad \rho(U) = \lambda (1 - F_N(p_N(U) + U)) + (1 - \lambda) (1 - F_B(p_B(U) + U)).$$

Note that the characterization of prices, given by (7), and the consumer's participation constraint ($U \geq 0$) still apply.

²³ An alternative and plausible explanation is that the nature of the search technology has changed, and, in particular, the quality of "targeted" search has improved. We consider this channel complementary to the one we study. They would reinforce each other and exacerbate the shift towards niche designs and the phenomena studied in this paper.

²⁴ This follows from Lemma 1 and Proposition 4 and does not require any restriction on $f_s(\cdot)$.

Given that all firms are identical, U_B and U_N , as defined in Proposition 3, coincide. We write $\bar{U} = U_B = U_N$. For $U > \bar{U}$, therefore, all firms prefer a niche design, whereas, for $U < \bar{U}$, all firms prefer a broad design. It is only at $U = \bar{U}$ that firms might mix. A mixed-strategy equilibrium can exist, however, over a wide range of search costs. This is immediate, by noting that at $U = \bar{U}$, expression (11) can be rewritten as

$$(14) \quad c = \lambda c_N + (1 - \lambda) c_B,$$

where c_B and c_N are the search-cost thresholds introduced in Proposition 3 and formally characterized in Equations (22) and (25) in the Appendix. Note that, in this case, c_B and c_N have interpretations as the expected consumer surplus from visiting a broad or, respectively, a niche firm.

If $c_N < c_B$, then the mixed-strategy equilibrium exactly fills the gap between the regions where all-broad and all-niche equilibria exist, and λ is linear and decreasing in c . If $c_N > c_B$, then in this region, there are, in principle, three equilibria: one all-broad, one mixed, and one all-niche. Note, however, that the mixed equilibrium in this case is unstable. Thus, for $c \in (c_B, c_N)$, only two pure equilibria remain. Finally, if $c_N = c_B$, the mixed-strategy equilibrium has no mass. It is easy to find examples of each of these three cases.²⁵

When search costs are sufficiently high (low) so that all firms choose the broad (niche) design, comparative statics are standard: As c falls, consumer surplus rises, profits decrease, and the sales of each firm stay constant at m .²⁶ Instead, when search costs are intermediate, results are more interesting.

PROPOSITION 5: *When all firms are ex ante identical, in the region (c_N, c_B) where both designs are offered, then as search costs fall: (i) consumer surplus is constant at \bar{U} ; (ii) there are more niche firms (λ increases); (iii) consumers search more (ρ decreases); (iv) every firm's profits increase; and (v) both long-tail and superstar effects arise.*

First, note that although a fall in search costs represents a direct benefit to consumers, this gain is exactly offset by the negative impact from searching more (ρ decreases) and from the increased preponderance of niche firms that provide less surplus in expectation ($c_B > c_N$). Next, since the consumer threshold is constant throughout the region, a firm's expected profit per consumer visit does not change. Given that there are more consumer visits (ρ decreases), however, profits increase.²⁷

Finally, we turn to market structure. Consistent with the long-tail evidence, we observe that as search costs fall, each niche firm sells more. In addition, there are

²⁵ When demands are linear, the ratio of consumer surplus to firm profits for a monopolist is constant at $1/2$. Therefore, two firms facing linear demands (regardless of their slopes) that earn the same profits must generate the same consumer surplus and, so, $c_N = c_B$. Instead, if demand is concave, the ratio of consumer surplus to profits is always lower than it would be in the linear case. Consequently, if F_N is linear and F_B is concave, then $c_N > c_B$ and multiplicity arises, whereas in the opposite case, with F_N concave and F_B linear, a unique equilibrium exists.

²⁶ See Proposition A2 in the online Appendix for a formal statement and proof.

²⁷ Although the probability of making a sale for any given visit stays constant for any given type of firm, this is consistent with more consumers visiting since the composition of firms changes. There are more niche firms as c falls, and niche firms sell less than broad firms.




FIGURE 1. DISTRIBUTION OF SALES AT DIFFERENT SEARCH COSTS

more niche firms and, since the total volume of sales is constant, it follows that the niche firms account for a greater proportion of overall sales. Note, also, that superstar effects are present. The “top” firm chooses a broad design and sells more as c goes down. The tail is niche throughout and also sells more as c goes down. The middle region, where the mix of broad and niche is changing, is the one that loses sales to both the head and the tail of the sales distribution. This is illustrated in Figure 1 below.

When search costs are low enough or high enough, all firms choose the same design and all of them sell m . Thus, as Figure 1 shows, sales are nonmonotonic in search costs. Profits are also nonmonotonic: they decrease in search costs when these are low or high, but increase in search costs in the intermediate region (as shown in Propositions 5 and B2 in the online Appendix).

VI. Uniformly Distributed Quality and Linear Demands

We once again consider heterogeneous firms, but impose further structure that allows us to derive additional analytic results. These highlight that the results of Section V extend naturally to more-general settings. We analyze the case where both the distribution of firm quality $v \sim U[L, H]$, and the distributions of consumer-product matches $F_s(\cdot)$ are uniform, leading to linear demand functions. In particular, the niche and broad product designs are, respectively, $\varepsilon \sim U[\underline{\theta}_N, \bar{\theta}_N]$ and $\varepsilon \sim U[\underline{\theta}_B, \bar{\theta}_B]$. We impose that $\underline{\theta}_N < \underline{\theta}_B$ and $\bar{\theta}_N > \bar{\theta}_B$. This ensures that these distributions represent demand rotations.

In this environment, we derive the following results.

PROPOSITION 6: *Under the conditions above, when all firms are active, then:*
 (i) *there is a unique equilibrium (U, V) for each search cost c . When different firms*

choose different design strategies, then as the search cost decreases: (ii) consumer surplus U increases; (iii) there are more niche firms (V increases); (iv) profits of the highest- and lowest-quality firms increase if and only if $\bar{\theta}_N - \bar{\theta}_B > H - L$; (v) the superstar effect arises; and (vi) the long-tail effect can, but need not, arise; a sufficient condition for it to arise is $\bar{\theta}_N - \bar{\theta}_B > H - L$.

It is worth highlighting that as the extent of vertical differentiation among firms diminishes (that is, as $H - L \rightarrow 0$), we recover the results of Proposition 5. This shows that the case of homogeneous firms is not knife-edged. In other words, one can regard the vertical component as a device to purify the mixed strategies that arise in the homogeneous case.

But, more importantly, Proposition 6 shows that while superstar effects are robust, the long-tail effect and the comparative statics of profitability depend on how $\bar{\theta}_N - \bar{\theta}_B$ compares to $H - L$. While $H - L$ captures the extent of vertical differentiation, $\bar{\theta}_N - \bar{\theta}_B$ captures horizontal differentiation: it measures the importance of changing from broad to niche designs in terms of the dispersion of match values. Thus, when designs have a relatively greater impact on horizontal differentiation, the competition-softening effect of firms switching to niche designs more than compensates for the intensified vertical competition that arises as search costs fall.

Note that if firms' types are very dispersed, then a low-quality firm must be forced out of the market when search costs are sufficiently low; following our definition, trivially, in such circumstances, long-tail effects cannot arise. Proposition 6, therefore, focuses on parameter ranges where all firms remain active even for low values of c .

We illustrate some results of Proposition 6 in the case that $\bar{\theta}_N - \bar{\theta}_B > H - L$. Specifically, consider $f_N(x) = \frac{1}{4}$ on $[-12, 4]$, $f_B(x) = \frac{1}{6}$ on $[-3, 3]$ and $h(x) = \frac{4}{3}$ on $[0, \frac{3}{4}]$. We use this example to demonstrate the nonmonotonicity of prices and profits, and the superstar and long-tail effects.

Figure 2 illustrates how prices vary with search costs for a particular firm (at $v = 0.5$). As one would anticipate, in general, prices increase with search costs. When the firm changes design from niche to broad, however, prices drop substantially, leading to prices that are nonmonotonic in search costs. The price pattern for other values of v is qualitatively the same.

Next, consider how firm profits for the worst firm, the best firm, and the industry's average profits, vary with search costs, as illustrated in Figure 3.²⁸ Note the two points where the derivative is discontinuous. These are the search-cost thresholds at which the equilibrium changes from all-niche or all-broad to one in which there is a mix of designs: below $c_N = 0.038$, all firms are niche, but as search costs increase, the high-quality firms gradually start switching to a broad design. At $c_B = 0.08$ and beyond, all firms choose a broad design. Figure 3 illustrates that profits may be nonmonotonic. The intuition is the, by now, familiar one that as search costs fall in the intermediate region, more firms choose a niche design and, thereby, soften competition.

Finally, we consider sales distributions. Figure 4 is the analogue of Figure 1 and plots the distribution of sales for two different search costs. Naturally, higher-quality firms sell more than low-quality firms do, regardless of the search costs. Comparing

²⁸ Since there is a mass 1 of firms, the graph of average profits also represents total industry profits.




FIGURE 2. PRICE AGAINST SEARCH COST




FIGURE 3. PROFITS AGAINST SEARCH COSTS




FIGURE 4. SALES AGAINST QUALITY (v) AT $c = 0.05$ AND $c = 0.06$

sales at different search costs, both the highest- and lowest-quality firms sell more at the lower level of search costs, illustrating that superstar and long-tail effects arise simultaneously. These are also illustrated at intermediate levels of search costs (where there is dispersion in designs offered) in Figure 5, which plots sales against search costs for the best and worst firms.




FIGURE 5. SALES AGAINST SEARCH COST FOR BEST AND WORST FIRMS

VII. Conclusions

There has been considerable focus on the Internet's influence on the kind of products offered and the distribution of their sales. In particular, academic and popular commentators have highlighted both long-tail and superstar effects for various industries (including publishing, media, and travel destinations, among others). This paper presents a simple and tractable model integrating consumer search and firms' strategic product-design choices that is useful to analyze these phenomena.

We show that, in equilibrium, different product designs coexist. More-advantaged firms prefer broad-market strategies, seeking a very broad design and choosing a relatively low price, while less-advantaged firms take a niche strategy with quirky products priced high to take advantage of the (relatively few) consumers who are well-matched to the product. Such design diversity arises even when all firms are ex ante homogeneous.

Prices and profits can be nonmonotonic in consumer search costs. There is an intuitive rationale for this: as search costs fall, and as long as the product designs remain unchanged, prices fall. At ever lower prices, however, the broad-market strategy becomes less appealing to firms, some of which adopt a niche strategy, charging a high price to the (few) consumers who are well-matched for their product. Moreover, the firms' decision to adopt a niche strategy acts as a form of differentiation that softens price competition and, so, effectively creates a positive externality on other firms. Indeed, this observation suggests a rationale for industry coordination: since profits can be nonmonotonic in search costs, as search costs fall exogenously, industries might benefit from reducing them further (for example, through industry-sponsored comparison sites).

Finally, our comparative statics analysis provides a demand-side explanation of the long-tail effect. As search costs fall, a greater proportion of firms choose the niche strategy. Consumers search to a much greater extent and, consequently, niche firms may account for a larger proportion of the industry's sales. In addition, lower search costs can simultaneously account for a superstar effect, with sales to both the head and tail of the sales distribution coming from middling firms whose designs change from broad to niche. Section VI suggests that for a long-tail effect to arise, there should be relatively low vertical differentiation among firms as compared to the potential scope for horizontal differentiation through design.

An aspect that our model did not address explicitly is the entry of new firms into the market. Empirically, entry is an important phenomenon, as highlighted by the increase of book titles, mentioned in the introduction. We abstract from entry for clarity of presentation. Firm entry can be endogenized, however, simply by assuming a fixed entry cost. Qualitatively, the results and intuitions of the paper would remain unchanged. Further, the Internet has had broader impacts that go beyond search costs, and long-tail and superstar phenomena may reflect changes to production costs. In this paper, we have focused on changes to the demand side to isolate their effects, as we believe they are economically significant.

APPENDIX: PROOFS

PROOF OF LEMMA 1:

First, note that since $f_s(x)$ is log-concave, $(1 - F_s(x))/f_s(x)$ is strictly decreasing in x (see, for example, Corollary 2 of Bagnoli and Bergstrom 2005). Suppose (for contradiction) that at some value of U , $p_{vs}(U)$ is increasing in U ; then, $p_{vs}(U) + U$ is also increasing in U , and so $(1 - F_s(p_{vs}(U) + U - v))/f_s(p_{vs}(U) + U - v) = p_{vs}(U)$ is decreasing in U , which provides the requisite contradiction. A similar argument ensures that $p_{vs}(U) + U$ is increasing in U , that $p_{vs}(U)$ is increasing in v , and that $p_{vs}(U) - v$ is decreasing in v .

PROOF OF PROPOSITION 1:

The optimal design is chosen to maximize $p_{vs}(U)(1 - F_s(p_{vs}(U) + U - v))$. Now, given that $p_{vs} + U - v$ is an affine transformation of p_s , it follows that $D_v(p_{vs}, s)$, as in (5), are rotation-ordered. The proof then follows immediately from Proposition 1 in Johnson and Myatt (2006, p. 761).

PROOF OF PROPOSITION 2:

For a fixed value of U , in principle, there may be more than one V solving equation (8). We show later that this is not the case. Consider one such solution and notice that

$$(15) \quad p_{VB}(U)(1 - F_B(p_{VB}(U) + U - V)) = p_{VN}(U)(1 - F_N(p_{VN}(U) + U - V))$$

$$(16) \quad \geq p_{VB}(U)(1 - F_N(p_{VB}(U) + U - V)).$$

It follows that

$$(17) \quad 1 - F_B(p_{VB}(U) + U - V) \geq 1 - F_N(p_{VB}(U) + U - V).$$

Similarly,

$$(18) \quad 1 - F_N(p_{VN}(U) + U - V) \geq 1 - F_B(p_{VN}(U) + U - V).$$

We use these facts to show that $p_{vN}(U) > p_{vB}(U)$. Suppose (for contradiction) that $p_{vN}(U) < p_{vB}(U)$. Note that since N and B are drawn from a family of demand rotations, it follows that there is some \tilde{x} such that $1 - F_N(x) > 1 - F_B(x)$ if and only if $x > \tilde{x}$. If $p_{vB}(U) + U - V > \tilde{x}$, then, $1 - F_N(p_{vB}(U) + U - V) > 1 - F_B(p_{vB}(U) + U - V)$ in contradiction to (17). If, instead, $\tilde{x} \geq p_{vB}(U) + U - V > p_{vN}(U) + U - V$, then (18) is contradicted. Thus, $p_{vN}(U) > p_{vB}(U)$ and from (8), trivially,

$$(19) \quad 1 - F_B(p_{vB}(U) + U - V) > 1 - F_N(p_{vN}(U) + U - V).$$

Define $\pi_{vs} := p_{vs}(1 - F_s(p_{vs} + U - v))$ with $s = B, N$. Since the price is chosen to maximize profits, by the envelope theorem, we have that $d\pi_{vs}/dv = p_{vs}f'_s(p_{vs} + U - v) = 1 - F_s(p_{vs}(U) + U - v)$ where the second equality follows from (7). Now, given (19), it follows that $d\pi_{vN}/dv < d\pi_{vB}/dv$. This ensures that $\pi_{vB} - \pi_{vN}$ always crosses zero from above and the uniqueness of $V(U)$ follows trivially.

Finally, by definition, $V(U)$ satisfies

$$\begin{aligned} p_{v(U)B}(U)(1 - F_B(p_{v(U)B}(U) + U - V(U))) \\ = p_{v(U)N}(U)(1 - F_N(p_{v(U)N}(U) + U - V(U))). \end{aligned}$$

Taking the derivative of both sides with respect to U , applying the envelope theorem, and using (7), we obtain

$$\begin{aligned} -(1 - F_N(p_{vN} + U - V(U)))\left(1 - \frac{dV(U)}{dU}\right) \\ = -(1 - F_B(p_{vB} + U - V(U)))\left(1 - \frac{dV(U)}{dU}\right), \end{aligned}$$

which, given (19), implies that $dV(U)/dU = 1$.

PROOF OF PROPOSITION 3:

If consumers use a $U = 0$ search rule, firms would react using a $V(0)$ strategy. Now, using (9), one can compute the search cost c_0 that delivers $(0, V(0))$ as an equilibrium:

$$(20) \quad \begin{aligned} c_0 = \int_{-\infty}^{V(0)} \left(\int_{p_{vN}(0)-v}^{\infty} (\varepsilon - p_{vN}(0) + v)f_N(\varepsilon) d\varepsilon \right) h(v) dv \\ + \int_{V(0)}^{\infty} \left(\int_{p_{vB}(0)-v}^{\infty} (\varepsilon - p_{vB}(0) + v)f_B(\varepsilon) d\varepsilon \right) h(v) dv. \end{aligned}$$

For all $c > c_0$, consumers would not search since this requires that the consumer search threshold, U , is below 0.

Next, consider the consumer stopping rule U_B that makes all firms prefer the broad strategy and the lowest-quality firm \underline{v} indifferent. This is the highest level of consumer search compatible with all firms offering a broad product and is characterized by

$$(21) \quad p_{vB}(U_B)(1 - F_B(p_{vB}(U_B) + U_B - \underline{v})) \\ = p_{vN}(U_B)(1 - F_N(p_{vN}(U_B) + U_B - \underline{v})).$$

If $U_B \leq 0$, clearly there is no all-broad equilibrium. If $U_B > 0$, using (9), one can compute the search cost c_B that delivers (U_B, \underline{v}) as an equilibrium:

$$(22) \quad c_B := \int_{-\infty}^{\infty} \left(\int_{p_{vB}(U) + U_B - v}^{\bar{\theta}_B} (v + \varepsilon - p_{vB}(U) - U_B) f_B(\varepsilon) d\varepsilon \right) h(v) dv.$$

Take, now, any $c \in [c_B, c_0)$. Suppose that all firms choose a broad design, then the consumer's optimal search threshold is implicitly defined, according to (9), by

$$(23) \quad c = \int_{-\infty}^{\infty} \left(\int_{p_{vB}(U) + U - v}^{\bar{\theta}_B} (v + \varepsilon - p_{vB}(U) - U) f_B(\varepsilon) d\varepsilon \right) h(v) dv.$$

Since the right-hand side is decreasing in U , it follows that the U that solves this equation is less than U_B . Proposition 2 shows that $V(\cdot)$ is increasing; thus, all firms do indeed prefer to choose the broad strategy. This proves that (U, \underline{v}) constitutes an equilibrium. A similar argument shows that there is no all-broad equilibrium if $c < c_B$.

Next, we consider equilibria in which all firms choose the niche design. Analogous to the all-broad case, one can define U_N and c_N :

$$(24) \quad p_{\bar{v}B}(U_N)(1 - F_B(p_{\bar{v}B}(U_N) + U_N - \bar{v})) \\ = p_{\bar{v}N}(U_N)(1 - F_N(p_{\bar{v}N}(U_N) + U_N - \bar{v})),$$

$$(25) \quad c_N = \int_{-\infty}^{\infty} \left(\int_{p_{\bar{v}N}(U) + U_N + -v}^{\bar{\theta}_N} (v + \varepsilon - p_{\bar{v}N}(U_N) - U_N) f_N(\varepsilon) d\varepsilon \right) h(v) dv,$$

and

$$\underline{c} = \begin{cases} c_N & \text{if } U_N > 0 \\ c_0 & \text{if } U_N \leq 0 \end{cases}$$

and show that an all-niche equilibrium exists if and only if $c' \leq \underline{c}$.

PROOF OF PROPOSITION 4:

Consider a stable equilibrium $(U, V(U))$. According to our definition, this means that $(\partial U/\partial V)(V(U)) < 1$, where $U(V)$ is implicitly defined by expression (9). Denote its left-hand side as $H(U, V)$. Note that

$$\frac{\partial U}{\partial V}(V(U)) = \frac{\frac{\partial H}{\partial V}(U, V(U))}{-\frac{\partial H}{\partial U}(U, V(U))} < 1 \Leftrightarrow \frac{\partial H}{\partial U}(V(U)) + \frac{\partial H}{\partial V}(V(U)) < 0.$$

Now, since $(\partial V/\partial U)(U) = 1$, it follows that stability is satisfied if and only if

$$(26) \quad \frac{\partial H}{\partial U}(V(U)) + \frac{\partial H}{\partial V}(V(U)) \frac{\partial V}{\partial U}(U) < 0.$$

Now, at an equilibrium, $\bar{H}(U) \doteq H(U, V(U)) = c$. Moreover, following (26), we see that $(\partial \bar{H}/\partial U)(U) < 0$. Finally, if c falls, $\bar{H}(U)$ needs to decrease as well, which implies that U needs to increase to restore equilibrium. Finally, using Proposition 2, we know that $V(U)$ increases as well.

PROOF OF PROPOSITION 5:

Following the argument in the text, part (i) is immediate and consumer surplus is constant at \bar{U} throughout this region. Next, part (ii) follows immediately from (14) since $c_B > c_N$. From the proof of Proposition 2, $1 - F_B(p_B(\bar{U}) + \bar{U}) > 1 - F_N(p_N(\bar{U}) + \bar{U})$; then, part (iii) follows immediately.

Firm profits for niche and broad firms are identical and given by $(m/\rho(U))p_N(\bar{U}) \times (1 - F_N(p_N(\bar{U}) + \bar{U})) = (m/\rho(U))p_B(\bar{U})(1 - F_B(p_B(\bar{U}) + \bar{U}))$. Following part (iii) of the proposition, part (iv) follows immediately.

Finally, sales for a broad and a niche firm are $(m/\rho(U))(1 - F_B(\bar{U} + p_B(\bar{U}))$ and $(m/\rho(U))(1 - F_N(\bar{U} + p_N(\bar{U}))$. Again, following part (iii) of the proposition, part (v) follows immediately.

PROOF OF PROPOSITION 6:

See online Appendix A.

REFERENCES

- Anderson, Chris. 2004. "The Long Tail." *Wired* 12(10): <http://www.wired.com/wired/archive/12.10/tail.html>.
- Anderson, Chris. 2006. *The Long Tail: Why the Future of Business Is Selling Less of More*. New York: Hyperion.
- Anderson, Chris. 2009. *The Longer Long Tail: How Endless Choice is Creating Unlimited Demand*. London: Random House Business Books.
- Anderson, Simon P., and Regis Renault. 1999. "Pricing, Product Diversity, and Search Costs: A Bertrand-Chamberlin-Diamond Model." *RAND Journal of Economics* 30(4): 719–35.
- Anderson, Simon P., and Regis Renault. 2009. "Comparative Advertising: Disclosing Horizontal Match Information." *RAND Journal of Economics* 40(3): 558–81.
- Bagnoli, Mark, and Ted Bergstrom. 2005. "Log-Concave Probability and Its Applications." *Economic Theory* 26(2): 445–69.
- Bakos, Yannis. 1997. "Reducing Buyer Search Costs: Implications for Electronic Marketplaces." *Management Science* 43(12): 1676–92.

- Bar-Isaac, Heski, Guillermo Caruana, and Vicente Cuñat.** Forthcoming. "Information Gathering Externalities in Product Markets." *Journal of Industrial Economics*.
- Bar-Isaac, Heski, Guillermo Caruana, and Vicente Cuñat.** 2010. "Information Gathering and Marketing." *Journal of Economics and Management Strategy* 19(2): 375–401.
- Bowker Industry Statistics.** 2009. "New Book Titles and Editions, 2002–2008." <http://bowker.com/bookwire/IndustryStats2009.pdf>.
- Brynjolfsson, Erik, Yu Jeffrey Hu, and Duncan Simester.** 2007. "Good-bye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=953587.
- Brynjolfsson, Erik, Yu Jeffrey Hu, and Michael D. Smith.** 2003. "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers." *Management Science* 49(11): 1580–96.
- Brynjolfsson, Erik, Yu Jeffrey Hu, and Michael D. Smith.** 2006. "From Niches to Riches: Anatomy of the Long Tail." *Sloan Management Review* 47(4): 67–71.
- Cachon, Gerard P., Christian Terwiesch, and Yi Xu.** 2008. "On the Effects of Consumer Search and Firm Entry in a Multiproduct Competitive Market." *Marketing Science* 27(3): 461–73.
- Diamond, Peter A.** 1971. "A Model of Price Adjustment." *Journal of Economic Theory* 3(2): 156–68.
- Elberse, Anita, and Felix Oberholzer-Gee.** 2007. "Superstars and Underdogs: An Examination of the Long Tail Phenomenon in Video Sales." *Marketing Science Institute* 4: 49–72.
- Goldmanis, Maris, Ali Hortacsu, Chad Syverson, and Onsel Emre.** 2010. "E-Commerce and the Market Structure of Retail Industries." *Economic Journal* 120(545): 651–82.
- Greco, Albert N.** 2004. *The Book Publishing Industry*. 2nd ed. New York: Routledge.
- Hervás-Drane, Andrés.** 2010. "Search, Information Provision, and Sales Concentration." Unpublished.
- Johnson, Justin P., and David P. Myatt.** 2006. "On the Simple Economics of Advertising, Marketing, and Product Design." *American Economic Review* 96(3): 756–84.
- Kuksov, Dmitri.** 2004. "Buyer Search Costs and Endogenous Product Design." *Marketing Science* 23(4): 490–99.
- Larson, Nathan.** Forthcoming. "Niche Products, Generic Products, and Consumer Search." *Economic Theory*.
- Lewis, Tracy R., and David E. M. Sappington.** 1994. "Supplying Information to Facilitate Price Discrimination." *International Economic Review* 35(2): 309–27.
- McCall, John J.** 1970. "Economics of Information and Job Search." *Quarterly Journal of Economics* 84(1): 113–26.
- Oestreicher-Singer, Gal, and Arun Sundararajan.** Forthcoming. "Recommendation Networks and the Long Tail of Electronic Commerce." *MIS Quarterly*.
- Tucker, Catherine, and Juanjuan Zhang.** 2011. "How Does Popularity Information Affect Choices? A Field Experiment." *Management Science* 57(5): 828–42.
- Watson, Randal.** 2009. "Search and Availability in Monopolistic Competition." Unpublished.
- Wolinsky, Asher.** 1986. "True Monopolistic Competition as a Result of Imperfect Information." *Quarterly Journal of Economics* 101(3): 493–511.