

MPRA

Munich Personal RePEc Archive

Niche products, generic products, and consumer search

Larson, Nathan

University of Virginia

2008

Online at <https://mpra.ub.uni-muenchen.de/32161/>
MPRA Paper No. 32161, posted 11 Jul 2011 19:58 UTC

Niche Products, Generic Products, and Consumer Search

Nathan Larson*
Department of Economics
University of Virginia
P.O. Box 400182
Charlottesville, VA 22904-4182

July 10, 2011

Abstract

We endogenize product design in a model of sequential search with random firm-consumer match value à la Wolinsky (1986) and Anderson and Renault (1999). We focus on a product design choice by which a firm can control the dispersion of consumer valuations for its product; we interpret low dispersion products as ‘generic’ and high dispersion products as ‘niché.’ Equilibrium product design depends on a feedback loop: when reservation utility is high (low), the marginal customer’s match improves (worsens) with more niche products, encouraging high (low) differentiation by firms. In turn, when firms offer more niche products, this induces more intense search; depending on search costs, this could raise or lower consumers’ reservation utility. Remarkably, when the match distribution satisfies a hazard rate condition, firm and consumer interests align: equilibrium product design always adjusts to the level that maximizes utility. When this condition is not met, either multiple equilibria (one niche, the other generic) or one asymmetric equilibrium (generic and niche firms coexist) can arise; we argue that the former is more likely for common specifications of consumer preferences.

Keywords: product differentiation; search; product design
JEL Codes D43 D83 L15

1 Introduction

Consider an entrepreneurial chemist choosing the scent profile for a new perfume. Understanding that consumer tastes are idiosyncratic, he might emphasize safe smells – say, vanilla or lavender – that most consumers would find pleasant and inoffensive. Alternatively, he could emphasize bold, exotic scents that some consumers would love and others would hate. Furthermore, he knows that a consumer cannot be sure of exactly how much she will like his perfume without at least making a trip to the store to sample it, and he faces competition from many other chemists who face the same choices that he does. How unique or generic should he make his product, how competitively should he price it, and how do consumer search costs factor into these decisions?

*I’d like to thank without implication Simon Anderson, Susan Athey, Abhijit Banerjee, Glenn Ellison, Bengt Holmstrom, and participants at the MIT theory and Industrial Organization lunches.

This paper addresses these questions by introducing endogenous product design into a canonical model of sequential consumer search. Broadly, we find a negative relationship between search costs and what we will call the “nichiness” of products: firms choose polarizing niche products when search costs are low and more generic products when search costs are high. In the low search cost case, when consumers are relatively selective, firms can soften price competition with an idiosyncratic product that provides very high value to a relatively small set of consumers. Alternatively, when firms make their products as generic as possible, consumers have little incentive to search for a better match; when search costs are high this turns out to be a better way to soften price competition.

The search model that we use is based on Wolinsky (1986) and Anderson and Renault (1999, henceforth AR) and features a continuum of firms and consumers. Each consumer’s value for a particular firm’s product depends on a match-specific taste shock with mean zero (plus a constant term that is common to all consumers). The firm’s product design choice will involve choosing the variance of the taste shock for its product; in the perfume example, this corresponds to the decision about how bland or provocative to make the scent. We will refer to a product with a high taste shock variance as *nichy*, or *idiosyncratic*, or *specialized*; for our purposes, these terms will all mean the same thing. In order to focus on the niche versus generic aspect of product design, we treat the average quality of a product – that is, the mean taste shock – as fixed.¹ In the game, firms first (simultaneously) choose prices and product nichiness. Consumers do not observe these choices, but they form expectations about their aggregate distribution and believe that they will face a random draw from this distribution at any given firm. Next consumers search. A consumer learns about her valuations for the firms’ products, and the prices they are charging, by visiting them randomly and sequentially, incurring a constant search cost with each visit. The optimal search strategy involves a cutoff rule: a consumer purchases from the first firm at which her surplus (valuation net of price) exceeds a threshold. The focus is on pure strategy outcomes, and our equilibrium concept, endogenous dispersion equilibrium (EDE), requires each firm to choose prices and nichiness optimally given correct expectations about consumers’ cutoff utilities, and each consumer to choose a utility cutoff rule optimally given correct expectations about the aggregate distribution of prices and nichiness.²

The model is best suited to describing what could be called ‘sample goods’ for which consumers have idiosyncratic tastes.³ These are similar to experience goods in the sense that a consumer must spend some time, money, or effort interacting with the good before she is able to evaluate how much she likes it. However, they differ from experience goods because the consumer does not need to purchase the product outright before evaluating it – she can sample it (at search cost c) instead. Goods that fit this profile include many consumer products for which tastes are personal, such as books, music, cars, and clothing. For example, a consumer can sample the work of an unfamiliar author for the time cost of a trip to the bookstore to browse through its pages.⁴ This can quickly give her a sense of how well this particular author’s style suits her own taste. In this case, product design is related to the author’s style and genre.

¹However, it is not difficult to incorporate a quality decision into the model of product design; we sketch an extension along these lines in Section 4.

²Here, “endogenous dispersion” is intended to refer to the fact that a firm can control the dispersion of consumer valuations for its product, not to price dispersion.

³I thank one of the referees for pointing me toward this terminology.

⁴Or increasingly, for the time cost of browsing through free excerpts of the book online.

An airport spy novel amounts to a generic product: it is no one's ideal product, but it serves most consumers relatively well in a pinch. On the other hand, other authors tend to provoke stronger reactions: readers either love the work or hate it.⁵ Their books would be called niche products in our model. With a car, a consumer can research many details in advance, but it is hard for her to be sure how much she will enjoy driving it without taking the time to do a test drive. In this case, a niche product could be one with particularly sporty handling: some drivers will like the responsiveness, but others will wish the ride were less bumpy. Clothing helps to illustrate how online markets relate to our model. Online markets would seem to offer less scope for sample goods, since one cannot handle the good before buying it. However, one could argue that lenient return policies for online clothing purchases (sometimes including free return shipping) have evolved to make online clothing more like a sample good: a consumer can sample an item's fit at a relatively low cost (time and shipping), and return it if she dislikes it.

Our first main finding is that a firm's optimal product is always extreme (Lemma 2 and Proposition 2). That is, a firm will create a product that disperses consumers' valuations either as much as possible (a niche product) or as little as possible (a generic product), depending on whether consumer utility is above or below a threshold level. In each case, the intuition is roughly that a firm gains by improving the match with its marginal customer. When consumers are relatively choosy, only customers with positive taste shocks will purchase, and these customers are made happier by a more distinctive product. However, for low enough utility, the marginal consumer is indifferent between settling for a negative taste shock and continued search. In this case, making the product more generic makes this marginal consumer less displeased with it, discouraging her from searching further.

Next we characterize the set of endogenous dispersion equilibrium at different levels of the search cost (Proposition 3). The progression is generally as follows. For low search costs, there is one EDE with maximally niche products. For intermediate search costs, there are either three EDEs – one maximally niche, one minimally niche, and an asymmetric EDE with both niche and generic firms – or there is just one asymmetric EDE. For yet higher search costs, there is one generic (minimally niche) EDE. For higher search costs, consumers prefer not to search, and there is no equilibrium. There are two caveats to this. First, for some parameters, equilibrium may fail sooner, so that the intermediate or high cost cases above may not appear. Second, it is possible for the intermediate case to vanish, so that the unique EDE shifts directly from niche to generic at a threshold search cost level. Thus, in a general sense, lower search costs are associated with nichier products. Social surplus is maximized when products are maximally idiosyncratic, regardless of search costs (Proposition 7), so product design represents a second channel through which falling search costs improve welfare.

For search costs in the intermediate range, the type of equilibrium that arises depends subtly on whether consumers or firms would capture more of the surplus gains associated with a shift from a generic to a niche market. If it is the consumers who capture more, then for the same search cost there can be both a generic EDE with consumers who settle for low utility and a niche EDE with consumers who demand high utility, reinforced by firms' incentive to offer nichier products to choosier consumers. Alternatively, the unique asymmetric equilibrium can arise if firms are able to capture so much of the gains from a wholesale shift to niche products that consumers' cutoff utility falls as a result. In this case, neither a completely generic market

⁵Joyce and Faulkner come to mind, but every reader will have their own examples.

nor a completely niche market is stable; in the former, consumers are so choosy that firms would rather offer niche products, while in the latter consumers are sufficiently accepting that firms would like to switch back to generic products. In this case, generic and niche firms must coexist for intermediate search costs.

We show that when firm and consumer interests are aligned in a particular sense, neither of these cases applies: the market has a unique equilibrium that shifts abruptly from generic to niche products as search costs fall below a threshold value (Proposition 4). This alignment occurs for a large family of consumer taste distributions (including the uniform and exponential) that satisfy a hazard rate condition. For these distributions, differentiation always adjusts endogenously to the level that maximizes consumers' equilibrium utility (Proposition 6).

The hazard rate condition roughly relates to whether a firm can extract a constant fraction of the total surplus from a match as consumers become choosier. We argue that if this fraction declines (rises) with choosier consumers, multiple symmetric equilibria (one asymmetric equilibrium) are more likely to arise. For taste distributions that are commonly used (such as the normal, logistic, extreme value, and generalized Pareto), this fraction is either constant or declining – as consumers are pushed further into the right tail of their tastes, their surplus rises faster than the equilibrium price. In contrast, the examples of asymmetric equilibria that we have found involve taste distributions with abruptly truncated right tails. Absent any compelling reason to expect this type of truncation in tastes, we regard the coexistence of “generic” and “specialized” firms as an interesting but probably uncommon outcome in our model.

The paper leads off with a connection to product differentiation. In the model, taste shocks are drawn independently across consumer-firm pairs. This will imply (Proposition 1) that a higher level of nichiness among products – that is, greater dispersion in consumer valuations for each product – can also be interpreted as a greater degree of differentiation between every pair of products. In this sense, our paper connects to prior work on the role of both search frictions and product differentiation in softening price competition. These topics have been studied extensively but usually separately, with seminal contributions by Diamond (1971) on the former and Perloff and Salop (1985) on the latter. The interaction between the two was studied in a unified model first by Wolinsky (1986) and later Anderson and Renault (1999). AR use this model to study entry and to derive reasonable comparative statics predictions about the response of prices to search costs, the heterogeneity of consumer tastes, and the level of product differentiation, which is taken to be exogenous. We build on AR in several ways. First, we introduce an endogenous product design choice for firms and demonstrate a formal sense in which it can “differentiate” consumers' values for different products. Then, by characterizing equilibria with endogenously niche or generic products, we are able to study how product design responds to search costs. Particularly for intermediate search costs, the subtle relationship between the distribution of consumer tastes and the response of product design would not have been obvious in a model without endogenous differentiation.

A recent spate of papers extends the AR model in a different direction by studying what happens if search is directed rather than random. In Arbatskaya (2007), Haan and Moraga González (2007), and Armstrong, Vickers, and Zhou (2008) consumers search in an order related to how prominent firms are or how much they advertise. The latter also study an extension with quality differentiation, finding that a higher quality firm has a greater incentive to make itself more prominent to consumers. Another strand of the literature builds on Butters' (1977)

monopoly model of consumers who learn about products from advertising. Christou and Vettas (2008) extend this model to differentiated product competition. They find that the manner in which consumers are informed can tend to generate firm profit functions that are not quasiconcave, so a firm’s optimal strategy can jump between high sales at a low price and low sales at a high price. While the causes are different, this is reminiscent of our finding that when firms optimize over both product design and price, their profits are not quasiconcave with respect to the product design choice. Finally, Anderson and Renault (2006) allow a firm to advertise information either about its price or about its product. This literature on directed search and advertising is complementary to our paper’s focus on endogenous product differentiation. Combining the two, in particular, studying product design when firms can advertise both price and product features, would be a natural subject for further study.

In contrast with the non-spatial approach to product differentiation in Wolinsky and AR, another branch of the literature studies differentiation that arises from a firm’s location, where this location might be in physical or product space. Here, endogenous location choice by firms is often the first stage in multi-stage competition that ends with the firms competing in prices. While we will not try to survey this literature here, two of the most celebrated results predict *extreme* levels of differentiation. Hotelling (1929) famously showed (among other results) that two firms with fixed and equal prices will choose to differentiate their products *minimally* in equilibrium (by choosing the same location). In contrast, d’Aspremont, Gabszewicz, and Thisse (1979) show that when the firms do choose prices, this result reverses – pure strategy equilibrium outcomes involve *maximal* differentiation. The intuition for the latter result is loosely that differentiating more is in each firm’s individual interest because the gains from softened price competition outweigh the losses from moving further away from the tastes of the median consumer. While this intuition is appealing, testing it in settings with many firms and alternative product spaces has been difficult because spatial models of product differentiation rarely “scale up” gracefully. By returning to a non-spatial model of differentiation à la Wolinsky and AR, we will be able to show that these examples illustrate a general principle of extreme differentiation and that an outcome of maximal or minimal differentiation in any particular case can be explained in a sensible way by consumer search costs.

Two recent papers in the marketing literature also touch on the interaction of endogenous product design and search costs. Kuksov’s (2004) setting is quite different from ours: he looks at spatial product differentiation between duopolists when consumers search for prices but know their product preferences in advance. He finds the same general pattern that we do – product differentiation rises as search costs fall – but otherwise the models are not easily compared. More closely related are Cachon, Terwiesch, and Xu (2008) who study both sequential and non-sequential search in a random utility framework similar to AR. Their firms choose how many products to offer, and a visiting consumer can purchase whichever of these products gives her the best (i.i.d.) match. Thus, in both models a firm chooses a parameter that affects the distribution of taste shocks realized by consumers who visit. Because every consumer’s best match at a firm improves (in expectation) as the firm adds more products, their model cannot disentangle the effect of this overall quality improvement from any effect related to horizontal differentiation, so in this respect, the two models address different questions. Furthermore, their model specializes to extreme value-distributed taste shocks, and some of their results

appear to depend delicately on this assumption.⁶ In contrast, we show that the shape of the taste distribution has a major influence on equilibrium outcomes for intermediate search costs. Finally, in recent work, Johnson and Myatt (2006) develop a model of product design, involving rotations of a firm’s demand curve, that has a similar flavor to our distinction between generic and niche products. As in our paper, they also find that a firm’s optimal design tends toward one of these two extremes. However, their focus is on advertising by firms, and they do not consider search.

The rest of the paper is laid out as follows. Section 2 introduces the model and develops useful partial equilibrium results. Section 3 characterizes equilibrium and contains the main results of the paper. Section 4 considers extending the options available to a firm in two important ways: by allowing investments in higher product quality, and by allowing a firm to offer a product line with more than one product. Section 5.

2 The Model

2.1 Consumer preferences and differentiation

The model is one of symmetric, non-spatial competition in horizontally differentiated products with sequential, random search by consumers. Much of the basic structure is shared with Wolinsky (1986) and Anderson and Renault (1999). There is a continuum of firms, indexed by $m \in [0, 1]$, each selling a single differentiated good. A consumer i has willingness to pay for good m given by

$$A_{im} = A_\mu + \sigma_m z_{im}$$

In this expression, A_μ should be interpreted as a consumer’s average valuation for the good, across all firms. (This is the same for all consumers.) The second term, $\sigma_m z_{im}$, reflects how much more or less than average consumer i likes the particular version of the good offered by firm m . In particular, firm m ’s good embodies an amount σ_m of a polarizing feature (also labeled m). A firm’s product design choice will be to choose the level of σ_m . The possible levels of that polarizing feature are represented by a positive interval: $\sigma_m \in [\sigma_L, \sigma_H]$, where $0 < \sigma_L < \sigma_H$. Consumer i ’s marginal utility for firm m ’s feature is given by z_{im} . Some consumers will be pleased and others displeased to have more of feature m ; specifically, these marginal utilities are distributed randomly with density $f(z_{im})$, independently across i and m . This incorporates two assumptions: consumer valuations are distributed symmetrically across all goods, and valuations are independent across goods – a consumer’s preference for one attribute (and its corresponding good) provides no information about his preference for other goods. Also define the corresponding cumulative distribution function $F(z)$. Both the distribution and density functions are assumed to be continuously differentiable, and the following condition is also imposed.

Condition 1 *The density function $f(z)$ is symmetric, logconcave, and has mean 0 and support $(-\infty, \infty)$.*

The zero mean assumption is used to isolate the choice of σ_m , which affects the dispersion of consumer valuations for a product, from product design choices (quality improvements) that

⁶For example, they find that when firms behave symmetrically, the intensity of consumer search does not change with the number of products per firm.

improve the average valuation A_μ . We focus on the choice of σ_m in this paper, but it is not difficult to extend the model to consider quality improvements as well – we sketch such an extension in Section 4. The assumption of an unbounded support is technically convenient but not essential. Symmetry is only necessary for the strong version of product differentiation introduced in Definition 1, while logconcavity helps both in defining product differentiation and later in guaranteeing an interior solution to a firm’s optimal pricing problem. Many commonly used distributions are logconcave, including the normal and uniform distributions (see Bagnoli and Bergstrom (2005) for examples). Logconcavity of the density implies logconcavity of the distribution function, and logconcave distributions have increasing hazard rates, facts that will be useful in the sequel.

Since this formulation of product differentiation can seem a bit abstract, a few concrete examples may help to fix ideas. One could think, for example, of a restaurant choosing its format. *Holding the quality level fixed*, a neutral format like modern American cuisine might generate relatively small taste differences among consumers, while a more polarizing format like a less familiar ethnic cuisine or a theme (rock and roll, medieval, etc.) might substantially enhance the experience for some diners and detract from it for others. We would interpret the modern American restaurant as a low σ choice and the other options as high σ choices. Similarly, an apparel firm choosing slogans for a tee shirt might opt either for something bland and inoffensive, or for something with niche appeal (bawdy humor, political slogans, etc) that will substantially enhance its value to some consumers but make it unwearable for others – again, these would be low σ or high σ choices, respectively. The practical implication of the independence of the match-specific shock across firms and consumers is that there is no shortage of directions along which firms might differentiate. In the restaurant example, there is effectively a limitless number of idiosyncratic restaurant formats (Mexican, Thai, Ethiopian, Indian, ...), each of which can come in a lower σ versions that tone the food down for average tastes and higher σ versions that cater to enthusiasts.⁷

There is another possible interpretation of the firm’s product differentiation choice which does not require assuming that some consumers dislike more of some features. In this interpretation, a product consists of a bundle of features. While consumers have positive marginal utility for all features, technological constraints require a firm that employs more of one feature to trade it off against less of some other feature. For example, one could think of cell phones. At a given cost point, current levels of miniaturization might allow a producer to incorporate a superb camera and a terrible music player, a very good music player and an adequate web browser, or decent, but not great, functionality for all three features. While all of these combinations might generate the same *average* consumer valuation, the more lopsided combinations (e.g. an excellent camera, with lip service paid to other features) may generate more of a split between consumers who really love the emphasized feature and consumers who care more about other things – this would be a high σ product. On the other hand, more balanced products that offer something to appeal to everyone will tend to generate consumer valuations with a

⁷This is, of course, a stylization. Furthermore, independence implies that there is no correlation or crowding out of tastes – my tastes for Indian and Ethiopian foods are unrelated to each other and are not statistically closer to each other than either is to my taste for Italian food. In this respect as well, the model (and non-spatial models generally) is an imperfect fit to reality. While this stylization is standard, and enormously useful in keeping the model tractable, exploring models with a more nuanced structure of match-specific taste shocks would certainly be of interest.

lower variance – these would be low σ products.⁸

To be mathematically precise about the sense in which higher σ corresponds to greater differentiation, we introduce the following definition. According to this definition, two goods become more or less differentiated as a typical consumer’s valuations for them grow further apart or closer together:

Definition 1 *Let $D_{mn}(k)$ be the probability that the difference in a consumer’s values for goods m and n is less than k : $D_{mn}(k) = \Pr(|A_{im} - A_{in}| \leq k)$. We will say that goods m' and n' are weakly more (less) differentiated than goods m and n if $D_{m'n'}(k) \leq D_{mn}(k)$ for all $k \geq 0$ ($D_{m'n'}(k) \geq D_{mn}(k)$) for all $k \geq 0$, with the inequality strict for some k . If the inequality is strict for all $k > 0$, we will say that m' and n' are strictly more (less) differentiated.*

Under this definition, goods do in fact become more differentiated as more of either polarizing feature is added:

Proposition 1 *Fix any three goods m , m' , and n , and assume Condition 1 holds. Then m' and n are strictly more differentiated than m and n if and only if $\sigma_{m'} > \sigma_m$.*

Proof. See the appendix. ■

Notice again that differentiation is non-spatial: by further customizing its product, a firm distances itself from all other firms, and by making its product more generic, it crowds all other firms. Furthermore, differentiation operates symmetrically among all of the goods: no two goods are intrinsically “closer” to each other than any other two. Loosely, Proposition 1 tells us that the fraction of consumers that could be induced to switch between products m and n by a price difference of k declines as σ_m or σ_n rises, for any price difference k .

While logconcavity and symmetry cover many important taste distributions, they exclude some interesting cases. To cover these cases, we introduce a second, weaker definition of horizontal differentiation. Suppose only that f has a zero mean and finite variance s^2 .

Definition 2 *The mean square taste difference between goods m and n is $S_{mn} \equiv E\left((A_{im} - A_{in})^2\right)$*

It is trivial to show that the mean square taste difference between goods m and n is increasing in both σ_m and σ_n :

$$S_{mn} = E\left((\sigma_m z_{im} - \sigma_n z_{in})^2\right) = (\sigma_m^2 + \sigma_n^2) s^2$$

An increase in S_{mn} indicates that consumer valuations for m and n are further apart on average, but they need not be further apart in the stricter pointwise sense of Definition 1. As a simple example of how these two definitions can diverge, consider the degenerate taste distribution F that places equal weight on $z = 1$ and $z = -1$. Suppose firm m chooses $\sigma_m = 1$ and firm n chooses $\sigma_n = 0$, so $S_{mn} = 1$. If firm n' chooses $\sigma_{n'} = 1$, then $S_{mn'} = 2$, so mean square taste differences between m and n' are greater than between m and n . However, for m and n , the realized difference $|\sigma_m z_m - \sigma_n z_n|$ always equals 1. For m and n' , this difference is sometimes 2 and sometimes 0, so some consumers will find m and n' less similar than m and n , while others will find them more similar. For this reason, we should not expect the implications for price competition to be quite as clear and unambiguous when only Definition 2 applies, but not Definition 1.

⁸Interested readers can consult the working version of this paper for a formal model along these lines.

2.2 Timing and Equilibrium

The game has two stages. In the first stage, each firm simultaneously chooses its strategy, once and for all. In general, a strategy for firm m will be a pair $(\sigma_m, p_m) \in [\sigma_L, \sigma_H] \times [0, \infty)$, where σ_m is the nichiness of the firm's product (chosen from a compact interval $[\sigma_L, \sigma_H]$) and p_m is its price. Both components of the strategy are chosen simultaneously. We focus on pure strategies for a single firm, but as is usual with a continuum of firms, asymmetric strategy profiles can be given a mixed strategy interpretation. Let us summarize the distribution of (σ, p) pairs in the firms' strategy profile with a function $P : [\sigma_L, \sigma_H] \times [0, \infty) \rightarrow [0, 1]$, where $P(\sigma, p)$ is the measure of firms m with strategies satisfying $\sigma_m \leq \sigma$ and $p_m \leq p$. Note that P is analogous to a cumulative distribution function.

In the second stage of the game, consumers search for products. Each consumer i has unit inelastic demand and realizes a net utility equal to $u_{im} = A_{im} - p_m$ if he purchases from firm m . Consumer i knows the taste shock distribution $F(z)$, but he does not know his realized taste shock z_{im} at any firm that he has not visited. He does not observe the firms' first stage actions, but he forms beliefs about those actions. To reflect the idea that a consumer has no basis for distinguishing between firms, we restrict these beliefs to treat firms anonymously. Formally anonymity will mean the following.⁹

Anonymous Consumer Beliefs

1. Consumer i forms a belief function $B_i : [\sigma_L, \sigma_H] \times [0, \infty) \rightarrow [0, 1]$, interpreted as his belief about the distribution of (σ, p) pairs among firms.
2. At any firm m that he has not previously visited, consumer i believes (σ_m, p_m) to be an independent random draw from distribution B_i , and he believes his taste shock z_{im} to be an independent random draw from F .

Our equilibrium concept will require consumers to hold consistent beliefs about the aggregate distribution of firms' actions; thus, later we will impose $B_i = P$. Firms' choices and the taste shocks only matter to a consumer to the extent that they affect his utility u_{im} from a purchase, so we can summarize a consumer's beliefs by a probability distribution

$$G_i(u) = \Pr(u_{im} \leq u \mid B_i(\sigma, p), F(z))$$

over the net utility available to him at a randomly chosen firm.

A consumer has the following options. He can quit the market immediately, walking away with utility 0, or at cost $c > 0$, he can visit a randomly selected firm m where he learns (σ_m, p_m, z_{im}) . Thus he learns his valuation for the firm's product $A_{im} = A_\mu + \sigma_m z_{im}$, its price, and therefore, his net utility draw u_{im} . He then has four options: he can purchase the product from m and leave the market, he can leave the market without purchasing, he can purchase (at no additional cost) from any previously visited firm, or he can continue to search. If he continues to search, he incurs cost $c > 0$ and visits a new firm randomly chosen from those he has not previously visited. This process continues until the consumer has left the market.

Notice that our notion of anonymous beliefs precludes a consumer from revising his expectations about the strategies of unvisited firms on the basis the information that he observed

⁹I am grateful to an editor for suggesting this terminology.

by visiting firm m . This is restrictive, but in the context of our model it is reasonable, given that firms act independently, and there is no common factor in firms' decisions, like a common cost shock, for consumers to learn about. The assumption also serves to pin down a consumer's subsequent beliefs if she ever were to observe a price outside of the support of B_i .

Since there is a continuum of firms, consumer i 's decision problem amounts to an optimal stopping problem with a stationary distribution. He can observe a sequence of utility draws from $G_i(u)$, at cost c for each draw, and must decide when to quit and take one of the utilities (that is, buy one of the products) that he has seen so far. It is a standard result that in this setting an optimal search strategy for consumer i can be expressed in terms of some stationary cutoff \bar{u}_i . That is, as soon as consumer i visits a firm that gives him utility greater than or equal to \bar{u}_i , he purchases and leaves the game; otherwise he continues to search. This encompasses the option to quit the game immediately: if $\bar{u}_i < 0$, the consumer 'buys' his outside option utility of 0 and exits. With this justification in mind, there is no loss of generality in restricting our analysis of consumer strategies to the set of stationary cutoff rules.

While different consumers have different taste shocks *ex post*, *ex ante* they are identical. In the model, a firm cannot control how many consumers show up at its front door, and the number of consumers who do happen to show up has no bearing on the strategy that maximizes its profit per consumer arrival. Furthermore, consumers actions do not directly affect other consumers' payoffs. For these reasons, we will remain vague about the total number of consumers. One can think of a single representative consumer or a continuum of *ex ante* identical consumers; this makes no difference to the results. Furthermore, Lemma 1 will show that for any given belief about firms, a consumer's best response threshold is unique. Thus, any sensible notion of equilibrium with consistent consumer beliefs must have all consumers choosing the same (pure strategy) threshold \bar{u} . To avoid the surplus notation of defining distributions over \bar{u}_i , we will simply restrict attention to equilibria in which consumers have symmetric cutoff rules and beliefs.

Finally, note that in a standard oligopoly model without search, the strategies of competing firms enter firm m 's payoff function directly. Here, because of the sequential search structure, they do not – the expected profit that firm m receives from a consumer visit depends only on the consumer's choosiness \bar{u}_i and on its own strategy. (Of course, \bar{u}_i will depend on the consumer's *beliefs* about the strategies of other firms. Similarly, the number of consumers who arrive at firm m 's front door will depend on both \bar{u}_i and on other firms' strategies, but this has no bearing on firm m 's own strategy choice.)

Now we are prepared to define our equilibrium concept, which we call an *endogenous dispersion equilibrium (EDE)* to emphasize the fact that the dispersion of consumers' taste shocks is a choice variable for firms. An assessment for the game is a collection $\{\boldsymbol{\sigma}, \mathbf{p}, \bar{u}, B\}$, where $\boldsymbol{\sigma} : [0, 1] \rightarrow [\sigma_L, \sigma_H]$, with $\boldsymbol{\sigma}(m) = \sigma_m$, specifies firms' nichiness choices, $\mathbf{p} : [0, 1] \rightarrow [0, \infty)$, with $\mathbf{p}(m) = p_m$, specifies firms' prices, \bar{u} is a threshold utility for consumers, and B is an anonymous belief function for consumers.

An *endogenous dispersion equilibrium* is an assessment satisfying the following conditions:

1. (Firms optimize) For all $m \in [0, 1]$, $(\boldsymbol{\sigma}(m), \mathbf{p}(m))$ maximizes firm m 's profit per consumer visit, given consumer cutoff rule \bar{u} .
2. (Consumers optimize) The cutoff utility \bar{u} maximizes a consumer's utility from search

(net of search costs), given anonymous belief function B .

3. (Aggregate consistency of beliefs) Let P be the the distribution of firms generated by (σ, \mathbf{p}) . Then $B = P$.

Below we discuss the firm and consumer decision problems in more detail. One notable result (Lemma 2) is that a firm will always prefer its product to disperse consumer valuations either as much as possible, or as little as possible. Then we combine the two decision problems and characterize equilibria of the model.

2.3 Consumer's problem

A consumer with beliefs $B(\sigma, p)$ who anticipates that the utility he would receive (net of price) from a purchase made at the next firm he visits is distributed according to

$$G(u) = \Pr(A_\mu + \sigma_m z_{im} - p_m \leq u \mid B(\sigma, p), F(z)) .$$

Consider a consumer whose current best offer in hand, including the option to quit without purchasing and accept 0, is \tilde{u} . (So we will have $\tilde{u} = 0$ for consumers who have not yet searched and consumers who have received only negative utility draws at the firms visited so far.) Suppose this consumer decides to search at one additional firm and then take the best available offer and leave the market. Relative to leaving the market now, this additional search benefits the consumer only if utility at the new firm is strictly greater than \tilde{u} , and costs c regardless. The expected net gain to conducting the additional search is

$$\int_{u > \tilde{u}} (u - \tilde{u}) dG(u) - c$$

Let the utility threshold \bar{u} be defined by

$$\int_{u > \bar{u}} (u - \bar{u}) dG(u) = c \tag{1}$$

Lemma 1 *For any consumer beliefs $B(\sigma, p)$, there is a unique utility threshold \bar{u} satisfying (1).*

Proof. Appendix. ■

For a consumer whose best current utility offer is strictly less than \bar{u} , the expected net gain from an additional search is positive, while a consumer holding an offer better than \bar{u} should take it and leave the market. Notice that this incorporates the participation constraint on search: if the \bar{u} that solves (1) is negative, then taking the best offer in hand, which might be to quit the market without purchasing and earn 0, always dominates continued search. Since the search environment is stationary, this means that if \bar{u} is negative, consumers will not be willing to search at all. (Put slightly differently, if $\bar{u} < 0$ satisfies (1), then $\int_{u \geq 0} u dG(u) < c$, so the net benefit of a single search is negative.) Alternatively, if $\bar{u} > 0$, then the optimal search strategy is to start searching, and accept the first utility offer that is weakly greater than \bar{u} .¹⁰

¹⁰In our model, the assumption that a consumer can costlessly purchase any product visited in the past is inessential – the results would not change if recalling old products were costly. This is true for two reasons.

We can write the expected gross gain from an additional search more explicitly in terms of beliefs B and the distribution of taste shocks F . Note that at a firm with strategy (σ, p) , the condition $u \geq \tilde{u}$ is equivalent to $z \geq \frac{\tilde{u} + p - A_\mu}{\sigma}$. Using this, define

$$\begin{aligned} L(\tilde{u}) &\equiv \int_{u \geq \tilde{u}} (u - \tilde{u}) dG(u) \\ &= \int_{(\sigma, p) \in S_B} \int_{\frac{\tilde{u} + p - A_\mu}{\sigma}}^{\infty} ((A_\mu + \sigma z - p) - \tilde{u}) dF(z) dB(\sigma, p) \end{aligned} \quad (2)$$

Integrating the interior integral by parts gives us the convenient representation

$$\begin{aligned} L(\tilde{u}) &= \int_{(\sigma, p) \in S_B} \sigma I\left(\frac{\tilde{u} + p - A_\mu}{\sigma}\right) dB(\sigma, p) \quad , \text{ where} \\ I(\tilde{z}) &\equiv \int_{\tilde{z}}^{\infty} 1 - F(z) dz \end{aligned} \quad (3)$$

where S_B denotes the support of the consumer's beliefs about (σ, p) . Thus, an alternative characterization of the unique optimal search cutoff (given beliefs B) is

$$\int_{(\sigma, p) \in S_B} \sigma I\left(\frac{\bar{u} + p - A_\mu}{\sigma}\right) dB(\sigma, p) = c \quad (4)$$

The integrand is the utility improvement expected by a consumer (with current best offer \bar{u}) from one additional search, conditional on visiting a type (σ, p) firm. The lefthand side takes a weighted average of these expected improvements over the distribution of firm strategies.

From (4), one can see a straightforward partial equilibrium effect of search costs on consumer behavior. Holding beliefs about firms constant, a fall in c requires an equilibrating decline in the lefthand side of (4), so the threshold utility \bar{u} at which consumers quit searching rises. Remember that $1 - F\left(\frac{\bar{u} + p - A_\mu}{\sigma}\right)$ is the probability of a purchase when a consumer visits a type (σ, p) firm; so as search costs fall and \bar{u} rises, these purchase probabilities decline across the board, indicating choosier behavior by consumers. The effect of a change in beliefs $B(\sigma, p)$ on the best response cutoff \bar{u} is more subtle; we defer an analysis of this until Section 3.

2.4 Firm's problem

Each firm simultaneously chooses a level of dispersion and a price $(\sigma, p) \in S = [\sigma_L, \sigma_H] \times [0, \infty)$ so as to maximize its profit per consumer visit, given the belief that all consumers search according to some common threshold rule \bar{u} .¹¹ The choice of σ can be thought of as one facet of a broader product design process in which the firm decides on a set of features to include

First, because the optimal cutoff strategy is stationary, if a product was not chosen when it was first visited, a consumer will never want to return to it later on. Second (and this is related), because there is a continuum of firms, a consumer is never forced to revisit old products because she has run out of new products to visit. Of course, the distinction between free and costly recall would become more important if the number of products were finite.

¹¹In principle, it would be more general to formulate firm beliefs about consumers as a probability distribution over utility thresholds, rather than assuming that beliefs are concentrated. However, since Lemma 1 establishes that consumers will concentrate on a single \bar{u} , there is no risk of overlooking equilibria by formulating firm beliefs as we do.

in its product.¹² We treat this product design process in reduced form by assuming that the firm controls parameters that affect consumers' willingness to pay for its product. Furthermore, we assume that the firm's product design decisions can be decomposed in terms of 'vertical' features, which affect consumers' mean valuation for the product, and 'horizontal' features that affect how dispersed valuations are around that mean. One could study the vertical component of product design by giving firms an additional choice to increase A_μ (at some cost). Our model shuts down this component of product design in order to focus on the horizontal dimension, but Section 4 sketches an extension that includes both.

By stipulating that $\sigma_L > 0$, we intend to capture the idea that (with the exception of pure commodities) most products have idiosyncrasies that appeal more to some consumers than to others. While a firm can choose to emphasize those idiosyncrasies (higher σ) or to downplay them (lower σ), it cannot eliminate them entirely. There is no cost associated with choosing σ , but of course the limits at σ_L and σ_H can be interpreted as the points at which reducing or increasing idiosyncrasy further becomes prohibitively costly.

After choosing σ and p , the firm can produce its product on demand at zero marginal cost. Firm m 's expected profit per consumer visit if it chooses (σ_m, p_m) is

$$\pi_m = p_m \Pr(u_{im} \geq \bar{u} \mid \sigma_m, p_m)$$

where $\Pr(u_{im} \geq \bar{u} \mid \sigma_m, p_m)$ is the probability that a consumer who arrives at firm m makes a purchase. Given the firm's strategy and its belief about consumers' cutoff rule, this probability is

$$\Pr(u_{im} \geq \bar{u} \mid \sigma_m, p_m) = 1 - F\left(\frac{\bar{u} - A_\mu + p_m}{\sigma_m}\right)$$

Thus, the firm solves

$$\begin{aligned} \max_{(\sigma_m, p_m) \in S} \pi_m(\sigma_m, p_m; \bar{u}) \quad \text{where} \\ \pi_m(\sigma_m, p_m; \bar{u}) = p_m \left(1 - F\left(\frac{\bar{u} - A_\mu + p_m}{\sigma_m}\right)\right) \end{aligned} \quad (5)$$

For firms, the choice of σ and p is simultaneous. However, it is analytically convenient to study the optimization in two steps. First, fix an arbitrary dispersion level σ_m and solve the price-setting problem:

$$\pi_m^*(\sigma_m; \bar{u}) \equiv \max_{p_m \in [0, \infty)} p_m \left(1 - F\left(\frac{\bar{u} - A_\mu + p_m}{\sigma_m}\right)\right) \quad (6)$$

The function $\pi_m^*(\sigma_m; \bar{u})$ identifies the greatest profit that can be achieved at each possible choice of σ_m . Then solve

$$\max_{\sigma_m \in [\sigma_L, \sigma_H]} \pi_m^*(\sigma_m; \bar{u}) \quad (7)$$

¹²In an earlier version of this paper, we show that our reduced form model of product design can be generated by explicitly modeling a product as a bundle of features. Consumer valuations are hedonic over these features, and a consumer's taste for different features is random. A firm can create a product that blends a little bit of a lot of different features. This corresponds to a low σ product – most consumers will have valuations near the mean, as their enthusiasm about some features will be balanced by lukewarm feelings about others. Alternatively (the high σ case), a firm can focus on providing high intensity for one or two features and ignoring others – this will lead to consumer valuations that are more dispersed.

to identify the optimal choice of σ_m .

To guarantee a unique interior solution to the price-setting component of the optimization, we impose the following:

Condition 2 (QC) *The firm profit function (5) is strictly quasiconcave in p_m (for any values of σ_m and \bar{u}).*

Following Caplin and Nalebuff (1991), one can show that logconcavity of $f(z)$ is sufficient (but not necessary) to ensure quasiconcavity of (5) in p_m , so Condition 1 implies (QC). Because we will spend some time later studying taste distributions that satisfy (QC) but are not logconcave, we mention the condition separately now.

Under condition (QC), (6) has a unique maximizing price (for each σ_m and belief \bar{u}) identified by the first order condition:¹³

$$\frac{\partial \pi_m(\sigma_m, p_m; \bar{u})}{\partial p_m} = \left(1 - F\left(\frac{\bar{u} - A_\mu + p_m}{\sigma_m}\right)\right) - \frac{p_m}{\sigma_m} f\left(\frac{\bar{u} - A_\mu + p_m}{\sigma_m}\right) = 0 \quad (8)$$

or $p_m = p(\sigma_m; \bar{u})$, with

$$p(\sigma; \bar{u}) \equiv \sigma_m \frac{1 - F((\bar{u} - A_\mu + p(\sigma; \bar{u}))/\sigma_m)}{f((\bar{u} - A_\mu + p(\sigma; \bar{u}))/\sigma_m)} \quad (9)$$

Next, turn to the choice of σ_m . We present the following result as a lemma.¹⁴

Lemma 2 *For any belief \bar{u} about consumer behavior, the maximizers of (7) form a subset of $\{\sigma_L, \sigma_H\}$. That is, the optimal level of dispersion is always extreme – either σ_L is optimal or σ_H is optimal (or possibly both).*

Proof. Suppose, toward a contradiction, that an interior choice $\hat{\sigma} \in (\sigma_L, \sigma_H)$ were optimal. The maximized profit (over both σ and price) would be $p_m(\hat{\sigma}; \bar{u})(1 - F(\hat{z}))$, where $\hat{z} = \frac{\bar{u} - A_\mu + p_m(\hat{\sigma}; \bar{u})}{\hat{\sigma}}$ is the taste shock of the firm's marginal consumer. This marginal taste shock must be either positive, negative, or zero. If $\hat{z} > 0$, then suppose the firm deviates to σ_H , leaving its price $p_m(\hat{\sigma}; \bar{u})$ unchanged. This reduces the marginal taste shock from \hat{z} to $z' = \frac{\bar{u} - A_\mu + p_m(\hat{\sigma}; \bar{u})}{\sigma_H}$, strictly improving both the chance of a purchase, and expected profit. Alternatively, suppose that $\hat{z} < 0$. In this case the firm could deviate to σ_L , again leaving its price at $p_m(\hat{\sigma}; \bar{u})$. Since $\sigma_L < \hat{\sigma}$ and the numerator of \hat{z} must be negative, this would also reduce the marginal taste shock, so the firm could strictly improve its chance of a purchase and its profit in this case as well. Finally, suppose that $\hat{z} = 0$. Consider a sequence of deviations: first switch the dispersion level to σ_H , leaving the price unchanged. This switch leaves the marginal taste shock at zero, and does not change the firm's expected profit. But note that in this strategy, $(\sigma_H, p_m(\hat{\sigma}; \bar{u}))$, the price is not set optimally. Next adjust the price from $p_m(\hat{\sigma}; \bar{u})$

¹³If f is logconcave, existence and uniqueness of a solution to (??) is guaranteed by the fact that $\frac{1-F(z)}{f(z)}$ is a decreasing function. If (QC) holds but f is not logconcave, use the fact that $p_m(1 - F(\frac{\bar{u} - A_\mu + p_m}{\sigma_m}))$ must tend to zero as $p_m \rightarrow \infty$. (The right tail $1 - F(z)$ must tend to zero faster than $\frac{1}{z}$ since, by assumption, $F(z)$ has a well-defined mean.) Then quasiconcavity implies a unique interior maximum, identified by the first order condition.

¹⁴This result is similar to a result of Johnson and Myatt (2006) in a context without search.

to $p_m(\sigma_H; \bar{u})$; by the strict quasiconcavity of the price-setting problem, this strictly improves expected profit.

In summary, regardless of \hat{z} , the firm can always earn strictly higher profits by using either σ_L or σ_H rather than $\hat{\sigma}$, contradicting the assertion that $\hat{\sigma}$ is optimal. ■

Define $p_H(\bar{u})$ and $p_L(\bar{u})$ to be the optimal price (determined by (9)) for a firm with a product of type σ_H or σ_L that anticipates a consumer cutoff strategy \bar{u} . Let $\pi_H(\bar{u})$ and $\pi_L(\bar{u})$ be the profit earned by a firm with product σ_H and price $p_H(\bar{u})$ (or σ_L and $p_L(\bar{u})$ respectively). By Lemma 2 and (QC) each firm's optimal strategy is simply to set either $(\sigma_L, p_L(\bar{u}))$ or $(\sigma_H, p_H(\bar{u}))$, depending on whether $\pi_L(\bar{u})$ or $\pi_H(\bar{u})$ is larger. We summarize this point formally.

Remark 1 Suppose that consumer threshold \bar{u} is part of an EDE assessment. Condition (QC) and Proposition 2 imply that firms' strategies in this assessment must satisfy $(\sigma_m, p_m) \in \{(\sigma_L, p_L(\bar{u})), (\sigma_H, p_H(\bar{u}))\}$ for all $m \in [0, 1]$.

The fact that firm profits are quasiconvex in σ_m has a fairly straightforward economic intuition. Because of search costs, a firm has temporary monopoly power over a visiting consumer. If we use terminology loosely by referring to 'quantity' when we really mean 'probability of sale,' then the firm essentially acts like a monopolist facing the demand curve [quantity] = $1 - F(\frac{\bar{u} - A_\mu + p_m}{\sigma_m})$. Note that the consumer's outside option \bar{u} acts like a demand shifter here. The firm's product design choice σ_m pivots this demand curve around the quantity $1 - F(0)$. A higher choice of σ_m makes this demand curve more vertical – it tilts out at (high price, low quantity) pairs, and tilts in at (low price, high quantity) pairs. A lower choice of σ_m has the opposite effect. Now suppose the firm has chosen an interior level of σ_m and priced optimally on its demand curve. Then consider shifting this choice of σ_m up or down. One of these two changes must tilt the firm's demand at its current price outward (and the other one shifts demand inward).¹⁵ Thus a firm can always improve its profit by shifting away from its interior level of σ_m , in whichever direction tilts its demand outward.

3 Equilibrium

A convenient implication of Lemma 2, and a corollary to Remark 2.4, is the following.

Remark 2 Suppose that consumer threshold \bar{u} is part of an EDE assessment. Then the consumer beliefs B in this assessment must be concentrated on the two point set $\{(\sigma_L, p_L(\bar{u})), (\sigma_H, p_H(\bar{u}))\}$.

This follows directly from the consistency of beliefs. In other words there cannot be an EDE in which firms choose (and consumers expect to face) more than two distinct (σ, p) pairs – namely, the ones listed in Remark 3. Therefore, from this point forward, without loss of generality, we restrict attention to assessments of the form $\{(\boldsymbol{\lambda}, \mathbf{p}), \bar{u}, (\boldsymbol{\lambda}^e, \mathbf{p}^e)\}$. In this expression, $\boldsymbol{\lambda} = (\lambda_L, \lambda_H)$ denotes the fraction of firms choosing σ_L and σ_H respectively (with $\lambda_L + \lambda_H = 1$), and $\mathbf{p} = (p_L, p_H)$ denotes the price set by a type σ_L or σ_H firm. The consumer threshold \bar{u} is unchanged. Consumer beliefs are now summarized more concisely by $(\boldsymbol{\lambda}^e, \mathbf{p}^e)$,

¹⁵Unless the firm has chosen the pivot quantity $1 - F(0)$.

where $\boldsymbol{\lambda} = (\lambda_L^e, \lambda_H^e)$ denotes consumer beliefs about the fraction of each type of firm, and $\mathbf{p}^e = (p_L^e, p_H^e)$ denotes consumer expectations about the price charged by each type of firm. Such a profile satisfies the definition of an EDE if the following hold.

Firm optimization:

$$\lambda_H \begin{cases} = 0 & \text{if } \pi_L(\bar{u}) > \pi_H(\bar{u}) \\ = 1 & \text{if } \pi_L(\bar{u}) < \pi_H(\bar{u}) \\ \in [0, 1] & \text{if } \pi_L(\bar{u}) = \pi_H(\bar{u}) \end{cases}, \text{ and}$$

A type σ_L or σ_H firm sets the price $p_L(\bar{u})$ or $p_H(\bar{u})$ that solves (9)

Consumer optimization:

$$\lambda_L^e \sigma_L I\left(\frac{\bar{u} + p_L^e - A_\mu}{\sigma_L}\right) + \lambda_H^e \sigma_H I\left(\frac{\bar{u} + p_H^e - A_\mu}{\sigma_H}\right) = c \quad (10)$$

Consistent beliefs:

$$\boldsymbol{\lambda}^e = \boldsymbol{\lambda} \text{ and } \mathbf{p}^e = \mathbf{p} = (p_L(\bar{u}), p_H(\bar{u}))$$

All equilibria are either symmetric – all firms choose the same σ and set the same price – or asymmetric, with a mixture of generic and niche firms. For the latter case, we can write the consumer optimization condition as the pair of conditions:

$$\begin{aligned} \lambda_L^e \sigma_L I(\bar{z}_L) + \lambda_H^e \sigma_H I(\bar{z}_H) &= c \\ A_\mu + \sigma_L \bar{z}_L - p_L^e &= \bar{u} = A_\mu + \sigma_L \bar{z}_H - p_H^e \end{aligned}$$

Written this way, \bar{z}_L and \bar{z}_H identify the minimum acceptable taste shock for a consumer when visiting a generic or a niche firm. The second line ensures that the consumer is holding out for equal utility levels at each type of firm, after adjusting for prices.

The type of equilibrium – generic, niche, or mixed – will depend on how firms' expectations about consumer choosiness affect product choice, and conversely, on how consumers' expectations about the product mix, as well as the search cost c , affect their willingness to search. We examine these in turn.

How does a firm's optimal level of σ depend on its expectation of \bar{u} ?

Proposition 2 answers this question unambiguously: the more selective consumers are expected to be, the stronger the incentives for a firm to switch from a generic to a niche product. First, we introduce the following condition.

Condition 3 (H) *The taste distribution satisfies $h(z) \equiv z - \frac{1-F(z)}{f(z)}$ strictly increasing, with $\lim_{z \rightarrow -\infty} h(z) = -\infty$ and $\lim_{z \rightarrow \infty} h(z) = \infty$*

We will give $h(z)$ an interpretation momentarily; for now we note that (H) is implied by logconcavity of f . Together, the combination of conditions (QC) and (H), which is weaker than logconcavity of f , suffices for most of the results that follow.

Proposition 2 *Suppose that (QC) and (H) hold. Fix A_μ , σ_L , and σ_H . There exists \tilde{u} such that a firm that anticipates a consumer cutoff rule \bar{u} will choose σ_L if $\bar{u} < \tilde{u}$, will choose σ_H if $\bar{u} > \tilde{u}$, and will be indifferent if $\bar{u} = \tilde{u}$.*

Intuitively, fixing \bar{u} , a firm with a generic product σ_L will tend to sell to a larger fraction of visiting consumers than a firm with a niche product σ_H . (That is, $1 - F(z_L(\bar{u})) > 1 - F(z_H(\bar{u}))$.)¹⁶ Suppose that firms' expectation of \bar{u} rises by du . A firm that adjusts to this change by reducing its price by du so as to maintain the same level of sales will endure a profit decline proportional to its sales volume – this hurts the generic firm more than the niche one. Of course, a firm could respond to more selective consumers with a mixture of price and quantity adjustments rather than just a price adjustment, but the envelope theorem implies that the decline in the firm's profit will be the same.

Proposition 2 implies that a mixed equilibrium, with both generic and niche firms, is only possible if consumers' equilibrium cutoff utility is \tilde{u} . If consumers are more or less selective than this in equilibrium, then all firms will be niche or generic, respectively.

How do a consumer's beliefs $(\boldsymbol{\lambda}^e, \mathbf{p}^e)$ about product design and prices affect her optimal utility cutoff \bar{u} ?

Write $\bar{u}(\boldsymbol{\lambda}^e, \mathbf{p}^e; c)$ for the consumer's optimal utility cutoff, given these beliefs and the search cost c . We are interested in the sign of $\frac{d\bar{u}(\boldsymbol{\lambda}^e, \mathbf{p}^e; c)}{d\lambda_H^e}$: does a greater prevalence of niche firms induce consumers to be more or less choosy? It turns out that we can analyze this question by looking at consumer utility in the boundary cases, when all firms have the same σ .

Lemma 3 *Let $\bar{u}_L(\mathbf{p}^e; c) = \bar{u}(\boldsymbol{\lambda}^e, \mathbf{p}^e; c)|_{(\lambda_L^e, \lambda_H^e)=(1,0)}$ and $\bar{u}_H(\mathbf{p}^e; c) = \bar{u}(\boldsymbol{\lambda}^e, \mathbf{p}^e; c)|_{(\lambda_L^e, \lambda_H^e)=(0,1)}$. For any $\boldsymbol{\lambda}^e$, $\frac{d\bar{u}(\boldsymbol{\lambda}^e, \mathbf{p}^e; c)}{d\lambda_H^e}$ has the same sign as $\bar{u}_H(\mathbf{p}^e; c) - \bar{u}_L(\mathbf{p}^e; c)$.*

The logic of the lemma is essentially the following. The expected benefit from an additional search that happens to reach a σ_L or σ_H firm is $\sigma_L I(\bar{z}_L)$ or $\sigma_H I(\bar{z}_H)$ respectively. Equation (10) states that the weighted average of these benefits must equal c ; however $\sigma_L I(\bar{z}_L)$ and $\sigma_H I(\bar{z}_H)$ need not be (and generally, will not be) equal to each other. If $\sigma_H I(\bar{z}_H) > c > \sigma_L I(\bar{z}_L)$, then an increase in the fraction of idiosyncratic firms improves the overall expected benefit from search, inducing the consumer to hold out for a higher \bar{u} . It turns out that that whenever this is true, we also have $\bar{u}_H(\mathbf{p}^e; c) > \bar{u}(\boldsymbol{\lambda}^e, \mathbf{p}^e; c) > \bar{u}_L(\mathbf{p}^e; c)$.

In order to compare consumer utility with all σ_L firms pricing at p_L^e , versus all σ_H firms pricing at p_H^e , it is useful to introduce an auxiliary function $v(\sigma; c)$ defined by

$$\sigma I\left(\frac{v(\sigma; c) - A_\mu}{\sigma}\right) = c \quad (11)$$

This function corresponds to the optimal consumer cutoff under the assumption that all firms choose strategy $(\sigma, 0)$. We are not interested a situation with zero prices *per se*, but $v(\sigma; c)$ provides a convenient way to express consumer utility in our two cases of interest. Given the implicit definitions of $\bar{u}_L(\mathbf{p}^e; c)$, $\bar{u}_H(\mathbf{p}^e; c)$, and $v(\sigma; c)$, we have $\bar{u}_L(\mathbf{p}^e; c) = v(\sigma_L; c) - p_L^e$ and $\bar{u}_H(\mathbf{p}^e; c) = v(\sigma_H; c) - p_H^e$. This means that the change in consumer search induced by a shift from a generic to a niche market can be characterized by

$$\bar{u}_H(\mathbf{p}^e; c) - \bar{u}_L(\mathbf{p}^e; c) = [v(\sigma_H; c) - v(\sigma_L; c)] - [p_H^e - p_L^e]$$

The difference between a consumer's choosiness when she expects a niche versus a generic market can be separated into a term related to product design σ and a price term. Lemma 4

¹⁶This is intuitive but not self-evident – Lemma 3 provides sufficient conditions for it to be true.

shows that the first term is unambiguously positive: if consumers expect the same price when all firms choose σ_H as they do when all firms choose σ_L , then they will hold out for higher utility in the nichier market. Moreover, if the upper bound σ_H on the nichiness of products increases, this utility gain grows at an increasing rate.

Lemma 4 *The function $v(\sigma; c)$ is strictly increasing and strictly convex in σ . Consequently, $v(\sigma_H; c) - v(\sigma_L; c) > 0$.*

Proof. For $v(\sigma; c)$ strictly increasing, let $z(\sigma; c)$ be defined by $\sigma I(z(\sigma; c)) = c$, so that $v(\sigma; c) = A_\mu + \sigma z(\sigma; c)$, and differentiate to get¹⁷

$$\begin{aligned} \frac{\partial v(\sigma; c)}{\partial \sigma} &= z(\sigma; c) + \frac{I(z(\sigma; c))}{1 - F(z(\sigma; c))} \\ &= z(\sigma; c) + E(z - z(\sigma; c) \mid z > z(\sigma; c)) \\ &= E(z \mid z > z(\sigma; c)) > E(z) = 0. \end{aligned}$$

For convexity, see the appendix. ■

The logic is essentially an option value argument. Increasing σ does not change the expected value $E(A_\mu + \sigma z)$ of a random consumer-firm match, but it does make good matches better and bad matches worse. However, a consumer is not affected by the deterioration of very bad matches ($z < z(\sigma; c)$) because she would not have accepted them anyway. Thus, the average effect of higher σ on the truncated set of matches she does accept is positive.

Together, Lemmas 3 and 4 also imply that a consumer who expects a mixture of σ_L and σ_H firms will unambiguously benefit if the share of σ_H firms rises, as long as σ_H firms are not expected to charge higher prices than σ_L firms. On the other hand, if $p_H^e > p_L^e$, it is not clear whether or not a shift toward more idiosyncratic firms benefits consumers or not – the answer depends in some sense on what share of the gains $v(\sigma_H; c) - v(\sigma_L; c)$ that firms are able to extract through higher prices. We will come back to this point later.

We can sum up the main partial equilibrium conclusions thus far as follows. Lower search costs encourage consumers to search longer and hold out for a higher cutoff utility \bar{u} . Firms that expect to face more discriminating consumers (higher \bar{u}) are induced to switch from generic (σ_L) to more idiosyncratic (σ_H) products. Consumers who expect more idiosyncratic products are in turn induced to hold out for even higher utility as long as the price premium demanded for those idiosyncratic products is not too large. We are now prepared to characterize equilibria of the model. Define a function $U(c, \sigma) = A_\mu + \sigma h(I^{-1}(\frac{c}{\sigma}))$.

Lemma 5 *$U(c, \sigma)$ is strictly decreasing in c , with $\lim_{c \rightarrow 0} U(c, \sigma) = \infty$ and $\lim_{c \rightarrow \infty} U(c, \sigma) = -\infty$. If an EDE exists in which all firms choose $\sigma = \sigma_S$, then the consumer utility cutoff in that EDE must be $U(c, \sigma_S)$.*

Proof. I is strictly decreasing and thus invertible. The monotonicity of U follows from I strictly decreasing and h strictly increasing. The limits follow from (H) and the fact that $I^{-1}(\frac{c}{\sigma})$ tends to ∞ or $-\infty$ as $c \rightarrow 0$ or $c \rightarrow \infty$ respectively. If the stipulated equilibrium exists, then applying (10) and (9), it must satisfy $\sigma_S I(\bar{z}) = c$ and $p = \sigma_S \frac{1 - F(\bar{z})}{f(\bar{z})}$, for $\bar{z} =$

¹⁷For the second line, reverse the integration by parts from Section 2.

$\frac{\bar{u}+p-A_\mu}{\sigma_S}$. The first equation requires that $\bar{z} = I^{-1}\left(\frac{c}{\sigma_S}\right)$. Using the equation for p , we have $\bar{u} = A_\mu + \sigma_S \left(\bar{z} - \frac{1-F(\bar{z})}{f(\bar{z})}\right) = A_\mu + \sigma_S h(\bar{z})$. Together, these show that equilibrium utility must be $U(c, \sigma_S)$. ■

Lemma 5 is of interest principally for the cases $\sigma_S = \sigma_L$ or σ_H , since we know that no other level of σ can be optimal for a firm. The next proposition characterizes equilibria of the model in terms of threshold values of $U(c, \sigma_L)$ and $U(c, \sigma_H)$. This highlights one of the technical difficulties introduced by the fact that optimal product design is extreme: the nature of equilibrium can depend on discrete differences between these two functions that are difficult to characterize in much generality. If an interior choice of σ had turned out to be optimal for firms, then the nature of the equilibrium would be related to a first order condition of $U(c, \sigma)$ with respect to σ , and this might lend itself to more definitive conclusions.

Proposition 3 *Fix A_μ , σ_L , and σ_H . Let \tilde{u} be the threshold utility, as described in Proposition 2, that makes a firm indifferent between σ_L and σ_H .*

- i) If $\tilde{u} < 0$, then*
 - a) Define c_1 by $U(c_1, \sigma_H) = 0$. An EDE with all firms choosing σ_H and consumer cutoff utility $U(c, \sigma_H) \geq 0$ exists iff $c \in (0, c_1]$.*
 - b) No other EDE exists for any value of c .*
- ii) If $\tilde{u} \geq 0$, then*
 - a) Define c_2 by $U(c_2, \sigma_H) = \tilde{u}$. An EDE with all firms choosing σ_H and consumer cutoff utility $U(c, \sigma_H) \geq \tilde{u}$ exists iff $c \in (0, c_2]$.*
 - b) Define c_3 and c_4 by $U(c_3, \sigma_L) = \tilde{u}$ and $U(c_4, \sigma_L) = 0$, with $c_3 \leq c_4$. An EDE with all firms choosing σ_L and consumer cutoff utility $U(c, \sigma_L) \leq \tilde{u}$ exists iff $c \in [c_3, c_4]$.*
 - c) An asymmetric EDE with some firms choosing σ_L , the remaining firms choosing σ_H , and consumer cutoff utility $\bar{u} = \tilde{u}$ exists if $c \in (\min(c_2, c_3), \max(c_2, c_3))$.*
 - d) No other EDE exists for any value of c .*
- iii) The cost thresholds c_1 and c_4 rise with an increase in the mean value of a product A_μ . However, thresholds c_2 and c_3 are invariant to A_μ .*

A few remarks may help to illuminate this characterization of equilibrium. First, note that for small enough search costs ($c \leq c_1$ if $\tilde{u} < 0$, or $c < \min(c_2, c_3)$ if $\tilde{u} \geq 0$), there is a unique equilibrium with all firms choosing niche products. This appears to reflect the feedback loop described above: cheap search encourages choosy consumers; this turns firms toward niche products which reinforces consumer choosiness. Second, consider large search costs. For c large enough ($c > c_1$ if $\tilde{u} < 0$ or $c > c_4$ if $\tilde{u} > 0$), no equilibrium with search exists because consumers do not anticipate realizing enough surplus from a purchase to cover their costs. The higher the mean value of a product A_μ , then (by part (iii)), the larger the range of costs c for which an equilibrium with search can be sustained. This is because the strength of competition (as measured inversely by search frictions), not A_μ , is the binding constraint on how much surplus a firm can extract from consumers *via* a higher price. If A_μ is sufficiently large (so that $c_4 > c_2$), then for moderately high search costs $c \in (\max(c_2, c_3), c_4)$, there is a unique equilibrium with all firms choosing generic products. This appears to be the flip side of the feedback loop discussed above: high c discourages consumer search, given this, firms prefer to offer generic products, and this discourages search further. In these interpretations, “appears”

is intended as a reminder that one segment of this loop – higher σ encouraging choosier search – only applies if prices do not rise too fast with σ , and we have not showed that this must happen.

For intermediate search costs between c_2 and c_3 , there are three main possibilities. (To simplify the discussion, assume A_μ large enough ($c_4 > c_2$) to be sure that an equilibrium exists.) If $c_2 > c_3$, then for $c \in (c_3, c_2)$ there are three EDEs: one nichy, one generic, and an asymmetric EDE with both σ_L and σ_H firms. In this case, the feedback loop does apply: it is what generates the multiplicity. In the σ_L equilibrium, consumers are not choosy enough to make it worthwhile for any single firm to deviate to σ_H . However, if consumers expect all firms to offer σ_H products, they will be sufficiently more discriminating that σ_H becomes optimal for firms. The asymmetric EDE in this case is unconvincing, as it fails standard notions of stability.¹⁸

Alternatively, if $c_2 < c_3$, then for search costs in the range $c \in (c_2, c_3)$ there is a unique, and asymmetric, equilibrium. If and when this case applies, the weak link in our feedback story must fail. That is, if consumers expect all products to be generic, they are still choosy enough that firms would prefer to deviate to σ_H and price substantially higher. However, if consumers expected all firms to shift to σ_H , the price hike would more than wipe out the potential utility gains described in Lemma 4, inducing consumers to settle for lower utility – low enough that firms would want to switch back to σ_L .

The third possibility is the knife edge case $c_2 = c_3$. In this case, as search costs fall from c_4 down to zero, the market shifts immediately from a unique generic equilibrium for $c > c_2 = c_3$ to a unique nichy equilibrium for $c < c_2 = c_3$. Given a distribution F and the other parameters, one can compute c_2 and c_3 (either analytically or numerically) to determine which of these three regimes applies; however there does not appear to be any straightforward and general classification. Computed examples demonstrate that all three regimes (multiple equilibrium, a unique asymmetric equilibrium, or the knife-edge transition) are possible. *A priori* one might imagine, based perhaps on assumptions about genericity, that the first two regimes would apply to more distributions and parameters than the third. Somewhat surprisingly, that is not necessarily true – many common distributions generate the knife edge case. The reason why has to do with how firms and consumers split the change in total surplus generated by shifting to more idiosyncratic products. It turns out that for these distributions, the split is such that consumer surplus and firm profits ‘agree’ about whether shifting to more idiosyncratic products is advantageous.

Definition 3 *The taste distribution $f(z)$ is inverse hazard rate invariant (IHRI) if its inverse hazard rate $v(z) = \frac{1-F(z)}{f(z)}$ satisfies the following condition for some positive constant K :*

$$E(v(z) \mid z \geq \bar{z}) = Kv(\bar{z})$$

This condition says the ratio of the inverse hazard rate’s value at a point to its average value to the right of that point does not depend on which point we choose. Distributions satisfying the IHRI condition include everything in the generalized Pareto family, such as the

¹⁸Informally, if consumers expected a slightly higher fraction of σ_H firms, their optimal cutoff utility would rise above \tilde{u} ; encouraging the shift toward σ_H . Similarly, a rise in the fraction of σ_L firms would tend to snowball toward the all- σ_L equilibrium.

uniform distribution, distributions with triangular densities, the exponential distribution, and distributions with right tails that obey power laws.¹⁹ If tastes are inverse hazard rate invariant, we have the following sharp result.

Proposition 4 *If $f(z)$ is inverse hazard rate invariant, then $c_2 = c_3$. That is, if an EDE exists, then it is generically unique and involves minimal differentiation by all firms for $c > c_2 = c_3$ and maximal differentiation by all firms for $c < c_2 = c_3$.*

While the following result is helpful in proving Proposition 6, it is also interesting in its own right.

Proposition 5 *If $f(z)$ is inverse hazard rate invariant with constant K , then firm profits in any equilibrium are equal to $\frac{c}{K}$.*

The next proposition clarifies the way in which consumers' and firms' (equilibrium) preferences about σ are aligned for an IHRI distribution, thereby generating the shift from a unique σ_L equilibrium to a unique σ_H equilibrium as c falls. Call the threshold cost $\bar{c} = c_2 = c_3$.

Proposition 6 *If $f(z)$ is inverse hazard rate invariant, then $U(c, \sigma_H) > U(c, \sigma_L) > \tilde{u}$ for $c < \bar{c}$ and $U(c, \sigma_H) < U(c, \sigma_L) < \tilde{u}$ for $c > \bar{c}$. Therefore, for any search cost c , consumers' EDE utility attains the maximum of $\{U(c, \sigma_H), U(c, \sigma_L)\}$.*

Proposition 6 tells us that given optimal pricing by firms, consumers prefer an all σ_H market over an all σ_L if and only if their cutoff utility is such that firms prefer σ_H too. (And similarly for σ_L .) What is it about the taste distribution that aligns firm and consumer interests in this way? Notice that the expectation over the right tail of $v(z)$ can be rewritten (integrating by parts) as follows:

$$E(v(z) | z \geq \bar{z}) = E(z - \bar{z} | z \geq \bar{z})$$

The righthand expression is related to the surplus that the *average* consumer receives when she finally buys, relative to the *marginal* consumer who buys. In this sense, consumer surplus is linked to the area under $v(z)$ to the right of \bar{z} . Meanwhile, a firm's optimal price is proportional to $v(\bar{z}) = \frac{1-F(\bar{z})}{f(\bar{z})}$. Thus the relationship between firm and consumer surplus is connected to the relationship between $v(\bar{z})$ and its right tail. If this relationship does not depend on \bar{z} , then when an individual firm changes σ (thereby shifting its marginal consumer), the change will tend to move its own profits and utility in lockstep.²⁰ Proposition 7 states that even though individual firms' gains from shifting σ may be aligned with improvements in consumer utility, those profit gains are fully dissipated when all firms follow suit.

When the taste distribution does not satisfy inverse hazard rate invariance, Proposition 4 can still offer intuition about whether to expect multiple symmetric equilibria or a single asymmetric one for intermediate search costs. Shifting from generic to niche products always tends to push consumers further into the right tail of the taste distribution. If $\frac{E(v(z)|z \geq \bar{z})}{v(\bar{z})}$ increases with \bar{z} , then based on the arguments above, as search costs decline, we might expect consumers to

¹⁹Of course, for distributions of the form $F(z) = 1 - z^{-a}$, we need the expectation in the definition to converge, so we must have $a > 1$.

²⁰This is quite informal; readers should consult the proof of Proposition 4 for a more rigorous statement of this connection.

begin to favor niche products earlier (for higher c) before it becomes more profitable for any individual firm to offer one. (That is, there may be c such that $U(c, \sigma_L) < U(c, \sigma_H) < \tilde{u}$.) In this case, multiple equilibria would arise. Alternatively, if $\frac{E(v(z) | z \geq \bar{z})}{v(\bar{z})}$ declines with \bar{z} , then firms may be more eager than consumers to switch to σ_H , because they are able to extract a larger share of the surplus from sales to consumers who are further into the right tail of the taste distribution. In this case, one could see asymmetric equilibria for intermediate search costs. For most commonly used distributions, the ratio $\frac{E(v(z) | z \geq \bar{z})}{v(\bar{z})}$ is either constant or increasing in \bar{z} . For example, the ratio is increasing for distributions often used in discrete choice settings such as the normal, logistic, and generalized extreme value distributions. On the other hand, it is rather difficult to construct distributions for which this ratio is decreasing.²¹

3.1 Efficiency

Neither consumers nor firms fully internalize the effects of their choices. Consumers do not account for the fact that by searching more assiduously they tend to encourage more idiosyncratic products, which may benefit other consumers. Meanwhile, individual firms do not consider how, collectively, their product design decisions induce changes in consumer selectivity that feed back into profits. Furthermore, the search frictions shift bargaining power from consumers to firms by depressing the consumer's outside option. In view of these distortions, it is perhaps surprising that under some circumstances the search equilibrium will turn out to be constrained efficient.

Throughout this section, we assume A_μ large enough that it is optimal to participate in search. Define total social surplus as the sum of consumer utility and firm profits, all measured on a per-consumer basis.

Proposition 7 *An assessment maximizes the total social surplus if i) all firms choose σ_H , ii) prices form a best response to consumer cutoff utility, and iii) consumer cutoff utility is a best response to σ_H and firms' prices. Therefore, for small search costs, the unique EDE is constrained efficient.*

To understand this result, consider the potential sources of inefficiency: distortions of prices, product design σ , or consumer search intensity away from their efficient levels. Because of the unit demand assumption, prices simply transfer surplus between consumers and firms; they cause no quantity distortions unless they are so high that they deter consumers from searching. Efficiency favors idiosyncratic products for reasons similar to Lemma 4: higher dispersion makes good firm-consumer matches better and bad matches worse, but the latter can be rejected in favor of continued search. As for consumer search, a consumer's objective function differs from a social planner's because the former considers not just the quality of a match, but also its price in deciding whether to buy or keep searching. However, because an all σ_H equilibrium involves no price dispersion across firms, the fact that consumers care about price does not distort their search process away from what a social planner would choose.

²¹Interested readers can consult our working paper for examples that fit this intuition. For multiple equilibria, $F(z) = \frac{e^{kz}}{1+e^{kz}}$ will work. For a unique, asymmetric equilibrium simple examples are hard to find, but the following will work: $F(z) = \tilde{F}(z + e^{-1})$ with support on $z \in [-e^{-1}, 1 - e^{-1}]$, where $\tilde{F}(x) = 1 - (1 - x)e^{-x}$ for $x \in [0, 1]$. Other distributions that generate the asymmetric equilibrium case are similarly labored, suggesting that this case should not be viewed as a common outcome in our model.

Along with the fact that equilibria with generic products will be constrained inefficient, this result suggests a new twist to the standard explanation of how a market with friction approaches a competitive benchmark as that friction shrinks. As search frictions diminish, our search equilibrium approaches a full information competitive equilibrium through two channels. The standard channel might be considered technological: as search costs decrease, the constrained optimal allocation of goods approaches the true optimum. The new channel, as outlined by Proposition 3, is more strategic: as consumer information improves, product differentiation endogenously shifts toward its efficient level. This also adds a nuance to results for homogeneous goods that predict that improvements in consumer information will drive prices down to marginal cost: here prices do not fall because differentiation provides an escape valve for the pressure applied by better information.

This result would be weaker if there were more sources of inefficiency in the model. For example, suppose that rather than unit demand, each consumer has a downward sloping demand curve.²² In this case, an equilibrium price above marginal cost would generate the standard downward distortion of quantities purchased. If it were true in this revised model (as is often the case in our current model) that prices are higher in a niche equilibrium than in a generic one, then a social surplus comparison could involve competing effects. A consumer might end up with a better-suited product in the niche equilibrium, but buy relatively less of it (due to the price distortion) than she would have had the equilibrium been generic. Other factors that could work against idiosyncratic products are costs for unsold inventory, or the possibility for a consumer of running out of options, if the number of firms were finite.

4 Extensions

In order to present the main results of the paper with as much clarity as possible, our model stripped away many elements of the product design process for firms. In this section, we discuss how two of these elements – quality improvements and multi-product firms, could be incorporated back into the model.

Firms that control both the mean and variance of consumer valuations

Our model focuses on product design choices that affect the dispersion of consumers' valuations for a product around a mean A_μ . We have set aside what may seem to be more basic questions: how much should a firm invest in improving the *average* value of its product to consumers, and how does consumer search affect this decision? We briefly sketch a version of the model that addresses these questions. Suppose that when a firm chooses σ and p , it also chooses $A_\mu \in [0, \infty)$. Just as with σ and p , consumers do not observe the values of A_μ that firms choose before searching, but consumers do form beliefs about A_μ that must be correct in equilibrium. One can interpret the pair (A_μ, σ) as a decomposition of the firm's product into a quality dimension (A_μ) about which consumers agree, and a taste dimension (σ) on which opinions differ. For example, for a car, higher σ might represent bolder styling that some consumers love and others hate, while higher A_μ might represent better fuel economy.²³ Suppose the cost of supplying a product of mean value A_μ is $w(A_\mu)$ per unit, with $w(0) = w'(0) = 0$

²²I thank a referee for bringing this point to my attention.

²³The assumption that a firm can choose A_μ and σ independently is a convenient fiction; in practice many changes to a product will probably affect A_μ and σ simultaneously.

and w strictly convex. A firm that anticipates consumers searching according to cutoff utility \bar{u} faces the revised profit maximization problem:

$$\max_{(\sigma, A_\mu, p)} (p_m - w(A_\mu)) \left(1 - F\left(\frac{\bar{u} - A_\mu + p}{\sigma}\right)\right).$$

The logic of Lemma 2 extends, so the firm's optimal level of σ will still be σ_L or σ_H . More interesting is the interplay between A_μ and p , which substitute one-for-one for one another in the firm's probability of a sale. Suppose the firm considers raising its price by a small amount ε and 'sterilizing' the price increase by also improving A_μ by ε . This change leaves the probability of a sale unaffected, and changes the profit margin on each unit sold by $\varepsilon - w'(A_\mu)\varepsilon$. If $w'(A_\mu) < 1$, this change improves profits, while if $w'(A_\mu) > 1$, a reduction in both p and A_μ would improve profits. Thus the optimal strategy for the firm must satisfy $w'(A_\mu) = 1$, but this pins down the firm's optimal quality level A_μ .²⁴ Importantly, the firm's optimal choice of A_μ is determined entirely by the marginal cost of improving quality; the expected consumer cutoff utility \bar{u} , and the firm's choices about σ and p do not affect this at all. Suppose A_μ^* is this optimal quality level, with $w'(A_\mu^*) = 1$. The model in which firms choose A_μ endogenously can effectively be treated as a special case of the original model in which A_μ is fixed exogenously at A_μ^* and firms face a production cost of $w(A_\mu^*)$ per unit. (While the original model was presented using a zero unit cost for firms, none of the results hinged on this assumption.)

As a result, it is perfectly acceptable to interpret the main results of the paper as applying to a situation in which product design affects both the mean and the variance of consumer values for a firm's product. We should note that the simplicity of the role played by the endogenous choice of A_μ would not necessarily carry through under alternative modeling assumptions. For example, in the case of the car, one might imagine that improving fuel economy entails not just per unit costs, but also some fixed costs such as research. We do not formally analyze the case in which choosing A_μ incurs a cost regardless of the number of units sold, but one might conjecture that in this case, a firm will tend to invest in higher A_μ when it is planning to produce a relatively generic (low σ) product that it expects to sell in relatively high quantities.

Multi-product firms

In practice, many firms offer not just one product, but a line of several related products. Part of the appeal of offering a broad product line may be that it allows a firm to hedge against uncertain consumer tastes – a consumer only needs to like at least one of the firm's products to make a sale. Furthermore, this hedging effect disproportionately helps firms that sell very polarizing (high σ) products, since these are precisely the firms that fail to convert most of their visitors to sales in the single product case. Thus, one might expect that in a model with multi-product firms, the market switches over from generic to niche products earlier (than in the single product model) as search costs fall.

To examine this logic, we sketch a version of the model in which each firm can offer up to N different products. For simplicity, we make the following assumptions. 1) A firm must set the same level of σ , and the same price p , for all of its products. Thus a firm strategy is a pair (σ, p) as before, plus a number of products $n \in \{1, \dots, N\}$.²⁵ 2) Each consumer draws a separate, independent taste shock z for each product offered by a firm. 3) Consumer search is

²⁴This fact can be seen easily, and a bit more formally, from the first order conditions for p and A_μ .

²⁵We do not consider the endogenous choice of A_μ here.

as described earlier *except* that when a consumer arrives at a firm, she learns her valuations for all of the products it offers (at no additional cost). She can buy any one of them or continue her search at another firm. She does not observe a firm’s choices of n before visiting it.²⁶ If she continues to search, she visits another randomly chosen firm.²⁷ 4) A consumer still demands at most one product. While these assumptions are restrictive – particularly Assumption 1, which rules out the possibility of a single firm offering both generic and niche products – they have the benefit of allowing us to frame the extension in terms of earlier results.

For additional simplicity, suppose that there are no additional costs associated with offering more than one product. In this case, we argue that it is self evident that each firm will choose $n = N$. For a visiting consumer, there is only one of a firm’s N products that she might actually purchase – that is the product for which she has the largest taste shock. Since a single taste shock has distribution $F(z)$, the largest of N independent taste shocks has distribution $F(z)^N$. Thus we can reinterpret this multiple product model as a version of our standard single product model in which the taste shock of a consumer for a firm has distribution $F(z)^N$ rather than $F(z)$. This distribution no longer has mean zero, so our first definition of differentiation (Proposition 1) no longer applies, but all of the other results carry through unchanged.²⁸ In particular, Proposition 2 applies: there is a threshold utility, which we denote $\tilde{u}(N)$, such that a firm will be indifferent between offering N relatively generic products ($\sigma = \sigma_L$) versus offering N idiosyncratic products ($\sigma = \sigma_H$) if it expects consumers to use the utility cutoff $\tilde{u}(N)$. A firm that expects consumers to demand more (less) utility than \tilde{u}_N will strictly prefer σ_H (σ_L). The intuition sketched above suggests that $\tilde{u}(N)$ should decline with N : a firm with a larger product line should be emboldened to offer nichier products. Proposition 8 confirms this intuition.

Proposition 8 *Suppose that $F(z)^N$ satisfies Condition 2 for all $N \geq 1$ and let $\tilde{u}(N)$ be the expected consumer cutoff utility at which a firm with N products is indifferent between choosing σ_L and σ_H . Then $\tilde{u}(N)$ is decreasing in N .*

A reasonable conjecture would be that the range of search costs that can support a high differentiation equilibrium grows with the number of products per firm, but we have not proved this. One might also wonder how multiple product firms affect the equilibrium regime (multiple equilibria or one asymmetric equilibrium) that prevails at intermediate search costs. This is difficult to determine analytically, but based on Section 3, we conjecture that the right tail behavior of the inverse hazard rate $v_N(z) = \frac{1-F(z)^N}{Nf(z)F(z)^{N-1}}$ may offer some insight. Numerical investigation suggests that the slope of $\frac{E(v_N(z)|z \geq \bar{z})}{v_N(\bar{z})}$ with respect to \bar{z} tends to increase with N for common distributions. For example, if $v(z)$ satisfies IHRI, then $\frac{E(v_N(z)|z \geq \bar{z})}{v_N(\bar{z})}$ tends to

²⁶Of course, in equilibrium, a consumer will form correct beliefs about the distribution of n across firms.

²⁷This implies that a multi-product firm and a single product firm are equally likely to be visited. One alternative version of random search would be for the consumer to pick a product randomly, and then visit the firm that produces it, observing the firm’s entire product line. In this version, multi-product firms would receive more consumer visits, augmenting the incentive for a firm to offer multiple products. As an example of this type of search, imagine a consumer who runs an online keyword search that reveals links to every available product. She clicks through on one product link randomly, and this leads her to its manufacturer’s web site, with information on the manufacturer’s entire product line.

²⁸Let $F_N(z) = (F(z))^N$, $v_N(z) = \frac{1-F_N(z)}{f_N(z)}$, and $h_N(z) = z - v_N(z)$. For earlier results to go through, F_N must satisfy Condition 2. This is straightforward if F is logconcave: F_N inherits the logconcavity of F , and this implies Condition 2. Otherwise, if F is not logconcave, then Condition 2 must be verified for F_N directly.

be increasing. Based on the arguments in Section 3, we speculate that this will tend to make multiple equilibria more common as the number of products per firm increases. A formal investigation of this is left to future work.

5 Concluding Remarks

In many ways, learning about differentiated products is a better motivation for the sequential search model than learning about prices is. With the advent of new information-gathering tools like the internet, for many goods it is difficult to justify the assumption that prices are hard to come by. It is the idiosyncrasies of those goods that make a persuasive case for search. For example, discovering whether one likes the scent of a new perfume, the tradeoff between portability and cramped typing on a small laptop computer, or the prose style of a new author generally requires taking some time to investigate the product. This may help to suggest why Brynjolfsson and Smith (2000), among others, find that pricing on the internet is not as competitive as the easy availability of price information might suggest; they suggest that the time costs of assessing fine print differentiated features like shipping and return policies or of testing the benefits of site personalizations provide real barriers to competition.²⁹ Our paper takes a first step toward understanding how firms take consumer search into account when designing products.

In order to focus on the interaction between differentiation and consumer search, the model presented here excludes many other issues that it would be desirable to incorporate into future analysis. As mentioned earlier, adding more channels for disseminating product information (such as advertising), and giving consumers the opportunity to use publicly available information to perform directed (rather than undirected) search, are two natural extensions. Based on the analysis here, we would conjecture that to the extent that both of these activities improve consumer information, they may tend to encourage greater product differentiation. However, as suggested by recent work (Bagwell and Ramey (1996), Anderson and Renault (2006)), the signaling introduced by advertising can have subtle effects, and details (such as whether prices are advertised, or product information, or both) may play a critical role. Understanding the endogenous relationship between how a product is designed (i.e., differentiation) and how it is marketed (price or product advertising) would be of great interest. We have also suppressed any discussion of product quality, but incorporating choices over vertical differentiation would be another logical step. The non-spatial taste shock model of differentiation used, and the assumption of a continuum of firms, exclude some interesting oligopoly issues, such as whether a firm should design its product to go head-to-head with some rival products while keeping its distance from others. Extensions that allow for this type of “selective crowding” would also be welcome.

²⁹Similarly, Hortaçsu and Syverson (2004) find that vertical product differentiation and search costs can explain substantial price dispersion in the seemingly homogeneous market for S&P 500 index funds. In our horizontal differentiation setting, soft competition is manifested in the level of prices rather than dispersion.

6 Proofs

Proof of Proposition 1

We have

$$\begin{aligned}
D_{mn}(k) &= \Pr(|\sigma_m z_m - \sigma_n z_n| \leq k) \\
&= \int_{-\infty}^{\infty} \Pr(z_n \in [\frac{\sigma_m z_m - k}{\sigma_n}, \frac{\sigma_m z_m + k}{\sigma_n}]) f(z_m) dz_m \\
&= \int_{-\infty}^{\infty} \{F(\frac{\sigma_m z_m + k}{\sigma_n}) - F(\frac{\sigma_m z_m - k}{\sigma_n})\} f(z_m) dz_m
\end{aligned}$$

To show both directions of the if and only if, it suffices to show that $D_{mn}(k)$ is strictly decreasing in σ_m for all k . Differentiating with respect to σ_m :

$$\frac{dD_{mn}(k)}{d\sigma_m} = \frac{1}{\sigma_n} \int_{-\infty}^{\infty} z_m \{f(\frac{\sigma_m z_m + k}{\sigma_n}) - f(\frac{\sigma_m z_m - k}{\sigma_n})\} f(z_m) dz_m$$

Note that $f(z)$ has a single peak at zero. (Logconcavity implies quasiconcavity, and symmetry about a zero mean tells us where the peak is.) It follows directly that the term in brackets is weakly negative when z_m is positive, and *vice versa*. Formally, if z_m is positive, then

$$\frac{\sigma_m z_m - k}{\sigma_n} \in [-\frac{\sigma_m z_m + k}{\sigma_n}, \frac{\sigma_m z_m + k}{\sigma_n}]$$

so by quasiconcavity

$$f(\frac{\sigma_m z_m - k}{\sigma_n}) \geq \min\{f(-\frac{\sigma_m z_m + k}{\sigma_n}), f(\frac{\sigma_m z_m + k}{\sigma_n})\} = f(\frac{\sigma_m z_m + k}{\sigma_n})$$

and similarly for z_m negative. It follows that the integrand is everywhere weakly negative. Furthermore the integrand must be strictly negative on a set of positive measure, so the derivative is strictly negative. The result follows.

Proof of Lemma 1

Following the text, we seek to show that $L(\bar{u}) = c$ has a unique solution, where the function L is defined in (2) and characterized in (3). First, we observe that $L(\bar{u})$ is differentiable with $L'(\bar{u}) < 0$. We have:

$$L'(\bar{u}) = - \int_{(\sigma, p) \in S_B} I\left(\frac{\bar{u} + p + A_\mu}{\sigma}\right) dB(\sigma, p)$$

The integrand is strictly positive for all (σ, p) , so $L'(\bar{u}) < 0$ as claimed. Next, note that $L(\bar{u}) < c$ for \bar{u} sufficiently large. To show this, note that for $\bar{u} > 0$, we have $\frac{\bar{u} + p + A_\mu}{\sigma} \geq \frac{\bar{u}}{\sigma_H}$ for all (σ, p) , so

$$\begin{aligned}
L(\bar{u}) &\leq \int_{(\sigma, p) \in S_B} \sigma_H I\left(\frac{\bar{u}}{\sigma_H}\right) dB(\sigma, p) \\
&= \sigma_H I\left(\frac{\bar{u}}{\sigma_H}\right)
\end{aligned}$$

But by choosing \bar{u} large enough, we can make the righthand side arbitrarily small, so $L(\bar{u}) < c$ for \bar{u} large enough.

Finally, we claim that $L(\bar{u}) > c$ for \bar{u} sufficiently negative. Let $\bar{p} = \sup \{p \mid B(\sigma_H, p) < 1\} < \infty$ be an upper bound on the prices the consumer expects to see. Let \hat{z} satisfy $F(\hat{z}) = \frac{1}{2}$, and choose \bar{u} small enough that $\max \left\{ \frac{\bar{u} + \bar{p} + A_\mu}{\sigma_L}, \frac{\bar{u} + \bar{p} + A_\mu}{\sigma_H} \right\} < \hat{z} - \frac{4c}{\sigma_L}$. By construction, we have $\left(\hat{z} - \frac{\bar{u} + \bar{p} + A_\mu}{\sigma} \right) > \frac{4c}{\sigma_L}$ for all $(\sigma, p) \in \text{supp}(B)$. Then,

$$\begin{aligned} L(\bar{u}) &\geq \int_{(\sigma, p) \in S_B} \sigma_L \int_{\frac{\bar{u} + \bar{p} + A_\mu}{\sigma}}^{\hat{z}} (1 - F(z)) dz dB(\sigma, p) \\ &\geq \int_{(\sigma, p) \in S_B} \sigma_L \int_{\frac{\bar{u} + \bar{p} + A_\mu}{\sigma}}^{\hat{z}} \frac{1}{2} dz dB(\sigma, p) \\ &\geq \int_{(\sigma, p) \in S_B} 2c dB(\sigma, p) \geq 2c, \end{aligned}$$

as claimed.

Together, $L'(\bar{u}) < 0$, $L(\bar{u}) < c$ for \bar{u} sufficiently large, and $L(\bar{u}) > c$ for \bar{u} sufficiently small suffice to show that there exists a unique solution to $L(\bar{u}) = c$.

The following lemmas are used in the proof of Proposition 4.

Lemma 6 *The marginal utility z of the marginal consumer at a firm with differentiation σ that is pricing optimally is increasing in the cutoff level of utility \bar{u} .*

Proof. Fixing σ , we need \bar{z} increasing with \bar{u} in the following expression, but this is just equivalent to Condition (H).

$$\bar{u} = A_\mu + \sigma \left(\bar{z} - \frac{1 - F(\bar{z})}{f(\bar{z})} \right) \quad (12)$$

■

Recall that we have defined $v(z) = \frac{1 - F(z)}{f(z)}$.

Lemma 7 *Let $u' = A_\mu - \sigma_H v(0)$ and $u'' = A_\mu - \sigma_L v(0)$. For $\bar{u} \in [u', u'']$, we have $z(\sigma_H; \bar{u}) > z(\sigma_L; \bar{u})$. (That is, for any utility level in this range, a firm choosing σ_H sells to fewer consumers than a firm choosing σ_L .)*

Proof. Note that $z(\sigma_H; u') = 0$, so Lemma 6 implies that $z(\sigma_H; \bar{u}) > 0$ for $\bar{u} > u'$. Similarly, $z(\sigma_L; u'') = 0$, so by Lemma 6, we have $z(\sigma_L; \bar{u}) < 0$ for $\bar{u} < u''$. ■

Lemma 8 *$\Delta(\bar{u})$ is negative for $\bar{u} < u'$ and positive for $\bar{u} > u''$.*

Proof. From the last lemma, both $z(\sigma_H; \bar{u})$ and $z(\sigma_L; \bar{u})$ are negative for $\bar{u} < u'$. Suppose $z(\sigma_H; \bar{u}) = \rho < 0$. Then the profit to choosing σ_H is $p^*(\sigma_H; \bar{u})(1 - F(\rho))$. Now imagine switching to σ_L but leaving the price unchanged. The firm's profit would then be $p^*(\sigma_H; \bar{u})(1 - F(\frac{\bar{u} + p^*(\sigma_H; \bar{u}) - A_\mu}{\sigma_L})) = p^*(\sigma_H; \bar{u})(1 - F(\frac{\sigma_H}{\sigma_L} \rho))$ – an improvement because $\frac{\sigma_H}{\sigma_L} \rho < \rho$. Switching to σ_L and pricing optimally would be better still. This proves the first part.

We also have both $z(\sigma_H; \bar{u})$ and $z(\sigma_L; \bar{u})$ positive for $\bar{u} > u''$. By similar logic, we can show that a firm using σ_L can always switch to σ_H and gain customers without changing its price, proving the second part. ■

Proof of Proposition 2

A firm that expects \bar{u} , chooses σ_L , and prices optimally earns

$$\pi_L(\bar{u}) = p_L(\bar{u}) \left(1 - F\left(\frac{\bar{u} - A_\mu + p_L(\bar{u})}{\sigma_L}\right) \right)$$

As consumers become more selective, this profit changes according to (applying the envelope theorem)

$$\frac{d\pi_L(\bar{u})}{d\bar{u}} = -\frac{p_L(\bar{u})}{\sigma_L} f\left(\frac{\bar{u} - A_\mu + p_L(\bar{u})}{\sigma_L}\right)$$

Using the equation (9) for the optimal price, we can write this as

$$\frac{d\pi_L(\bar{u})}{d\bar{u}} = -(1 - F(z_L(\bar{u})))$$

where $z_L(\bar{u}) = \frac{\bar{u} - A_\mu + p_L(\bar{u})}{\sigma_L}$. Similarly, for a firm choosing σ_H , we have $\frac{d\pi_H(\bar{u})}{d\bar{u}} = -(1 - F(z_H(\bar{u})))$. Define $\Delta(\bar{u}) = \pi_H(\bar{u}) - \pi_L(\bar{u})$ to be the difference between the profits to σ_H and σ_L (assuming optimal pricing in either case). Lemma 4 (in the appendix) demonstrates that there are utility levels u' and u'' such that $\Delta(\bar{u})$ is negative for $\bar{u} < u'$ and positive for $\bar{u} > u''$. Thus for consumer utility sufficiently low (high) σ_L (respectively σ_H) is optimal. To show that there is a unique $\tilde{u} \in (u', u'')$ at which the optimal product choice switches from σ_L to σ_H , note that

$$\frac{d\Delta(\bar{u})}{d\bar{u}} = F(z_H(\bar{u})) - F(z_L(\bar{u}))$$

Lemma 3 demonstrates that for the same u' and u'' , any $\bar{u} \in [u', u'']$ satisfies $z_H(\bar{u}) > z_L(\bar{u})$. Thus $\frac{d\Delta(\bar{u})}{d\bar{u}}$ is strictly positive on $[u', u'']$, which suffices to prove the claim.

Proof of Lemma 3

Using (10),

$$\frac{d\bar{u}(\boldsymbol{\lambda}^e, \mathbf{p}^e; c)}{d\lambda_H^e} = \frac{\sigma_H I(z_H) - \sigma_L I(z_L)}{\lambda_L(1 - F(z_L)) + \lambda_H(1 - F(z_H))}$$

where

$$z_S = \frac{\bar{u}(\boldsymbol{\lambda}^e, \mathbf{p}^e; c) + p_S^e - A_\mu}{\sigma_S}, \quad S = L, H.$$

So $\text{sgn}\left(\frac{d\bar{u}(\boldsymbol{\lambda}^e, \mathbf{p}^e; c)}{d\lambda_H^e}\right) = \text{sgn}(\sigma_H I(z_H) - \sigma_L I(z_L))$. Equation (10) tells us that a convex combination of $\sigma_L I(z_L)$ and $\sigma_H I(z_H)$ equals c . Thus, $\sigma_H I(z_H) > \sigma_L I(z_L)$ is equivalent to $\sigma_H I(z_H) > c > \sigma_L I(z_L)$. Likewise, $(\sigma_H I(z_H) < \sigma_L I(z_L)) \Leftrightarrow (\sigma_H I(z_H) < c < \sigma_L I(z_L))$, and $(\sigma_H I(z_H) = \sigma_L I(z_L)) \Leftrightarrow (\sigma_H I(z_H) = c = \sigma_L I(z_L))$.

Next, note that $\bar{u}_H(\mathbf{p}^e; c)$ and $\bar{u}_L(\mathbf{p}^e; c)$ are defined by

$$\begin{aligned}\sigma_H I(\tilde{z}_H) &= c, \text{ where } \tilde{z}_H = \frac{\bar{u}_H(\mathbf{p}^e; c) + p_H^e - A_\mu}{\sigma_H}, \text{ and} \\ \sigma_L I(\tilde{z}_L) &= c, \text{ where } \tilde{z}_L = \frac{\bar{u}_L(\mathbf{p}^e; c) + p_L^e - A_\mu}{\sigma_L}\end{aligned}$$

Recall that I is a strictly decreasing function, so $\sigma_H I(z_H) > c > \sigma_L I(z_L)$ is true if and only if $\tilde{z}_H > z_H$ and $\tilde{z}_L < z_L$. But this in turn is equivalent to $\bar{u}_H(\mathbf{p}^e; c) > \bar{u}(\boldsymbol{\lambda}^e, \mathbf{p}^e; c)$ and $\bar{u}_L(\mathbf{p}^e; c) < \bar{u}(\boldsymbol{\lambda}^e, \mathbf{p}^e; c)$. Thus, $\frac{d\bar{u}(\boldsymbol{\lambda}^e, \mathbf{p}^e; c)}{d\lambda_H^e} > 0$ is equivalent to $\bar{u}_H(\mathbf{p}^e; c) - \bar{u}_L(\mathbf{p}^e; c) > 0$. The equivalences for $\frac{d\bar{u}(\boldsymbol{\lambda}^e, \mathbf{p}^e; c)}{d\lambda_H^e} = 0$ and $\frac{d\bar{u}(\boldsymbol{\lambda}^e, \mathbf{p}^e; c)}{d\lambda_H^e} < 0$ follow in exactly the same way.

Proof of Lemma 4

For $v(\sigma; c)$ strictly convex, note that $\frac{\partial v(\sigma; c)}{\partial \sigma} = z(\sigma; c) + \frac{I(z(\sigma; c))}{1-F(z(\sigma; c))}$ depends on σ only through $z(\sigma; c)$. Differentiate to get

$$\begin{aligned}\frac{\partial^2 v(\sigma; c)}{\partial \sigma^2} &= \frac{\partial}{\partial z} \left(z + \frac{I(z)}{1-F(z)} \right) \cdot \frac{\partial z(\sigma; c)}{\partial \sigma} \\ &= \frac{f(z)I(z)}{(1-F(z))^2} \Big|_{z=z(\sigma; c)} \frac{\partial z(\sigma; c)}{\partial \sigma}\end{aligned}$$

Next, differentiate the definition of $z(\sigma; c)$ implicitly to get $\frac{\partial z(\sigma; c)}{\partial \sigma} = \frac{1}{\sigma} \frac{I(z(\sigma; c))}{1-F(z(\sigma; c))}$. So $\frac{\partial^2 v(\sigma; c)}{\partial \sigma^2} = \frac{1}{\sigma} \frac{f(z)I(z)^2}{(1-F(z))^3} \Big|_{z=z(\sigma; c)}$ which is strictly positive.

Proof of Proposition 3

By Lemma 2, no differentiation level σ except σ_L or σ_H can ever be a best response for a firm. Furthermore, (QC) implies that in any equilibrium, all firms that choose the same value of σ will set the same price, as determined by (9). Thus any potential equilibria can be classified depending on whether all firms choose σ_L , all firms choose σ_H , or some fraction λ_L of firms choose σ_L and the rest choose σ_H . We will take these three cases in turn.

EDE with all firms choosing σ_L

Note that the threshold \tilde{u} is determined entirely by the parameters A_μ , σ_L , and σ_H and the distribution F . An assessment with all firms choosing σ_L is an EDE if (by Lemma 5) the consumer utility cutoff satisfies $\bar{u} = U(c, \sigma_L)$, if $U(c, \sigma_L) \geq 0$ (consumers search rather than exiting immediately), and if the choice of σ_L is optimal for a firm, given $\bar{u} = U(c, \sigma_L)$ – this last condition requires $U(c, \sigma_L) \leq \tilde{u}$. (These conditions build in the requirement of correct beliefs.) If $\tilde{u} < 0$, the requirements that $U(c, \sigma_L) \geq 0$ and $U(c, \sigma_L) \leq \tilde{u}$ are incompatible; thus no all- σ_L EDE exists. This is one step toward showing part (i.b). If $\tilde{u} \geq 0$, then the equilibrium requirements can be satisfied if and only if $0 \leq U(c, \sigma_L) \leq \tilde{u}$. By Lemma 5, there exist c_3 and c_4 such that $U(c_3, \sigma_L) = \tilde{u}$, $U(c_4, \sigma_L) = 0$, and $c_3 \leq c_4$. Then because $U(c, \sigma_L)$ is strictly decreasing in c , $0 \leq U(c, \sigma_L) \leq \tilde{u}$ is satisfied if and only if $c \in [c_3, c_4]$. This demonstrates (ii.b).

EDE with all firms choosing σ_H

The logic for this case is similar to the previous one. Any equilibrium with all firms choosing σ_H must have consumer cutoff utility given by $U(c, \sigma_H)$. For consumers to search rather

than exit, we must have $U(c, \sigma_H) \geq 0$, and for firms to prefer σ_H over σ_L , we must have $U(c, \sigma_H) \geq \tilde{u}$. If $\tilde{u} < 0$, the first constraint (consumer participation) is the one that binds. In this case, Lemma 5 guarantees that there exists $c_1 > 0$ such that $U(c_1, \sigma_H) = 0$, and $U(c, \sigma_H) \geq 0$ if and only if $c \in (0, c_1]$. This suffices for (i.a). Alternatively, if $\tilde{u} \geq 0$, then the second constraint (firm optimality) binds. Again, by Lemma 5, there exists c_2 such that $U(c_2, \sigma_H) = \tilde{u}$, and $U(c, \sigma_H) \geq \tilde{u}$ if and only if $c \in (0, c_2]$. This suffices for (ii.a).

Asymmetric EDE

Consider an assessment in which a fraction λ_L of firms choose (σ_L, p_L) , the remaining λ_H firms choose (σ_H, p_H) , and the consumer cutoff utility is \bar{u} . The equilibrium condition for firm optimality requires that $\bar{u} = \tilde{u}$; otherwise either σ_L or σ_H would be strictly preferred. Consumer participation requires $\bar{u} \geq 0$, so if $\tilde{u} < 0$, no such equilibrium can exist. Remember that \tilde{u} is determined entirely by parameters. Given $\bar{u} = \tilde{u}$, price optimality for firms requires that

$$p_L^* = \sigma_L \frac{1 - F(z_L)}{f(z_L)} \text{ and } p_H^* = \sigma_H \frac{1 - F(z_H)}{f(z_H)}, \text{ where } z_L^* = \frac{\tilde{u} + p_L^* - A_\mu}{\sigma_L} \text{ and } z_H^* = \frac{\tilde{u} + p_H^* - A_\mu}{\sigma_H}$$

Thus, prices p_L^* and p_H^* and marginal taste shocks z_L^* and z_H^* in such an equilibrium are entirely pinned down by the parameters A_μ , σ_L , and σ_H . The equilibrium condition for consumers, given correct expectations about firms, is then

$$\lambda_L \sigma_L I(z_L^*) + \lambda_H \sigma_H I(z_H^*) = c \quad (13)$$

But $\sigma_L I(z_L^*)$ and $\sigma_H I(z_H^*)$ are pinned down by A_μ , σ_L , and σ_H . If $\sigma_L I(z_L^*)$ and $\sigma_H I(z_H^*)$ are both strictly larger or both strictly smaller than c , then (13) cannot be satisfied, and no asymmetric EDE exists. If $\sigma_L I(z_L^*)$ and $\sigma_H I(z_H^*)$ lie on opposite sides of c , then there is exactly one pair $(\lambda_L^*, \lambda_H^*)$, with $\lambda_L + \lambda_H = 1$, for which (13) can be satisfied. In this case, there is exactly one asymmetric EDE, described by $(\lambda_L^*, \lambda_H^*)$ and the strategies above. (More precisely, there is one such equilibrium, up to relabelings of which firms choose σ_L or σ_H .)

Now, recalling the definition of $U(\cdot)$, for $U(c_3, \sigma_L) = \tilde{u}$ we have $\sigma_L I(z) \Big|_{z = \frac{\tilde{u} + p - A_\mu}{\sigma_L}} = c_3$ and $p = \sigma_L \frac{1 - F(\bar{z})}{f(\bar{z})} \Big|_{z = \frac{\tilde{u} + p - A_\mu}{\sigma_L}}$. By inspection, the p and z that satisfy these conditions are just p_L^* and z_L^* . Thus we have $\sigma_L I(z_L^*) = c_3$, so $\sigma_L I(z_L^*) \geq c$ if and only if $c_3 \geq c$. Similarly, for $U(c_2, \sigma_H) = \tilde{u}$, we have $\sigma_H I(z_H^*) = c_2$, and so $\sigma_H I(z_H^*) \geq c$ if and only if $c_2 \geq c$. Thus, if $c > \max(c_2, c_3)$ or $c < \min(c_2, c_3)$, then (13) cannot be satisfied by any (λ_L, λ_H) . If $c \in (\min(c_2, c_3), \max(c_2, c_3))$, then $\sigma_L I(z_L^*)$ and $\sigma_H I(z_H^*)$ lie on opposite sides of c . In this case, the assessment is an EDE if and only if $\lambda_L c_3 + \lambda_H c_2 = c$, or equivalently, for $(\lambda_L^*, \lambda_H^*) = \left(\frac{|c_2 - c|}{|c_3 - c_2|}, \frac{|c_3 - c|}{|c_3 - c_2|} \right)$. (This equilibrium degenerates to full weight on σ_L or σ_H at the boundary cases $c = \min(c_2, c_3)$ or $c = \max(c_2, c_3)$.) This establishes part (ii.c).

This enumeration exhausts all the possible EDEs, so we are done with parts (i) and (ii). For part (iii), we claim that \tilde{u} and $U(c, \sigma)$ both move one-for-one with changes in A_μ . To clarify the relationship, rewrite these variables as $\tilde{u}(A_\mu)$ and $U(c, \sigma; A_\mu)$ to emphasize their dependence on A_μ . The claim is that $\tilde{u}(A'_\mu) - A'_\mu = \tilde{u}(A_\mu) - A_\mu$ and $U(c, \sigma; A'_\mu) - A'_\mu = U(c, \sigma; A_\mu) - A_\mu$ for any A_μ and A'_μ . To see this for $\tilde{u}(A_\mu)$, note that the profit equivalence that defines it,

$\pi_L(\tilde{u}(A_\mu)) = \pi_H(\tilde{u}(A_\mu))$ can be written (using (9))

$$p(1 - F(z)) \Big|_{z = \frac{\tilde{u}(A_\mu) - A_\mu + p}{\sigma_L}, p = \sigma_L \frac{1 - F(z)}{f(z)}} = p(1 - F(z)) \Big|_{z = \frac{\tilde{u}(A_\mu) - A_\mu + p}{\sigma_H}, p = \sigma_H \frac{1 - F(z)}{f(z)}}$$

which depends only on the quantity $(\tilde{u}(A_\mu) - A_\mu)$, not on $\tilde{u}(A_\mu)$ and A_μ separately. For $U(c, \sigma; A_\mu)$ we have $U(c, \sigma; A_\mu) - A_\mu = \sigma h\left(I^{-1}\left(\frac{c}{\sigma}\right)\right)$ which is obviously invariant to changes in A_μ . Together, these imply that $U(c, \sigma; A_\mu) - \tilde{u}(A_\mu) = U(c, \sigma; A'_\mu) - \tilde{u}(A'_\mu)$ for any A_μ and A'_μ , so changes in A_μ do not affect c_2 and c_3 . On the other hand, for c_1 defined by $U(c_1, \sigma_H; A_\mu) = 0$, since U is decreasing in c , it is straightforward that c_1 must rise if A_μ rises. (And similarly for c_4 .)

Proof of Proposition 4

It suffices to show that if $U(c_2, \sigma_H) = U(c_3, \sigma_L) = \tilde{u}$, then $c_2 = c_3$. Start with firms. A firm that anticipates consumer utility \tilde{u} and chooses σ_L will set a price, and therefore the marginal taste shock $z_L(\tilde{u})$ that it sells to, according to $p_L(\tilde{u}) = \sigma_L \frac{1 - F(z_L(\tilde{u}))}{f(z_L(\tilde{u}))}$, where $z_L(\tilde{u})$ satisfies $\tilde{u} = A_\mu + \sigma_L z_L(\tilde{u}) - p_L(\tilde{u})$. The last equation can be written $\tilde{u} = A_\mu + \sigma_L h(z_L(\tilde{u}))$. This firm earns profit $\sigma_L \frac{(1 - F(z_L(\tilde{u})))^2}{f(z_L(\tilde{u}))}$. Similarly, a firm that chooses σ_H instead will sell to a marginal consumer determined by $\tilde{u} = A_\mu + \sigma_H h(z_H(\tilde{u}))$ and earn profit $\sigma_H \frac{(1 - F(z_H(\tilde{u})))^2}{f(z_H(\tilde{u}))}$. Because firms are indifferent between σ_L and σ_H at \tilde{u} , we have

$$\begin{aligned} \tilde{u} = A_\mu + \sigma_L h(z_L(\tilde{u})) &= A_\mu + \sigma_H h(z_H(\tilde{u})) \quad \text{and} \\ \sigma_L \frac{(1 - F(z))^2}{f(z)} \Big|_{z=z_L(\tilde{u})} &= \sigma_H \frac{(1 - F(z))^2}{f(z)} \Big|_{z=z_H(\tilde{u})} \end{aligned}$$

Now consider consumers. If c_3 satisfies $\tilde{u} = U(c_3, \sigma_L) \equiv A_\mu + \sigma_L h\left(I^{-1}\left(\frac{c_3}{\sigma_L}\right)\right)$, then (given the previous line and strict monotonicity of h), we have $I^{-1}\left(\frac{c_3}{\sigma_L}\right) = z_L(\tilde{u})$, and thus $c_3 = \sigma_L I(z_L(\tilde{u}))$. By the same steps for $U(c_2, \sigma_H)$, we have $c_2 = \sigma_H I(z_H(\tilde{u}))$. Now use inverse hazard rate invariance. Suppose that $\frac{E(v(z) | z \geq \bar{z})}{v(\bar{z})} = \frac{I(\bar{z})}{(1 - F(\bar{z}))^2 / f(\bar{z})} = K$, or equivalently, $I(z) = K \frac{(1 - F(z))^2}{f(z)}$. Plugging this in, we have

$$c_3 = K \sigma_L \frac{(1 - F(z))^2}{f(z)} \Big|_{z=z_L(\tilde{u})} \quad \text{and} \quad c_2 = K \sigma_H \frac{(1 - F(z))^2}{f(z)} \Big|_{z=z_H(\tilde{u})}$$

But the righthand sides are equal, so we have $c_L = c_H$ as claimed.

Proof of Proposition 5

Given Proposition 4, all equilibria involve all firms choosing σ_L or all firms choosing σ_H . Fix a common level of dispersion $\sigma \in \{\sigma_L, \sigma_H\}$ for firms. In an equilibrium, the marginal consumer type \bar{z} is uniquely defined by $\sigma I(\bar{z}) = c$. A firm's optimal price is $p = \sigma v(\bar{z})$, and its

equilibrium profit is

$$\begin{aligned}
p(1 - F(\bar{z})) &= \sigma v(\bar{z})(1 - F(\bar{z})) \\
&= c \frac{v(\bar{z})(1 - F(\bar{z}))}{I(\bar{z})} \\
&= c \frac{v(\bar{z})}{E(v(z) | z \geq \bar{z})} \\
&= \frac{c}{K}
\end{aligned}$$

since the righthand side, by assumption, does not vary with \bar{z} .

Proof of Proposition 6

Define $z_L(c)$ to be the marginal consumer in an all σ_L assessment in which firms price optimally with respect to consumers, and consumers search optimally with respect to firms. Consumer utility in this profile is $U(c, \sigma_L) = A_\mu + \sigma_L h(z_L(c))$ where $z_L(c)$ is determined by $\sigma_L I(z_L(c)) = c$. (If $U(c, \sigma_L) \leq \tilde{u}$, so that σ_L is optimal for firms, this assessment is an EDE. Otherwise, it is not.) Define the function $J(z) = \frac{h(z)}{I(z)}$ so that we can write consumer utility in this profile as

$$U(c, \sigma_L) = A_\mu + cJ(z_L(c))$$

Define $z_H(c)$ similarly for an all σ_H assessment; then we also have $U(c, \sigma_H) = A_\mu + cJ(z_H(c))$. We will now demonstrate three claims about $z_L(c)$, $z_H(c)$, and $J(z)$.

Claim 1 $z_L(c) < z_H(c)$

This follows directly from $I' < 0$ and $\sigma_H > \sigma_L$.

Claim 2 $z_L(c)$ and $z_H(c)$ are both strictly decreasing in c .

Again, this follows from $I' < 0$.

Claim 3 If $f(z)$ is inverse hazard rate invariant, then $J(z)$ is strictly quasiconvex, with a minimum at $z = 0$.

Let $\frac{E(v(z) | z \geq \bar{z})}{v(\bar{z})} = \frac{I(\bar{z})f(\bar{z})}{(1-F(\bar{z}))^2} = K$, as usual. We have $J(z) = \frac{z}{I(z)} - \frac{1}{K} \frac{1}{1-F(z)}$, and therefore

$$\begin{aligned}
J'(z) &= \frac{I(z) + (1 - F(z))z}{I(z)^2} - \frac{1}{K} \frac{f(z)}{(1 - F(z))^2} \\
&= \frac{I(z) + (1 - F(z))z}{I(z)^2} - \frac{1}{I(z)} \\
&= \frac{(1 - F(z))z}{I(z)^2}
\end{aligned}$$

so $J(z)$ is strictly decreasing for $z < 0$ and strictly increasing for $z > 0$.

From Lemma 5, there exists some \bar{c} such that $U(\bar{c}, \sigma_L) = U(\bar{c}, \sigma_H) = \tilde{u}$. Thus we have $J(z_L(\bar{c})) = J(z_H(\bar{c}))$, and by Claim 3, $z_L(\bar{c}) < 0 < z_H(\bar{c})$. Then for any $c > \bar{c}$, we have $z_H(c) \in (z_L(c), z_H(\bar{c}))$ by Claims 1 and 2, and also $J(z_L(c)) > J(z_L(\bar{c})) = J(z_H(\bar{c}))$ by Claims 2 and 3. But then, $J(z_H(c)) < \max(J(z_L(c)), J(z_H(\bar{c}))) = J(z_L(c))$ by strict quasiconvexity. Proceeding similarly for $c < \bar{c}$, we have $z_L(c) \in (z_L(\bar{c}), z_H(c))$

by Claims 1 and 2, $J(z_H(c)) > J(z_L(\bar{c}))$ by Claims 2 and 3, and therefore, $J(z_L(c)) < \max(J(z_L(\bar{c}), J(z_H(c))) = J(z_H(c))$ by Claim 3. We conclude that $J(z_L(c)) \gtrless J(z_H(c))$ if and only if $c \gtrless \bar{c}$. But then because $U(c, \sigma_H) - U(c, \sigma_L)$ has the same signs as $J(z_H(c)) - J(z_L(c))$, we conclude that $U(c, \sigma_H) - U(c, \sigma_L)$ is positive if $c < \bar{c}$, negative if $c > \bar{c}$, and zero if $c = \bar{c}$. Because the EDE selects σ_H and σ_L precisely when $c < \bar{c}$ and $c > \bar{c}$ respectively, EDE utility is always equal to $\max(U(c, \sigma_H), U(c, \sigma_L))$. Because $U(c, \sigma)$ is decreasing in c , we have $U(c, \sigma_H) > U(c, \sigma_L) > \tilde{u}$ for $c < \bar{c}$ and $U(c, \sigma_H) < U(c, \sigma_L) < \tilde{u}$ for $c > \bar{c}$, as claimed.

Proof of Proposition 7

Each consumer will eventually purchase exactly one good, so payments to firms are just transfers that cancel out of the total surplus. Then the surplus is simply the gross expected utility of consumers, $A_\mu + \sigma z$. For a fixed distribution of levels of differentiation across firms, the socially optimal pattern of search is exactly the search strategy that consumers would choose were they trying to maximize σz rather than $\sigma z - p$. Suppose that levels of differentiation are distributed with density $\phi(\sigma)$ across firms. A consumer trying to maximize σz uses a cutoff value \bar{v} given by

$$\int_{\sigma_L}^{\sigma_H} \int_{\bar{v}/\sigma}^{\infty} \phi(\sigma) f(z) (\sigma z - \bar{v}) dz d\sigma = c \quad (14)$$

We can write this as $\int_{\sigma_L}^{\sigma_H} \phi(\sigma) \sigma \int_{\bar{v}/\sigma}^{\infty} f(z) (z - \bar{v}/\sigma) dz d\sigma = \int_{\sigma_L}^{\sigma_H} \phi(\sigma) \left(\sigma \int_{\bar{v}/\sigma}^{\infty} 1 - F(z) dz \right) d\sigma = c$.

Then observe that the inside term, $\sigma \int_{\bar{v}/\sigma}^{\infty} 1 - F(z) dz$, is increasing in σ : to see this, differentiate with respect to σ to get

$$\begin{aligned} \int_{\bar{v}/\sigma}^{\infty} 1 - F(z) dz + \frac{\bar{v}}{\sigma^2} (1 - F(\bar{v}/\sigma)) &= \\ \frac{1 - F(\bar{v}/\sigma)}{\sigma} \left(\int_{\bar{v}/\sigma}^{\infty} \frac{1 - F(z)}{1 - F(\bar{v}/\sigma)} dz + \frac{\bar{v}}{\sigma} \right) &= \\ \frac{1 - F(\bar{v}/\sigma)}{\sigma} (E(z - \frac{\bar{v}}{\sigma} | z \geq \frac{\bar{v}}{\sigma}) + \frac{\bar{v}}{\sigma}) &= \\ \frac{1 - F(\bar{v}/\sigma)}{\sigma} E(z | z \geq \frac{\bar{v}}{\sigma}) > \frac{1 - F(\bar{v}/\sigma)}{\sigma} E(z) = 0 \end{aligned}$$

Therefore, as $\phi(\sigma)$ transfers weight from lower to higher values of σ , the left-hand side of (14) increases. In order to preserve the equality, \bar{v} must increase as well. Thus, whenever $\phi(\sigma)$ does not place full weight on σ_H , \bar{v} , and hence total surplus, can be raised by placing more weight on σ_H , so the total surplus is maximized when all firms use $\sigma = \sigma_H$ and when consumers use a search rule that satisfies $\sigma_H \int_{z \geq \bar{v}/\sigma_H} 1 - F(z) dz = c$. But when $p = p^*(\sigma_H)$ is constant across all firms, we can define $\bar{u} = A_\mu - p^*(\sigma_H) + \bar{v}$ and write this search rule equivalently as $\sigma_H \int_{z \geq (\bar{u} + p^*(\sigma_H) - A_\mu)/\sigma_H} 1 - F(z) dz = c$, which is precisely the search rule that consumers use in an equilibrium in which all firms choose σ_H . Thus, if there is an all σ_H EDE for search cost c , then it is constrained efficient.

Proof of Proposition 8

A firm that expects consumer cutoff utility u and plans to offer products with dispersion level σ will optimally choose its price to satisfy (9) using distribution function $F_N(z) \equiv F(z)^N$.

That is, its price solves

$$p = \sigma \frac{1 - F_N \left(\frac{u+p-A_\mu}{\sigma} \right)}{f_N \left(\frac{u+p-A_\mu}{\sigma} \right)}$$

where $f_N = F'_N$. Write $p_L(u; N)$ for the optimal price set by a firm choosing $\sigma = \sigma_L$, let $\pi_L(u; N)$ be the profit earned by such a firm, and let $z_L(u; N) = \frac{u+p_L(u; N)-A_\mu}{\sigma_L}$ be the taste shock of the marginal consumer for such a firm. (So $\pi_L(u; N) = (1 - F_N(z_L(u; N)))p_L(u; N)$.) Define $p_H(u; N)$, $\pi_H(u; N)$, and $z_H(u; N)$ analogously for a firm choosing $\sigma = \sigma_H$. By definition, we have $\pi_L(\tilde{u}(N); N) \equiv \pi_H(\tilde{u}(N); N)$. In what follows, it is useful to treat N as a continuous parameter. As a matter of mathematics, the profit functions $\pi_L(u; N)$ and $\pi_H(u; N)$ are perfectly well defined at non-integer values of N , even though have no obvious economic interpretation at these values. Apply the implicit function theorem to the identity that defines $\tilde{u}(N)$ to get

$$\frac{d\tilde{u}(N)}{dN} = \frac{\frac{\partial \pi_L(\tilde{u}(N); N)}{\partial N} - \frac{\partial \pi_H(\tilde{u}(N); N)}{\partial N}}{\frac{\partial \pi_H(\tilde{u}(N); N)}{\partial u} - \frac{\partial \pi_L(\tilde{u}(N); N)}{\partial u}} \quad (15)$$

The proof of Proposition 2 demonstrates that $\pi_H(u; N) - \pi_L(u; N)$ is strictly increasing in u at $u = \tilde{u}(N)$, so the denominator is strictly positive. Thus, it suffices to show that the numerator is negative. That is, increasing the size of the product line slightly, holding consumer utility $\tilde{u}(N)$ fixed, improves profits relatively more at a σ_H firm than at a σ_L firm. To show this, write $\pi_L(u; N)$ in the form

$$\pi_L(u; N) = \left(1 - F \left(\frac{u + p_L(u; N) - A_\mu}{\sigma_L} \right)^N \right) p_L(u; N)$$

Recognizing that $p_L(u; N)$ is optimized with respect to u and N , apply the envelope theorem to get

$$\begin{aligned} \frac{\partial \pi_L(u; N)}{\partial N} &= -p_L(u; N) F \left(\frac{u + p_L(u; N) - A_\mu}{\sigma_L} \right)^N \ln F \left(\frac{u + p_L(u; N) - A_\mu}{\sigma_L} \right) \\ &= -p_L(u; N) (F(z_L(u; N)))^N \ln F(z_L(u; N)) \end{aligned}$$

For brevity, write $\mathcal{Z}_L = z_L(\tilde{u}(N); N)$, $\mathcal{Z}_H = z_H(\tilde{u}(N); N)$, and $\Pi = \pi_L(\tilde{u}(N); N) = \pi_H(\tilde{u}(N); N)$. Then we have

$$\begin{aligned} \frac{\partial \pi_L(\tilde{u}(N); N)}{\partial N} &= -p_L(\tilde{u}(N); N) F(\mathcal{Z}_L)^N \ln F(\mathcal{Z}_L) \\ &= \xi(F(\mathcal{Z}_L)) \Pi \end{aligned}$$

where $\xi(x) = -\frac{x^N}{1-x^N} \ln x$. In a similar manner, we can show that $\frac{\partial \pi_H(\tilde{u}(N); N)}{\partial N} = \xi(\mathcal{Z}_H) \Pi$. The proof of Proposition 2 also shows that $\mathcal{Z}_H > \mathcal{Z}_L$, and of course $F(\mathcal{Z}_L), F(\mathcal{Z}_H) \in (0, 1)$, so to demonstrate that the numerator of (15) is negative, it suffices to show that ξ is a strictly

increasing function on $(0, 1)$. This is relatively straightforward. We have

$$\xi'(x) = -\frac{x^{N-1}}{(1-x^N)^2} (1-x^N + \ln x^N)$$

To evaluate the term in parentheses, define $y = 1 - x^N$ (with $y \in (0, 1)$, since $x \in (0, 1)$) and take a series expansion of the log term to get

$$\begin{aligned} 1 - x^N + \ln x^N &= y + \ln(1 - y) \\ &= y - \left(y + \frac{y^2}{2} + \frac{y^3}{3} + \dots \right) < 0 \end{aligned}$$

So $\xi'(x) > 0$ for $x \in (0, 1)$, as claimed. This completes the proof that $\tilde{u}(N)$ is strictly decreasing in N . *A fortiori*, this result holds at the integer values of N that have economic meaning, so we have $\tilde{u}(1) > \tilde{u}(2) > \tilde{u}(3) > \dots$, and so forth.

References

- [1] ANDERSON, S. AND R. RENAULT (2006): “Advertising Content,” *American Economic Review*, 96, 93-113.
- [2] ANDERSON, S. AND R. RENAULT (1999): “Pricing, Product Diversity, and Search Costs: A Bertrand-Chamberlin-Diamond Model,” *RAND Journal of Economics*, 30, 719-35.
- [3] ARBATSKAYA, M. (2007): “Ordered Search,” *RAND Journal of Economics*, 38, 119-127.
- [4] ARMSTRONG, M., J. VICKERS, AND J. ZHOU (2008): “Prominence and Consumer Search,” University of Oxford Discussion Paper.
- [5] D’ASPROMONT, J., J. GABSZEWICZ, AND J.-F. THISSE (1979): “On Hotelling’s ‘Stability in Competition,’” *Econometrica*, 47, 1145-1150.
- [6] BAGNOLI, M. AND T. BERGSTROM (2005): “Log-Concave Probability and Its Applications,” *Economic Theory*, 26, 445-469.
- [7] BAGWELL, K. AND G. RAMEY (1996): “Coordination Economies, Sequential Search and Advertising,” Northwestern University, Center for Mathematical Studies in Economics and Management Science, Discussion Paper 1148.
- [8] BRYNJOLFSSON, E. AND M. SMITH (2000): “Frictionless Commerce? A Comparison of Internet and Conventional Retailers,” *Management Science*, 46.
- [9] BUTTERS, G. (1977): “Equilibrium Distributions of Sales and Advertising Prices,” *Review of Economic Studies*, 44, 465-491.
- [10] CACHON, G., C. TERWIESCH, AND Y. XU (2008): “On the Effects of Consumer Search and Firm Entry in a Multiproduct Competitive Market,” *Marketing Science*, 27(3), 461-473.

- [11] CAPLIN, A. AND B. NALEBUFF (1991): “Aggregation and Imperfect Competition: On the Existence of an Equilibrium,” *Econometrica*, 59, 25-59.
- [12] CHRISTOU, C. AND N. VETTAS (2008): “On Informative Advertising and Product Differentiation,” *International Journal of Industrial Organization*, 26, 92–112.
- [13] DIAMOND, P. (1971): “A Model of Price Adjustment,” *Journal of Economic Theory*, 3, 156-68.
- [14] HAAN, M. AND J. MORAGA GONZALEZ (2007): “Competing for Attention in a Model of Search,” mimeo.
- [15] HORTAÇSU, A. AND C. SYVERSON (2004): “Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry,” *Quarterly Journal of Economics*, 119, 403-456.
- [16] HOTELLING, H. (1929): “Stability in Competition,” *Economic Journal*, 39, 41-57.
- [17] KUKSOV, D. (2004): “Buyer Search Costs and Endogenous Product Design,” *Marketing Science*, 23(4), 490–499.
- [18] JOHNSON, J., and D. MYATT (2006): “On the Simple Economics of Advertising, Marketing, and Product Design,” *American Economic Review*, 96(3): 756–784.
- [19] PERLOFF, J. AND S. SALOP (1985): “Equilibrium with Product Differentiation,” *Review of Economic Studies*, 52, 107-20.
- [20] WOLINSKY, A. (1986): “True Monopolistic Competition As a Result of Imperfect Information,” *Quarterly Journal of Economics*, 96, 493-511.