Search, Information Provision, and Sales Concentration*

Andres Hervas-Drane†

August 14, 2010

Abstract

We present a model to assess the impact of increased product information provision on the firm's demand. Consumers face a search problem within an assortment of horizontally differentiated products supplied by a monopolist. They may search for a product match by drawing products from the assortment or by seeking product recommendations. We analyze the firm's incentives to supply product information, facilitate consumer-to-consumer communication, and implement personalization mechanisms such as recommender systems. Our model explains how these forms of information provision reduce consumer search costs and affect the concentration of sales. We account for recent developments in online retail, provide insights on their strategic implications for the firm, and contribute to the debate on their impact on sales concentration. The model is suited for experience good markets such as music, cinema, literature and video game entertainment.

Keywords: Product Discovery, Word of Mouth, Personalization, Recommender Systems, Long Tail

JEL Classification: C78, D42, D83, L15, M31

*I am deeply grateful to Ramon Casadesus-Masanell for excellent advice and support throughout this research. I thank Pablo Casas-Arce, Yeon-Koo Che, Peter Coles, Chris Dellarocas, Anita Elberse, Andrei Hagiu, Haim Mendelson, Natalija Novta, David Parkes, Michael Riordan and Francisco Ruiz-Aliseda. I also thank seminar participants at Harvard University, MIT Sloan, Real Colegio Complutense, Erasmus School of Economics, Pompeu Fabra, IESE, Harvard Business School, Stanford GSB, the Toulouse Internet Conference, Columbia University, Télécom ParisTech, and the CRES Strategy Conference at Olin Business School. The main body of this research was conducted at Harvard University. An earlier version of this paper was circulated under the title 'Word of Mouth and Taste Matching: A Theory of the Long Tail.' Financial support from the HBS Division of Research, the Real Colegio Complutense, the NET Institute, and the Caja Madrid Foundation are gratefully acknowledged.

†Universitat Pompeu Fabra (andres.hervas@upf.edu)

1 Introduction

The expansion of electronic commerce in recent years is transforming the retail landscape. Consumers are gaining access to a larger variety of products than ever before, and the trend has been most noticeable in product categories such as books, music, and films, where assortment sizes have increased dramatically. Online retailers have also become valuable venues for consumers to obtain information about products. Amazon, for instance, offers extensive product information such as book excerpts, music clips, and movie trailers, has become a platform for consumer-to-consumer communication by listing consumer reviews, ratings, and wish-lists, and has invested heavily in personalization features. We should expect these changes to affect the product discovery process of consumers.[1] Amazon's founder, Jeffrey P. Bezos, observed this in 1997. 'Today, online commerce saves customers money and precious time. Tomorrow, through personalization, online commerce will accelerate the very process of discovery.'[2] Examples of the growing importance of personalization abound. In 2009, Netflix awarded a million dollars in a public contest to improve the quality of its movie recommendations, which drive over 60 percent of its movie rentals, and chief product officer Neil Hunt stressed their strategic importance. 'Accurately predicting the movies Netflix members will love is a key component of our service.'[3] However, because we lack a formal understanding of how personalization affects consumers' product discovery process, it is unclear why firms have incentives to invest in it and how it affects demand.

Electronic commerce is also shifting the concentration of sales within product categories. Some observers have proposed that online distribution will increase the market share of products catering to niche audiences, increasing their participation in the sales mix with respect to traditional distribution channels. The main argument has been that traditional distribution limited the availability of products with a low market share due to logistical constraints. If some consumers can now access their preferred products online, which were previously unavailable, this should reduce the concentration of sales. But recent studies suggest that factors beyond product availability are

[1] For example, analysts estimate that by 2014 over half of retail sales in the US will be influenced by online research. See Forrester Research's 'US Online Retail Forecast 2009 To 2014,' March 5 2010.

[2] See Amazon.com's 1997 letter to shareholders.

[3] See 'The screens issue. If you liked this, you're sure to love that,' The New York Times, November 23 2008, and 'Netflix Awards $1 Million Prize and Starts a New Contest,' New York Times, September 21 2009.

contributing to drive down sales concentration online. Brynjolfsson et al. [8] and Elberse and Oberholzer-Gee [16] examine online and offline sales concentration for a clothing retailer and a large sample of video titles, controlling for differences in product availability, and continue to find lower sales concentration online. Both studies suggest that the online channel is triggering changes in consumption patterns, but the drivers of these changes are not well understood.

This paper presents a model that can rationalize the above facts and explain how they are interconnected. Our approach focuses on consumers' product discovery process and how it is affected by information provision. By presenting a novel search model, we explain how the provision of product information, consumer-to-consumer communication, and personalization reduce search costs in the market and affect the concentration of sales. It is profitable for firms to invest in these forms of information provision because it allows them to appropriate a larger share of consumer surplus. We also show that product information and consumer-to-consumer communication increase sales concentration, benefitting mass market products and mainstream consumers the most. Product market shares, on the one hand, enjoy increasing returns to appealing to a larger share of the consumer population. The benefits derived by consumers, on the other hand, exhibit increasing returns to the prevalence of their product preferences in the population. We show that introducing personalization in the market, a distinctive feature of the online channel, reduces sales concentration by eliminating these asymmetries. To the best of our knowledge, no previous theoretical work has explored the links between consumer-to-consumer communication, personalization, and sales concentration.

Because our model explains why a firm can profit from increased information provision that reduces consumer search costs, it can explain recent trends observed in the marketplace with the deployment of personalization mechanisms. Major online retailers have pioneered these trends by implementing recommender systems in their storefronts. These systems generate recommendations by exploiting preference similarity across consumers (e.g. 'customers who bought this item also bought...'), and this is achieved by analyzing data on consumer preferences originating from product purchases, consumer demographics, browsing activity, product reviews, product ratings, and product wish-lists. Other industry players in a position to track and exploit consumer activity such as brick-and-mortar retailers, social networking sites, and financial intermediaries are also investing to deploy them. Our model sheds light on the mechanisms that enable firms supplying superior personalization to generate value and sustain a competitive advantage.

We consider a market of horizontally differentiated products supplied by a monopolist at a common price. The monopolist may be an online retailer or content provider offering a large product assortment. Consumer preferences partition the product space into preferred and non-preferred products, and consumers only derive utility from the consumption of the former. But consumers cannot identify their preferred products within the assortment when they arrive to the market, all products are ex-ante identical and the value of each product can only be determined by sampling it. However, sampling products is costly as it requires time and attention, and thus consumers face a search problem to locate preferred products. A product *match* is achieved when a consumer searches and identifies a product that belongs to her preferred set, and consumers form an expectation of their search costs to locate a match when deciding to participate or not in the market.

To enrich the demand side of the market, we let consumers differ in their product preferences and sampling costs. Consumers search for a match by sampling products, and may either draw products directly from the assortment or seek product recommendations. We show that improved product information that allows consumers to sample products before purchase reduces search costs in the market, eliminating unsuccessful purchases and increasing firm profits. Consumers seeking product recommendations from other consumers learn from those that previously located a match by drawing products from the assortment, and we find that consumers choose to seek and follow recommendations because this increases their probability of locating a product match. Mainstream consumers, those whose product preferences are more prevalent in the population, benefit the most because recommendations are more likely to originate from others that share their preferences, thus enjoying a larger probability of locating a match. Niche consumers with less prevalent product preferences benefit less, and may not seek recommendations. When the firm introduces personalization, consumers obtain personalized product recommendations that account for their product preferences. This improves the probability of locating a match for all consumers, but has a larger impact on niche consumers and this reduces the concentration of sales.

The construction is well suited for experience goods such as books, music, films, or video games.

The satisfaction derived from these products is hard to anticipate. It can be argued that however informed a consumer may be on the discernible characteristics of a product, such as genre, characteristics or plot, personal judgment requires direct exposure. In addition, these product categories are well suited to horizontal differentiation settings, as preferences are largely idiosyncratic and consumers tend to not agree on their preferred products. For this reason, consumer-to-consumer communication arises in our model because of preference correlation in the population, and not because of information acquisition about superior products.

1.1 Literature

Little theoretical work has examined the mechanisms driving the product discovery process of consumers within large assortments. Product differentiation models, for example, cannot readily explain how changes in the information structure affect the composition of sales. Some instances have explored heterogeneous consumer preferences with location models, such as Bakos [7]. But in this case the equilibrium is symmetric for all consumer types and sellers, and no sales concentration is predicted by the model. Search models have mainly focused on price dispersion, by considering homogeneous goods offered by multiple sellers. These models are suited for settings where price dominates the search and provide little insight on sales concentration across heterogeneous product assortments. Anderson and Renault [4] have analyzed a case where sellers compete with differentiated products. Consumer valuations for different products are realized randomly during search, however, so sellers make no strategic choices affecting sales concentration and the equilibrium is again symmetric.

Our analysis is based on the monopoly case, where there is a unique seller in the market supplying an assortment of products. Our findings on how information provision affects consumer search costs and sales concentration are easier to convey in this setting. Some contributions in the literature have considered the incentives of a monopolist to invest in information provision. Johnson and Myatt [22] analyze the incentives to supply informative advertising by examining its impact through rotations on the demand curve, and find that the degree of information disclosed by a monopolist depends on how widespread is the appeal of the advertised product. The mechanisms underlying demand rotation that constitute the main focus of our analysis are not explored, however.

Anderson and Renault [5] examine the hold-up problem that arises when a monopolist supplies informative advertising without committing to prices, as consumers expect the firm to charge high prices once they have incurred search costs. The hold-up problem does not arise in our framework, as we consider a market where consumers observe prices before searching. Other contributions have considered scenarios where the monopolist profits from strategic information disclosure that increases search costs. Hagiu and Jullien [19] examine the incentives to divert search of a monopolist acting as a gateway to independent sellers. In that setting, the incentives to divert search mainly arise from the strategic interaction between the monopolist and the sellers.

The observation that online distribution could trigger changes in the concentration of sales was proposed by Anderson [3] and coined as the *long tail*, referring to the increase in the tail of the sales distribution. The subject has become a matter of academic debate, with a growing strand of the literature turning to the empirical evidence. Elberse and Oberholzer-Gee [16] report decreasing sales concentration within a sample of video titles over a five year period. Their dataset covers both online and offline retail channels, and they conclude that the changes observed are driven by demand side effects and online retailing. Brynjolfsson et al. [8] analyze the sales distribution of a clothing retailer offering the same product assortment across two separate channels with equal prices and terms of sale, both catalog and online. They find that sales concentration is lower online, even when considering consumers that purchase through both channels. They also present a search model with advertising where consumers incur search costs to learn about products that have not been advertised, and sales concentration depends on how the size of the advertised and non-advertised product pools compare. Fleder and Hosanagar [17] analyze the impact of recommender systems on sales concentration with simulations where consumers and products are located on a 2-axis space. They examine several scenarios, and find that the recommender tends to increase concentration across most of them. Feng and Zhang [34] analyze sales data from the video game market and the impact of online consumer reviews. They find their impact to be stronger for niche products where alternative sources of information are scarce, potentially contributing to lower sales concentration online.

Artistic markets exhibit highly concentrated sales distributions with a minority of bestselling titles. The phenomenon is widely acknowledged in music, cinema and books, and has sometimes been referred to as ‘hit culture’. A series of papers in the economics literature have analyzed these markets, pioneered by Rosen’s [29] famous superstars model as well as later contributions, such as MacDonald [26]. This literature has, for the most part, explained the phenomena by assuming a dispersion of talent among producers – greater talent commands higher profits and market shares than lesser talent. While this approach provides valuable insights on artistic markets, it is unclear that talent alone can explain the distribution of sales. Consumers generally acknowledge that differences in talent are important, yet they have a hard time describing what defines talent or evaluating it. Artistic quality may not be measurable independently of taste. Producers widely recognized as talented do not appeal to all consumers, while lesser talented artists generally have a niche audience of followers. Our analysis suggests that mainstream appeal and the added effects of search costs may well be an alternative route to stardom.

The paper is organized as follows. The next section introduces the building blocks of our search model and then solves the equilibrium for the simplest instance of search. We then proceed to enrich search strategies in steps to isolate their impact on the market. In Section 3 we introduce evaluations and allow consumers to learn the utility they derive from products before purchasing them. In Section 4 we introduce consumer-to-consumer communication in the model, and let consumers seek product recommendations from others. In Section 5 we consider personalization and let the firm supply personalized product recommendations to consumers. We discuss extensions to the model in Section 6, and conclude in Section 7.

2 Search framework

2.1 The model

Consider a market where a monopolist supplies an assortment of products. The assortment consists of a continuum of products of measure one. We partition the product space into N product pools, which can be understood as product varieties. For simplicity, we assume that product pools are of equal size, so the measure of each product pool within the product space is $1/N$. It is important to stress that the purpose of this partition is to define consumer preferences, and that there are no discernible product characteristics that allow products to be classified by pools. The monopolist quotes a common price p for all products in the assortment and incurs a transaction cost t per unit sold. The single price restriction implies that the monopolist cannot price discriminate consumers, which simplifies the analysis and allows us to isolate demand-side effects. We discuss richer pricing strategies in Section 6.

In this market there is a unit mass of consumers. Preferences over products are simplified to a binary classification: a consumer may derive positive utility from a product or not. In the first case the consumer derives utility u from consumption, and in the second case the consumer derives zero utility. Consumers exhibit unit demand, and may participate in the market to purchase and consume a preferred product or remain out. Consumers are heterogeneous in their product preferences, and we take the view that the most significant difference across consumers is their selection of preferred products. In particular, we assume that all consumers agree on some products, which exhibit universal appeal, but differ in their remaining subset of preferred products. We consider $T = N - 1$ consumer types, and let consumers of type t prefer products pertaining to product pools t and N. So products in pool N are mass market products, while products in the remaining pools appeal only to a subset of the population. We will refer to T as a measure of *taste fragmentation*, since the larger the value of T, the more differentiated the product space is for consumers. We assume throughout that $T \geq 3$, and therefore $N \geq 4$.

The analysis is of interest when consumer types differ in their prevalence in the population, and we denote by s^t the share of consumers of type t. Without loss of generality, we order types in increasing prevalence, where $s^1 < s^2 ... < s^T$. Thus consumer types become increasingly mainstream in t (or less niche), as their preferences are more widespread in the population. Similarly, product pools also become more mainstream in t, as they appeal to a larger share of the population.

When entering the market, consumers observe the level of prices p and taste fragmentation in the population, T. However, they arrive uninformed about products and cannot identify their preferred product pools within the assortment. All products are ex-ante identical, and as a result consumers face a search problem in order to locate a preferred product. A consumer can become informed about products by sampling them. A product *match* is achieved when a preferred product is identified. Sampling products is costly, and we let sampling costs be uniformly distributed in the consumer population, independently of product preferences, where the cost of consumer i is given by $c^i \sim U[0, \bar{c}]$. Thus sampling a product which does not yield a match incurs disutility c^i, and sampling and consuming a product match yields utility $u - c^i$. Consumers always incur a sampling cost before consumption. For experience goods, this can be interpreted as the time investment required to experience the good. We will assume $\bar{c}$ is sufficiently high to ensure the market remains uncovered. This simplifies our analysis by avoiding corner solutions in the pricing game, as a positive mass of consumers will not participate in the market in equilibrium.[4]

To summarize our model:

- There are $N = T + 1$ product pools in the assortment, and all products are priced at p.
- There are T consumer types, and consumers of type t derive utility u from products in pools t and N, and zero from the remaining.
- The share of consumers of type t in the population is given by s^t, where $s^t < s^{t+1}$.
- Sampling costs are uniformly distributed in the consumer population, $c^i \sim U[0, \bar{c}]$.

2.2 Search benchmark

We start our analysis by considering the simplest instance of search in our model. This is a two-stage game where the monopolist first chooses the price level in the market p. In the second stage, consumers may search for a match by sequentially drawing and purchasing products from the assortment. Consumers can only become informed about products by purchasing them first, as there is no product information available, and consumers incur price p and sampling cost c^i on each draw. The following graph depicts the sequential search process faced by consumers:

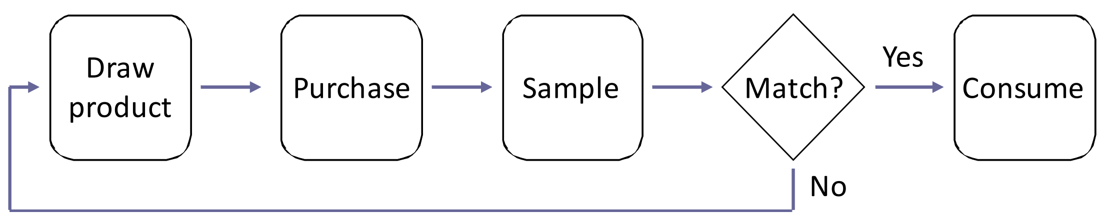

We solve search by assuming uniform sampling from the assortment. This is consistent with the fact that products are ex-ante identical for consumers. We define *search costs* as the sum of

[4] This requires $\bar{c} > \frac{1}{2}(u - t - r)$ throughout our analysis, where r is the cost of seeking a recommendation introduced in Section 4.

the costs incurred by a participating consumer to locate a product match, excluding the cost of purchasing the product. Search costs are endogenous in our framework, and we will show that they depend on the information available to consumers during their search. Consumers will form a rational expectation of their search costs when deciding to participate or not in the market, and will only participate if their search costs added to the price to be paid for their preferred product does not exceed utility u.

Before proceeding, it is useful to define a *sales distribution*. A sales distribution assigns a market share to each product pool in the assortment, and these are obtained by dividing the aggregate sales of products pertaining to each pool over the total sales across the assortment. This will be useful to evaluate the impact of different search strategies on the market, since the sales distribution allows us to isolate variations in the concentration of sales (or market share variations) from volume effects driven by shifts in consumer participation. In particular, we are interested in analyzing how changes in the information available to consumers affect the concentration of sales. To compare concentration across sales distributions we will apply the following property. Consider an ordering of product pools in increasing market share order, such that the product pool with rank 1 has the lowest market share and the rank N pool has the highest. A market share transfer from a low rank pool to a higher rank pool that preserves the ranks is said to *increase* concentration. Conversely, a rank-preserving transfer from high to low rank pools is said to *reduce* concentration. All concentration indexes in the literature satisfy this property, including for example the Gini index.[5]

Consumer search strategy. We proceed by backwards induction and consider the search problem faced by consumers in the second stage given a price level p. The only feasible search strategy is to sequentially purchase and sample products until a match is located. Denote by β the match probability for a consumer on each draw. A consumer of type t will only obtain a product match when drawing a product from pools t or N. Since products are drawn uniformly from the assortment, the probability of drawing from any given pool is $1/N$. Hence,

$$\beta = \frac{2}{N}, \tag{1}$$

[5] See Hall and Tideman [20] for an analysis of the desirable properties for a measure of concentration.

and each purchase is a Bernoulli trial with success probability β, which is common for all consumers. The expected utility of a new purchase for a consumer i with sampling cost c^i is

$$u_s^i = \beta u - c^i - p, \tag{2}$$

given that utility u is only derived with probability β but price p and sampling cost c^i are incurred on each purchase. The expected utility of a purchase does not depend on a consumer's type, but will vary across consumers depending on their sampling cost c^i. The utility of a successive draw, however, is constant throughout the search for any given consumer. Hence a consumer either searches until a match is obtained or does not participate in the market. We can identify the consumer of each type which is strictly indifferent between both alternatives by equating u_s^i to zero. Denote this indifferent consumer by c_s^i,

$$c_s^i = \beta u - p. \tag{3}$$

Only consumers with a sampling cost $c^i \leq c_s^i$ choose to search, and participation is homogeneous across types. Consumers with a higher sampling cost prefer not to participate in the market. The search process for any consumer finalizes once a match is located; searching for a second match cannot be optimal given that product prices are homogeneous and search is costly.

Firm pricing. We next turn to the first stage of the game and solve the firm's problem. The consumer participation constraint for all types (3) is a function of price level p. Note that for the firm to sustain positive prices and face demand, so that $c_s^i > 0$, we require $t < \beta u$. If the monopolist's transaction costs are high or taste is very fragmented (high T), then $t \geq \beta u$ and no feasible transaction is profitable, so the market breaks down. We need only consider the case where $t < \beta u$. Given that search is a Bernoulli process and each trial has success probability β, the expected number of purchases a consumer requires for a match is β^{-1}. So consumers of all types with $c^i \leq c_s^i$ participate in the market and each consumer executes β^{-1} purchases on average. Firm profits given the aggregate demand for all product pools are

$$\pi_s = \frac{c_s^i}{\bar{c}} \beta^{-1}(p-t) = \frac{(u\beta - p)(p-t)}{\bar{c}\beta}. \tag{4}$$

Solving for the firm's optimal price we obtain

$$p_s = \frac{u\beta + t}{2}. \tag{5}$$

Sales distribution. We next characterize the distribution of sales across products. A participating consumer may purchase several non-preferred products until a match is located, due to failed draws during her search, but will only purchase a single preferred product. Denote by D_p and D_{np} a consumer's expected demand for preferred and non-preferred product pools respectively. The probability of purchasing a non-preferred product on each draw is given by $1 - \beta$. The probability of purchasing j non-preferred products before purchasing a preferred product is given by $(1 - \beta)^j \beta$. If we consider all possible search sequences, and given that each consumers has two preferred product pools,

$$D_p = \frac{1}{2} \sum_{j=0}^{\infty} (1 - \beta)^j \beta = \frac{1}{2}. \tag{6}$$

And since there are $N - 2$ non-preferred product pools, the expected demand for each of these pools is

$$D_{np} = \frac{1}{N-2} \sum_{j=0}^{\infty} j(1 - \beta)^j \beta = \frac{1}{2}. \tag{7}$$

So $D_p = D_{np}$, and each participating consumer's demand for preferred and non-preferred products pools coincides. Hence the sales distribution is uniform, and the concentration of sales is minimum.

Social welfare. We next derive social welfare SW_s, defined as the sum of consumer surplus and firm surplus. Every participating consumer generates social surplus u net of the transaction and sampling costs involved in the search. Since every consumer purchases on average β^{-1} products to locate a match,

$$SW_s = \frac{c_s^i}{\overline{c}}(u - \beta^{-1}t) - \int_0^{c_s^i} \beta^{-1} c^i dc^i. \tag{8}$$

Proposition 1 *When consumers search and there is no information provision, search costs are* *high due to due to unsuccessful purchases. The firm discounts prices to account for the low success rate, sales concentration is minimum, and the market breaks down if transaction costs are high or taste is very fragmented.*

When there is no information provision about products in the market, consumers need to incur unsuccessful purchases to locate a match. Consumers anticipate this and will not participate in the market if it does not pay off, given product prices and the number of expected purchases required. As a result, the firm discounts prices by β, the match probability faced by participating consumers on each draw from the assortment and which determines their willingness to participate. If transactions costs t are high or if taste is very fragmented, high T which implies a low β, no profitable price for the monopolist faces positive demand and the market breaks down.

All product pools enjoy equal market shares when consumers search without product information, so the sales distribution across product pools is uniform. This is due to the unsuccessful purchases incurred by consumers, which ensure that every participating consumer exhibits uniform demand (in expectation) over all products. If products that appealed to no consumers were present in the assortment, they would enjoy an equal market share to the rest. But their presence would reduce the monopolist's profits, increasing the rate of unsuccessful purchases and lowering prices. Thus when no product information is available to consumers, the market shares of products are not informative of consumer preferences.

3 Evaluations

We next analyze the market when the monopolist provides product information that allows consumers to sample products prior to purchase. Note that consumers, when taking prices as given, strictly prefer to sample products before purchase because it avoids unsuccessful purchases. But their ability to do so depends on the monopolist, who may invest to supply product previews such as book excerpts, music clips, or movie trailers, or provide other means for consumers to experience products before purchase.

We introduce product evaluations in our two-stage game. In the first stage, the monopolist chooses the price level in the market, p. In the second stage, consumers may search for a match by sequentially drawing and sampling products from the assortment. Since consumers can sample products before purchase, they incur sampling cost c^i on each draw but will only execute a purchase at price p when they locate a match.[6] Consumers now face the following sequential search process:

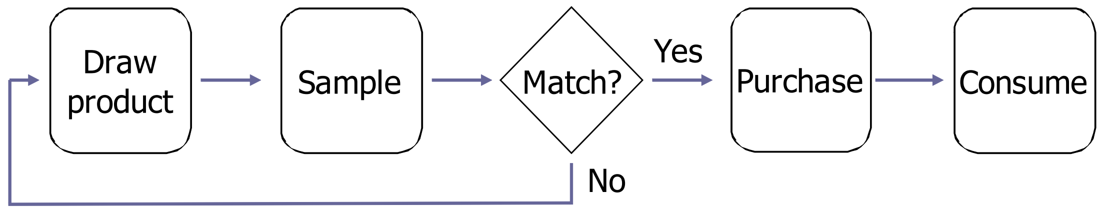

Consumer search strategy. Consider the consumer's problem in the second stage given a price level p. The probability of a match when drawing and sampling a product is given by β. The expected utility of a new product evaluation for an unmatched consumer is

$$u_e^i = \beta(u - p) - c^i, \tag{9}$$

given that consumers only purchase if a match is located but incur sampling cost c^i on every draw. The expected utility does not depend on a consumer's type, but will vary across consumers depending on their sampling cost. The utility of a successive draw, however, is constant throughout the search for any given consumer. Hence we can identify the consumer of each type which is strictly indifferent between evaluating products and not participating by equating u_e^i to zero. We denote the indifferent evaluator by c_e^i,

$$c_e^i = \beta(u - p). \tag{10}$$

Only consumers with a sampling cost $c^i \leq c_e^i$ choose to search, and participation is homogeneous across types. Consumers with a higher sampling cost prefer not to participate in the market. The search process for any consumer finalizes once a match is located; searching for a second match cannot be optimal.

Firm pricing. We next turn to the firm's problem given the consumer participation constraint for all types (10). Given that every participating consumer now purchases only once, firm profits are

[6] Our model of evaluations is equivalent to a market where consumers can realize costless returns of undesired products.

$$\pi_e = \frac{c_e^i}{\bar{c}}(p - t) = \frac{\beta(u - p)(p - t)}{\bar{c}}. \tag{11}$$

Solving for the firm's optimal price we obtain

$$p_e = \frac{u + t}{2}. \tag{12}$$

Sales concentration. Next we characterize the sales distribution with evaluations, denoted by σ. Let s_e^t be the share of consumers of type t among the mass of consumers that searches with evaluations. We proceed by characterizing separately the sales distribution generated by each consumer type σ^t, where $\sigma_n = \sum_t s_e^t \sigma_n^t$.

To characterize σ^t, note that consumers only purchase when they locate a product match, so the sales distribution generated by consumer of type t must equal their distribution of matches over products. Note that all consumers of type t are identically and independently distributed in the sampling outcome, as every product evaluation is independent of past evaluations and those of other consumers. Thus σ^t is independent of the market participation of consumers of type t, and we can derive σ^t by characterizing the distribution of matches over products for a single evaluation of a consumer of type t. To do so, it is useful to define indicator function λ based on consumer preferences. Let $\lambda_n^t = 1$ if $n = t$ or $n = N$, and $\lambda_n^t = 0$ otherwise. The probability that a consumer of type t matches a product in pool n is equal to $(1/N)\lambda_n^t$, and the probability of a match over all products is given by β. This implies

$$\sigma_n^t = \frac{(1/N)\lambda_n^t}{\beta} = \begin{cases} \frac{1}{2} & \text{if } n = t \text{ or } n = N \\ 0 & \text{otherwise} \end{cases} \tag{13}$$

We can now derive σ,

$$\sigma_n = \sum\nolimits_t s_e^t \sigma_n^t = \begin{cases} \frac{s_e^t}{2} & \text{if } n \in (1, N-1) \\ \frac{1}{2} & \text{if } n = N \end{cases} \tag{14}$$

And since participation is homogeneous across all consumer types, $s_e^t = s^t$ and the market share of product pools is increasing in n. Hence introducing evaluations strictly increases the concentration of sales in the market.

Social welfare. We next derive social welfare with evaluations, SW_e. Every participating consumer generates social surplus u net of transaction cost t and sampling costs, and each consumer samples on average β^{-1} products to locate a match,

$$SW_e = \frac{c_e^i}{\overline{c}}(u-t) - \int_0^{c_e^i} \beta^{-1} c^i \; dc^i. \tag{15}$$

It is easy to show that social welfare is higher with evaluations as long as sampling costs in the population are low, that is $SW_e > SW_s$ if and only if $\overline{c} \leq 4$. In particular, firm profits are always higher with evaluations, $\pi_e > \pi_s$, but the impact of evaluations on consumer surplus is only positive as long as $\overline{c} \leq 2$.

Proposition 2 *When the firm provides product information that allows consumers to evaluate products prior to purchase, this reduces search costs and increases firm profits. Product evaluations increase product prices, consumer participation, and the concentration of sales in the market.*

Evaluations allow consumers to purchase only products they match with, and this has two separate effects on the firm. On the one hand, there are more consumers ready to participate at every price level, since evaluations reduce search costs by ensuring there are no unsuccessful purchases. On the other hand, every participating consumer now realizes a unique purchase, once a match is located. These effects rotate the demand curve, expanding demand in the higher price range and contracting it in the lower range. The firm raises product prices and no longer discounts them by β, as there are no unsuccessful purchases. As a result, firm profits are strictly higher with evaluations.

Evaluations may be costly to implement for the firm if additional resources or infrastructure are required to supply product information. When transaction costs are high or taste is fragmented, $t < \beta u$, evaluations enable markets that would otherwise break down due to unsuccessful purchases. In these cases, a certain degree of information provision is necessary for the market to function and the firm has strong incentives to implement evaluations. The profitability of evaluations decreases quickly when taste becomes less fragmented, as $\beta \rightarrow 1$ and consumers incur few unsuccessful purchases in the absence of evaluations. Hence we should expect evaluations to be implemented when consumer taste is fragmented. The firm's incentives to implement evaluations also increase with match utility u and decrease with sampling costs $\bar{c}$, as higher sampling costs reduce market participation. Consumers are better off with evaluations only when sampling costs in the population are low.

Our analysis reveals that the firm benefits from lowering consumers' sampling costs. When introducing mechanisms to enable product evaluations the firm may also impact the sampling costs of consumers, and casual evidence suggests that firms invest in doing so. Many bookstores, for example, provide a comfortable environment and cafeteria services for their customers to browse books. Online retailers invest in the infrastructure required to directly stream book excerpts, music clips and movie trailers from their product pages. According to our model, this provides incentives for more consumers to search within the assortment, allowing the firm to sustain higher prices and increase profits.

Evaluations also increase the concentration of sales in the market. This result is notable because improved product exploration online has been suggested to reduce sales concentration, due to the possibility of more consumers venturing into niche products. But our analysis suggests otherwise, and the explanation is simple. Consumer participation increases with evaluations, but consumers no longer purchase products they do not match with and this increases the concentration of sales. As product pools differ in their appeal to the consumer population, when sales are realized by informed consumers there is a market share shift from pools that appeal to a small share of the population to those that appeal to a larger share, benefitting mass market products the most. But we should note that evaluations can also increase the sales volume of all products in the assortment, independently of the concentration shift. Thus niche products could increase their sales without increasing their market shares. This case arises when the participation increase due to evaluations is very large, when taste is fragmented and β is low, or transaction costs t are high.

4 Recommendations

We next introduce consumer-to-consumer communication in our model. This communication can be understood to take place online or offline. In the first case, the firm provides a platform for sampling consumers to actively publish their product recommendations and consumers seeking recommendations browse them. In the second case, consumers seeking recommendations observe which consumers have identified preferred products and request product references from them. While certain aspects of the search process will vary in each context, we attempt to formalize the core element of consumer-to-consumer communication in our model – the exchange of information among consumers about which products they like.

We introduce recommendations by adding a third stage to the game. In the first stage, the monopolist chooses the price level in the market, p. Consumers willing to participate then choose between two available search strategies. In the second stage, consumers may search for a match with evaluations by sequentially drawing and sampling products from the assortment. In the third stage, consumers may search for a match by browsing product recommendations provided from those consumers that searched before them in the second stage. Instead of drawing products from the product space, consumers searching with recommendations draw product references from the mass of consumers that searched with evaluations. A consumer providing a recommendation identifies the product she matched with.[7] The consumer seeking recommendations may then draw and sample the identified product at cost c^i. The sequential search process when seeking recommendations is as follows:

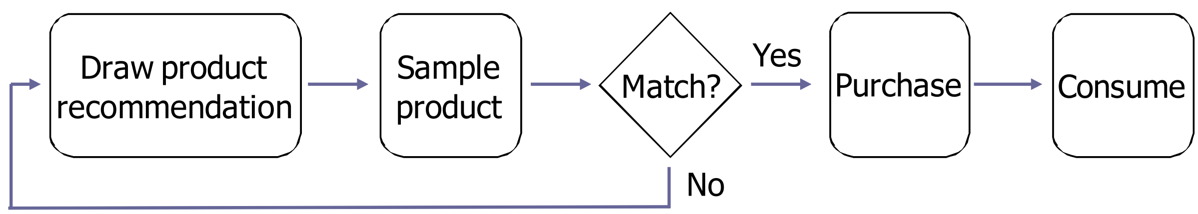

Recommendations are assumed to be drawn uniformly from the mass of consumers that searched with evaluations. This ensures that product recommendations in the market are representative of the evaluating population's preferences. We assume consumers seeking recommendations form a correct expectation of the share of evaluating consumers of their type, s_e^t. Past search experience, for example, could enable consumers to forecast it correctly. In equilibrium this will determine their match probability with recommendations. Each recommendation draw incurs a fixed cost r,

[7] Note that recommendations about which products not to sample, e.g. negative reviews or ratings, are not valuable for consumers engaged in search. As there is a continuum of products in the assortment, a consumer discarding a finite number of products cannot increase her probability of locating a match. The result carries over to discrete product spaces when they are sufficiently large.

since an additional step in the search is required to obtain information from others. To ensure that recommendations hold in the market, we need to assume $r < (u-t)/4$. Consumers providing recommendations freely identify their product match. We further discuss the implications of this assumption in Section 6.

Consumer search strategy. Consider the problem of an unmatched consumer in the third stage when the price level in the market is p. Product recommendations are drawn from the mass of consumers that searched with evaluations in the second stage. Note that that the sales distribution generated by evaluating consumers σ (14) carries over from our previous analysis, and describes the distribution of matches over product pools for the mass of evaluating consumers (although s_e^t will differ with recommendations). The expected probability of a match for a consumer of type t seeking recommendations, denoted by α^t, is given by

$$\alpha^t = \sigma_t + \sigma_N = \frac{1 + s_e^t}{2}. \tag{16}$$

The expression is a function of the share of evaluating consumers of type t. Thus the match probability when seeking recommendations will differ across types. As $\partial \alpha^t / \partial s_e^t > 0$, the larger the share of evaluating consumers of a consumer's own type, the larger her match probability when drawing a recommendation. We proceed by assuming that a positive mass of evaluating consumers of each type exists. Given that $s_e^t > 0$ and $N \geq 4$, it can be shown that $\alpha^t > \beta$ for all types.

The expected utility of seeking a new recommendation for consumer i of type t is

$$u_r^{t,i} = \alpha^t (u - p) - r - c^i, \tag{17}$$

as every recommendation draw incurs cost r in addition to sampling cost c^i. Note that the $u_r^{t,i}$ differs both across types due to α^t and within types depending on c^i. So while seeking recommendations yields a higher probability of a match on each draw, it is also more costly due to r. The utility of a successive draw, however, is constant throughout the search for any given consumer. Hence we can identify the consumer of type t which is strictly indifferent between seeking recommendations and not participating by equating $u_r^{t,i}$ to zero. We denote the indifferent recommendation seeker of type t by $c_r^{t,i}$, where

$$c_r^{t,i} = \alpha^t(u-p) - r. \tag{18}$$

Unmatched consumers of type t with a sampling cost $c^i \leq c_r^{t,i}$ choose to search with recommendations in the third stage, and those such that $c^{t,i} > c_r^{t,i}$ prefer to stay out of the market.

We next turn to the second stage of the game and analyze the decision to search with evaluations. As consumers anticipate that they may search with recommendations in the third stage, they decide which search strategy to pursue (if any) by comparing the expected utility of both. Given that the number of draws required for a match differs between both strategies, as $\alpha^t > \beta$ for all types, consumers need to evaluate the expected costs incurred to locate a match with both. Note that this comparison holds at any point of the search process for an unmatched consumer, as the expected utility of both search strategies is unaffected by past search history. This implies that no consumer that chooses to search with evaluations will abort the search in order to search with recommendations.

To identify the indifferent evaluator of type t, denoted by $c_e^{t,i}$, we equate the expected utility derived from both search strategies in order to locate a match, $u_r^{t,i} = u_e^i$. Note that u_e^i (9) carries over from our previous analysis and is type-independent. The expected number of draws required for a match with evaluations and recommendations are given by $1/\beta$ and $1/\alpha^t$ respectively. The indifferent evaluator of type t is then

$$\begin{aligned} u - p - \frac{r + c_e^{t,i}}{\alpha^t} &= u - p - \frac{c_e^{t,i}}{\beta} \\ c_e^{t,i} &= \frac{\beta r}{\alpha^t - \beta}. \end{aligned} \tag{19}$$

Consumers of type t with an evaluation cost $c^i \in [0, c_e^{t,i})$ prefer to search with evaluations in the second stage over seeking recommendations. For consistency, we require a positive mass of consumers of type t to seek recommendations in equilibrium, so $c_e^{t,i} < c_r^{t,i}$ must hold. As $c_r^{t,i}$ is decreasing in price level p for each type, we can identify the boundary price $\overline{p}^t$ by equating $c_e^{t,i} = c_r^{t,i}$,

$$\overline{p}^t = u - \frac{r}{\alpha^t - \beta}. \tag{20}$$

If no consumers of type t are willing to search with recommendations, consumers of this type will search only with evaluations and the indifferent evaluator of type t is given by $c_e^{t,i} = c_e^i$ as in (10), following our previous analysis. Note that participation is homogeneous across types that search only with evaluations.

We can now characterize consumer's search strategy. If $p < \overline{p}^t$, consumers of types t with sampling cost $c^i \in [0, c_e^{t,i})$ search with evaluations, and those with sampling cost $c^i \in [c_e^{t,i}, c_r^{t,i})$ seek recommendations. If $p \geq \overline{p}^t$, consumers of type t with sampling cost $c^i \in [0, c_e^i)$ search with evaluations. All remaining consumers stay out of the market.

We next characterize in more detail the composition of search strategies across types. Clearly, all types participate in the market, so there is always a positive mass of evaluators of each type. For those types that search with recommendations, note that $c_e^{t,i}$ is given by an implicit equation as α^t is a function of s_e^t, which in turn depends on the mass of evaluating consumers of all types, including the type considered. So the equilibrium participation of types that search with recommendations is defined by a system of implicit equations, one equation for each type. We next argue that the solution to this system satisfies that $c_e^{t,i}$ and s_e^t are decreasing and increasing in t, respectively, for types that search with recommendations. We show this by contradiction.

Assume recommendations hold for two types, t and $t+1$. First, consider the case $c_e^{t,i} = c_e^{t+1,i}$. This requires that $\alpha^t = \alpha^{t+1}$ by (19), which then implies that $s_e^t = s_e^{t+1}$ by (16). But on the other hand, since there is a larger share of consumers of type $t+1$ in the population, $s^t < s^{t+1}$ and $c_e^{t,i} = c_e^{t+1,i}$ both imply $s_e^t < s_e^{t+1}$, which is a contradiction. Next, consider the case $c_e^{t,i} < c_e^{t+1,i}$. This requires that $\alpha^t > \alpha^{t+1}$ by (19), which implies that $s_e^t > s_e^{t+1}$ by (16). But in this case $s^t < s^{t+1}$ and $c_e^{t,i} < c_e^{t+1,i}$ imply that $s_e^t < s_e^{t+1}$, which again is a contradiction. Hence the only feasible solution must satisfy $c_e^{t,i} > c_e^{t+1,i}$ and $s_e^t < s_e^{t+1}$ for types t and $t+1$.

We can now draw some conclusions for all types. Among the mass of consumers searching with evaluations and among the mass of consumers searching with recommendations, the shares of consumers of type t, denoted by s_e^t and s_r^t respectively, are increasing in t. To be sure, note that

$c_e^{t,i}$ is constant across types that search with evaluations only, and that if type t searchers with recommendations but type $t-1$ does not, $s_e^{t-1} < s_e^t$ must hold. So, since s_e^t is increasing in t, then α^t must also be increasing in t. The latter implies that $c_r^{t,i}$ and $\overline{p}^t$ are increasing in t, so s_r^t must also be increasing in t. Thus, in equilibrium, types with a large population share (higher t) have more incentives to search with recommendations than types with a low population share (lower t), and if recommendations hold for type t in equilibrium they must also hold for types $j > t$.

Firm pricing. We next turn to the first stage of the game and analyze the firm's pricing problem. Given a price level p in the market, we have established that only types t such that $p < \overline{p}^t$ search with recommendations. So the number of consumer types that search with recommendations decreases (in a step-wise fashion) with prices, and if prices are sufficiently high, $p \geq \overline{p}^T$, no types search with recommendations. Let t^r be the marginal type seeking recommendations given p, such that $\overline{p}^{t^r-1} \leq p < \overline{p}^{t^r}$ (recall that $\overline{p}^t$ is increasing in t). Firm profits can be written as

$$\pi_r = [\sum_{t=1}^{t_r-1} \frac{c_e^i}{\overline{c}} s^t + \sum_{t=t_r}^{T} \frac{c_r^t}{\overline{c}} s^t](p-t). \tag{21}$$

The firm's demand curve is composed of $T+1$ linear components, is continuous, (non-strictly) convex, and non-differentiable at $\overline{p}^t$ for $t \in (1,T)$. Each component of the demand curve describes a concave profit curve. Each profit curve lies above the rest in its own price range, and intersects with the curves of neighboring ranges at the price points $\overline{p}^t$ that separate components.

Define $\widehat{\alpha}^t$ as the following population-weighted match probability given the search strategies across of types when $t^r = t$,

$$\widehat{\alpha}^t = \frac{\sum_{t=t_r}^{T} s^t}{\sum_{t=1}^{t_r-1} s^t \beta + \sum_{t=t_r}^{T} s^t \alpha^t}, \tag{22}$$

where $\widehat{\alpha}^t > 0$. For each component of demand such that $t_r \in (1,T)$ we can derive the maximum of the corresponding profit curve from (21), denoted by $\widehat{p}^t$, where

$$\widehat{p}^t = \frac{u + t - r\widehat{\alpha}^t}{2}. \tag{23}$$

For the component in which $t_r = T+1$, consumers search only with evaluations and $\widehat{p}^{T+1} = p_e$

as in (12).

To identify the profit maximizing solution p_r, the firm need only evaluate profits at well defined maximums. Given the component-linearity and convexity of the demand curve, it follows that $\widehat{p}^t$ is increasing in t (so $\widehat{\alpha}^t$ must be decreasing in t). Well defined maximums are those such that $\overline{p}^{t-1} \leq \widehat{p}^t < \overline{p}^t$. In addition, whenever multiple maximums are well defined, it follows that they pertain to contiguous ranges. Our restriction on r ensures that the firm's solution falls in the range $p_r < p^T$ and recommendations hold in equilibrium for some consumer types.[8]

Sales concentration. We next analyze the impact of product recommendations on sales concentration. Denote the sales distribution in the market with recommendations by ρ, and let s^t_{er} be the share of consumers of type t among all participating consumers (with subindex er to denote that this includes both consumers searching with evaluations and recommendations). We argue that the introduction of recommendations increases the concentration of sales, and show this in two steps. Consider the sales distribution in the market with evaluations only, σ (14). To analyze how ρ differs from σ, we first account for the shift in consumer participation driven by recommendations while keeping fixed the per-type sales distribution (the *participation effect*). In doing so, we derive a participation-adjusted sales distribution $\overline{\rho}$, where $\overline{\rho}_n = \sum_t s^t_{er}\sigma^t_n$. In the second step, we account for the change in the sales distribution generated by consumers seeking recommendations (the *mass market effect*) to obtain ρ, where $\rho_n = \sum_t s^t_{er}\rho^t_n$ and ρ^t is the sales distribution generated by consumers of type t in the market.

To account for the participation shift, we can directly write $\overline{\rho}$ using σ^t (13),

$$\overline{\rho}_n = \sum\nolimits_t s^t_{er}\sigma^t_n = \begin{cases} \frac{s^t_{er}}{2} & \text{if } n < N \\ \frac{1}{2} & \text{if } n = N \end{cases} . \tag{24}$$

To see how $\overline{\rho}$ differs from σ, denote the marginal type that searches with recommendations by t^r, such that types $t < t^r$ search only with evaluations and types $t \geq t^r$ search with both evaluations and recommendations. We have established that participation is homogeneous for types $t < t^r$ and

[8] This requires the maximum for the component without recommendations to not be well defined, $\widehat{p}^{T+1} < \overline{p}^T$, which implies $r < \frac{1}{2}(u-t)(\alpha^T - \beta)$. Given that in equilibrium $\alpha^T > (1+1/T)/2$ and $\beta = 2/(T+1)$, it follows that $\alpha^T - \beta$ is increasing in T and $Lim_{T\to\infty}\ \alpha^T - \beta = 1/2$. So $r < (u-t)/4$ is sufficient to ensure recommendations hold in equilibrium.

given by c_e^i, while participation for types $t \geq t^r$ is given by $c_r^{t,i}$, where $c_r^{t,i} - c_e^{t,i} > 0$ and increasing in t. So s_{er}^t is constant for types $t < t^r$, and larger and increasing in t for types $t \geq t^r$. Inspection of $\overline{\rho}$ (24) and σ (14) reveals that this implies: (1) a market share transfer from product pools $n < t^r$ to pools $n \in (t^r, T)$, and (2) a market share transfer from pool n to pool $n+1$ within product pools $n \in (t^r, T)$. Since both transfers shift market share from low to high ranked product pools according to sales rank, the participation shift unambiguously increases concentration.

We next account for the shift in the per-type sales distribution generated by recommendation seekers. Note that ρ^t can be decomposed into sales driven by consumers of type t searching with evaluations, σ^t, and those searching with recommendations, which we denote by μ^t (which is only defined for consumer types that search with recommendations). To characterize the shift we next analyze how μ^t differs from σ^t.

To characterize μ^t, note that every recommendation draw is independent from past draws, so all consumers of type t seeking recommendations are identically and independently distributed. Thus μ^t is independent of the mass of consumers of type t seeking recommendations, and we need only characterize the distribution of matches for a single recommendation draw. The probability that a consumer of type t matches with product pool n when drawing a recommendation is given by $\lambda_n^t \sigma_n$, and the probability of a match over all products is given by α^t. This implies

$$\mu_n^t = \frac{\sigma_n \lambda_n^t}{\alpha^t} = \begin{cases} \frac{s_e^t}{1+s_e^t} & \text{if } n = t \\ \frac{1}{1+s_e^t} & \text{if } n = N \\ 0 & \text{otherwise} \end{cases} \tag{25}$$

Note that $\mu_t^t < 1/2$ and $\mu_N^t > 1/2$, so μ^t differs from σ^t in that $\mu_t^t < \sigma_t^t$ and $\mu_N^t > \sigma_N^t$. Since this implies a transfer from low to high ranked product pools according to sales rank, the sales distribution shift generated by recommendation seekers unambiguously increases concentration. Thus we conclude that the exchange of product recommendations strictly increases the concentration of sales in the market.

Social welfare. With respect to the equilibrium without recommendations derived in Section 3, whenever recommendations hold in equilibrium for some consumer types we have established that: (1) consumer participation is higher, and (2) prices are lower. This implies that recommendations strictly increase firm profits and consumer surplus, unambiguously increasing social welfare.

Proposition 3 *The exchange of product recommendations among consumers reduces search costs and increases firm profits. Product recommendations lower product prices, increase consumer participation benefitting consumers with widespread preferences the most, and increase the concentration of sales in the market.*

Product recommendations allow consumers to benefit from those that searched before them, reducing search costs in the market. Consumers seeking recommendations increase their probability of a match by gathering information about which products to sample. But since recommendations are costly, only consumers with high sampling costs seek recommendations, and consumers with low sampling costs prefer to search with evaluations. This implies that product recommendations rotate and expand the firm's demand in the low price range. Demand in the high price range is unaffected because no consumers seek recommendations when product prices are high, given that consumers seeking recommendations are those with high sampling costs and their willingness to participate is lower. As a result, the firm discounts prices to account for the value of recommendations in the market, increasing profits and ensuring recommendations are exchanged in equilibrium. With respect to the market with no recommendations, equilibrium prices are lower and consumer participation is higher.

The value of consumer-to-consumer communication increases with the fragmentation of taste in the market, as higher fragmentation lowers the match probability of consumers when sampling products from the assortment. This renders product recommendations more attractive, particularly due to the high probability that they lead to mass market products. The share of participating consumers that seek recommendations increases with consumption utility u and the fragmentation of taste T, and decreases with recommendation cost r. The introduction of product recommendations in the market increases consumer surplus. Since the firm lowers prices, all consumers are better off.

Similarly to lowering sampling costs for consumers, facilitating the exchange of product recommendations by lowering the cost of obtaining them has the potential to expand markets. This provides incentives for the firm to play an active role in the process, an opportunity fueled by the online environment. Online retailers such as Amazon have implemented mechanisms to facilitate consumer-to-consumer communication, becoming valuable resources for consumers in the process. Chevalier and Mayzlin [11] analyze the impact of online book reviews at two major online retailers. They find that reviews increase the relative sales at the retailer they are posted on. The findings are consistent with our model, and suggest that part of the market growth spurred by electronic commerce may be attributable to facilitating consumer-to-consumer communication alone.

Consumer-to-consumer communication also has important effects on the concentration of sales. This is driven by the fact that recommendations end up being exchanged in the market between consumers of different types, who differ in their product preferences. This cross-type exchange has an asymmetric impact across consumer types and across product pools. We decompose the impact in two effects, a *mass market effect* and a *participation effect*. The mass market effect follows from the fact that all consumers agree on mass market products. Consumers seeking recommendations are more likely to match with mass market products than those searching with evaluations, as successful cross-type recommendation can only yield a match with those products. This effect increases the market share of mass market products. The participation effect is driven by the fact that some product preferences are more widespread in the consumer population. In equilibrium, more recommendations originate from consumers with widespread preferences, as a larger mass of these consumers choose to search with evaluations. Thus the benefit consumers derive from recommendations increases with the prevalence of their taste in the population. As a result, consumers with widespread preferences exhibit higher participation, and a higher share of them search with recommendations. In addition, seeking recommendations may not pay off for consumers with uncommon preferences if their share in the population is sufficiently low, and those consumers may search only with evaluations. This effect increases the market shares of product pools with widespread appeal and decreases that of pools with low appeal.

Both the mass market and the participation effect increase the concentration of sales, and the shift in concentration grows with the share of consumers searching with recommendations. The mechanisms driving these effects have also been identified in the empirical literature. Leskovec et al. [25] analyze a large dataset originating from an online person-to-person recommendation network, and find that recommendations for products which are recommended more often also exhibit a higher success rate.[9] Hence products enjoy increasing returns to appealing to a larger share of the consumer population, reinforcing their market shares. This implies that market shares overestimate the appeal of best-selling products and underestimate that of lesser performing products. The result is reminiscent of the double jeopardy effect discussed by Ehrenberg et al. [13], where small brands perform comparatively worse than large brands. Our model suggests that the exchange of product information among consumers could be an explanatory factor.

The findings are also consistent with those reported for product popularity feedback. Salganik et al. [30] study the concentration of demand over a set of rare songs offered to test subjects on the Internet, with some treatments reporting to subjects the popularity of songs and others not. They find that popularity feedback increases both concentration and the unpredictability of popularity in the outcome. Tucker and Zhang [33] analyze a dataset containing the click-through rates of a webpage indexing marriage agencies, both when popularity is reported to users and when it is not. They find that both concentration and consumer participation increase with popularity feedback. This suggests that the exchange of product recommendations, and perhaps word-of-mouth processes more generally, have similar effects to popularity feedback.[10]

5 Personalized recommendations

We next analyze the impact of personalization in the market. The firm introduces a mechanism that generates personalized product recommendations for consumers. Our model has shown that the exchange of recommendations between consumers with different product preferences is inefficient because they don't always result in a product match. Thus the firm could improve the efficiency of product recommendations if she could generate personalized recommendations that always yield a match. Recommender systems implemented by online retailers strive to achieve this, using collaborative filtering techniques to identify preference similarity across consumers and recommending

[9] The study also finds that the number of distinct senders and receivers involved in the exchange of recommendations for a given product reduces the success rate. The authors argue that this is related to the structure of the recommendation network, a dimension of the problem which is not captured in our model.

[10] Note that the source of product recommendations is endogenous in our model, that is, the composition of consumers supplying recommendations. The supply of recommendations in equilibrium does not match the overall market shares of products, but instead matches the market shares generated by sales to consumers searching with evaluations. This differs from popularity indexes based on overall market shares, but turns out to have similar properties.

products that are preferred by consumers with similar preferences.[11] We introduce personalized recommendations in the third stage by assuming that consumers seeking recommendations always identify a preferred product, drawn uniformly from their set of preferred products. The setup is equivalent to a recommender system that perfectly matches consumers in the recommendations exchange based on their type.

Consumer search strategy. Recommendations always yield a match when they are personalized. Therefore $\alpha^t = 1$ for all t, and the match probability with personalization no longer depends on the composition of types among evaluating consumers. The impact of introducing personalization in the market follows from our analysis in the previous section taking into account that $\alpha^t = 1$. This homogenizes across types the utility of recommendations $u_r^{t,i}$(17), the indifferent recommendation seeker $c_r^{t,i}$ (18), the indifferent evaluator $c_e^{t,i}$ (19), and the boundary recommendation price $\overline{p}^t$. To account for the fact that they no longer depend on t, we denote them by u_r^i, c_r^i, c_e^i, and $\overline{p}$ respectively.

If prices are above the boundary recommendation price, $p \geq \overline{p}$, all types search only with evaluations and the market configuration is equivalent to that of Section 3. If $p < \overline{p}$, all types search with personalized recommendations. In this case $0 < c_e^i < c_r^i$ holds for all types, and there is a positive mass of consumers of each type willing to search with evaluations. This also implies that $s_e^t = s_r^t = s_{er}^t = s^t$, and participation is homogeneous across types.

Firm pricing. The firm's profit function π_r (21) carries over by taking into account that there is now a unique non-differentiability at $\overline{p}$. The demand curve has two linear components; either $p \geq \overline{p}$ and $t^r = T + 1$, or $p < \overline{p}$ and $t^r = 1$. The maximum of the profit curve in the range $p < \overline{p}$ is given by

$$p_p = \frac{u + t - r}{2}, \tag{26}$$

since $\widehat{\alpha}^1 = 1$ given that $\alpha^t = 1$ for all types (we need only consider the case $t^r = 1$ in the range $p < \overline{p}$). The firm's profit maximizing price is p_p, given that our restriction on r ensures that $p_p < \overline{p}$.[12]

[11] A taxonomy of recommender systems and an overview of the related computer science literature are presented by Adomavicius and Tuzhilin [2]. For a brief discussion on the economics of recommender systems, see Resnick and Varian [28].

[12] The maximum of the profit curve in the range $p \geq \overline{p}$ is given by p_e in (12). For the solution to be in the range

Sales concentration. We next argue that personalization of product recommendations reduces sales concentration. Consider the participation shift, given by $\overline{\rho}$ in (24). Since $s_{er}^t = s_e^t$ with personalization, $\overline{\rho} = \sigma$ and the participation shift does not alter concentration with respect to evaluations. Next, consider the sales distribution shift generated by recommendation seekers. With personalization, recommendations draw uniformly from the set of preferred products of each recommendation seeker. This implies that $\mu_n^t = 1/2$ for $n = \{t, N\}$, and zero otherwise. Thus $\mu^t = \sigma^t$ and recommendation seekers do not alter concentration with respect to evaluations. We conclude that $\rho = \sigma$ and sales concentration with personalized recommendations is equivalent to that derived in Section 3 with evaluations only.[13]

Social welfare. With respect to non-personalized recommendations in the previous section, consumer participation increases in the price range $p < \overline{p}$, unambiguously increasing firm profits in equilibrium. But the impact of personalization on consumer surplus is extremely complex to characterize, unfortunately. Personalization reduces search costs for consumers, but in addition may increase or decrease prices, rendering the net effect on consumer surplus ambiguous. To illustrate this, consider consumer surplus in the market when taking α and β as exogenous,

$$CW_p = \frac{c_r^i}{\overline{c}} u - \int_0^{c_e^i} \beta^{-1} c^i \, dc^i - \int_{c_e^i}^{c_r^i} \alpha^{-1}(c^i + r) \, dc^i. \tag{27}$$

In this scenario, prices are increasing in α, and it can be shown that $\partial CW_p / \partial \alpha < 0$ if sampling costs $\overline{c}$ are sufficiently high.

The impact on consumer surplus in our model is more complex, as α and β differ across types in the equilibrium with no personalization, and the sign and intensity of the price change depend on $\widehat{\alpha}^t$ (22) in prices without personalization p_r. Thus $p_p < p_r$ and $p_p > p_r$ are possible. Due to the complexity of $\widehat{\alpha}^t$ we are unable to pin down the exact behavior of prices in order to draw clear-cut conclusions, but the above suggests that consumer surplus will increase whenever $\widehat{\alpha}^t$ or sampling

$p < \overline{p}$ we require that $p_e < \overline{p}$, which implies $r < \frac{1}{2}(u-t)(1-\beta)$. This always holds given our assumption $r < (u-t)/4$. In addition, this equilibrium marks the highest consumer participation predicted in the model. For the market to remain uncovered in equilibrium, we require $c_r^i < \overline{c}$, which given p_p implies $\overline{c} > \frac{1}{2}(u-t-r)$. This lower boundary on $\overline{c}$ ensures the market is uncovered in all equilibria derived in our analysis.

[13] To see why this is equivalent to a recommender system that matches consumers based on their type, note that in that case recommendations are drawn exclusively from evaluating consumers of the same type. Then μ^t is given by $\mu_n^t = \sigma_n^t \lambda_n^t / \alpha^t = \sigma_n^t$, and the same result follows. For consistency, this requires a positive mass of consumers of each type to search with evaluations, a condition we have shown is satisfied in equilibrium.

costs $\bar{c}$ are low.

Proposition 4 *When the firm supplies personalized product recommendations, this reduces search costs and increases firm profits. Personalization may increase or decrease product prices, always increases consumer participation benefitting consumers with uncommon preferences the most, and reduces the concentration of sales in the market.*

Personalized product recommendations increase the value of recommendations for consumers, ensuring they always yield a match and thereby reducing search costs in the market. This rotates and expands the firm's demand in the low price range, the range in which consumers seek recommendations and benefit from personalization. More consumers are now ready to participate by seeking recommendations, and are ready to do so with higher prices. This ensures that firm profits always increase, and the firm adjusts prices to account for the higher value of recommendations in the market. This may increase or decrease prices, and the sign and intensity of the change depends on the precise market configuration in equilibrium when there is no personalization.

With personalization, the value consumers derive from recommendations no longer depends on how prevalent their preferences are in the population. So consumers with widespread preferences no longer enjoy an advantage over their peers. Thus consumers with uncommon preferences benefit the most from personalization, and since there is no longer an asymmetric benefit from recommendations across the consumer population, participation becomes homogeneous across types. The impact on consumer surplus, however, is ambiguous. Consumers searching with recommendations benefit from lower search costs, but a price increase could offset this benefit. Inspection of equilibrium prices with and without personalization suggests that consumer surplus increases when taste fragmentation T and sampling costs $\bar{c}$ are low, and the cost of recommendations r is high.[14]

Because personalization eliminates the exchange of product recommendations among consumers of different types, it reduces sales concentration in the market. To see this, consider the effects driving concentration in the absence of personalization. On the one hand, there is no longer a mass market effect. As there are no cross-type recommendations, consumers seeking recommendations no longer have a higher probability of matching with mass market products than matching with

[14] Note that consumer surplus is strictly higher than with evaluations only, and consumers always prefer personalized recommendations to no recommendations.

the remaining of their preferred products. This implies that personalization shifts market share from mass market products to all other product pools. On the other hand, there is no longer a participation effect. Again, since there are no cross-type recommendations, consumers with widespread preferences no longer derive higher value from recommendations than others and do not participate comparatively more in the market. This shifts market share from products that appeal to a large share of the population to those that appeal to a lower share. As a result, personalization always reduces sales concentration and renders it equivalent to that derived in Section 3 with evaluations only.

Our results suggest that personalization features can contribute to explain the lower sales concentration that has been reported in the online channel. Recent empirical research confirms that personalization plays an important role in the process. Oestreicher-Singer and Sundararajan [27] examine the sales concentration of book categories on Amazon accounting for the co-purchase patterns reported across books. The algorithms that Amazon uses to generate personalized recommendations are based on these patterns. They find that sales concentration is lower among categories with denser co-purchase networks, where personalization should be expected to be more accurate. Using similar measures of concentration and co-purchase patterns, Ehrmann and Schmale [14] report the same findings on Amazon's Germany site. Thus personalization may be one of the differential factors contributing to variations in sales concentration between the online and offline retail channels.

6 Discussion

In this section we review some of the assumptions made in our analysis and discuss the robustness of our findings.

Consumer preferences. The preference structure we have considered assumes that all consumers agree on some products but disagree on the remaining. This structure attempts to simplify the analysis while capturing interactions among consumers with different product preferences, and is broadly consistent with the empirical evidence available. Goel et al. [18] examine movie data from Netflix and music data from Yahoo Music and find that a large majority of users consume both mass market and niche products. Tan and Netessine [32] and Elberse [15] report similar findings.

Our setup could be simplified by assuming that consumer types fully disagree on their preferred products, ruling out the mass market product pool. Our main findings on search costs and sales concentration would continue to hold. But consumer-to-consumer communication would no longer exhibit a mass market effect, only the participation effect would be present.

Our setup could be enriched by considering partial overlapping in the product pools preferred by different consumer types, allowing some product pools to be preferred by several (but not all) consumer types. This would increase the complexity of consumer-to-consumer communication by increasing the interdependencies between types, as product pools that are preferred by several types would also benefit from cross-type recommendations. However, we do not expect these preference structures to deter from our main findings.

Our setup could also be enriched by considering different levels of utility across product pools. For example, consumers may derive higher utility from niche products catering to their type than from mass market products.[15] This would introduce more complex stopping rules in consumers' search strategies. In our setup consumers always stop searching once they locate a product match, so the stopping rule is simple and common for all consumers. If different products yield different levels of utility, a consumer may be willing to continue searching after locating a product match if the expected utility gain to be obtained from a better product match offsets the expected search costs required to obtain it. Unfortunately, more complex stopping rules significantly increase the complexity of the analysis. The precise mechanisms are beyond the scope of our analysis, but the general implications are clear – increasing the relative utility consumers obtain from a product pool will tend to increase its market share in equilibrium.

Assortment composition. We have assumed the assortment supplied by the monopolist to be exogenous. For example, we have assumed for simplicity that all product pools are of equal size. But one may expect product pools to vary in size, particularly if the firms or producers supplying the products can control product design to cater to different consumer types. Increasing the relative size of a product pool within the assortment would benefit consumers who derive utility from it,

[15] We should note that the empirical evidence is inconclusive on this point. Goel et al. [18] analyze consumer product ratings for movies and music. For movies, niche products are rated below mass market products, but the trend is reversed for music.

increasing their match probability when searching with evaluations β and reducing those of other consumers. Rendering evaluations more attractive for these consumers would also increase their match probability with recommendations α in equilibrium. As a result, increasing the comparative size of a product pool in the assortment will increase its market share by reducing the search costs for consumers who prefer those products.

If the monopolist were to choose the composition of the assortment to maximize profits, she would consider the marginal profitability of each stocked product. The monopolist would supply products that appeal to a large share of the population and discard those that appeal to few, minimizing search costs over the whole population. In our setup, the monopolist would maximize profits by supplying only mass market products. The fact that smaller assortments may be more attractive for consumers has been explored in the literature. Iyengar and Lepper [21] and Chernev [10] report experiments suggesting that more choice is not always preferred by consumers. Kuksov and Villas-Boas [23] formalize the findings with a search model where consumers anticipate higher search costs when facing larger assortments. But when assortments are too small, consumers anticipate they will not locate a good product match. This variety-seeking effect is absent from our setup. Our goal has been to abstract from supply-side decisions and focus on the empirically relevant case, given the large product assortments offered by major online retailers.

Pricing. We have assumed the monopolist quotes a single price across the whole assortment, so our search model has priorized product relevance over price. While the assumption may seem restrictive, price dispersion across titles is arguably low in the markets considered. Consumers perceive the prices of music, movies, books or videogames to be largely homogeneous, and it is not uncommon for best-sellers and obscure titles to share similar price tags. Moreover, major online retailers such as Amazon and Apple apply single price schemes across their digital content catalogs.

Relaxing this restriction implies that the monopolist can price discriminate consumers in the market. For instance, the monopolist could charge different prices to consumers searching with evaluations and consumers searching with recommendations. In this case, it can be shown that the monopolist would charge consumers searching with evaluations price p_s derived in Section 3. And since the monopolist no longer needs to pool consumers with different search strategies when setting prices, she would reduce prices for recommendation seekers below those derived in Sections 4 and

5. This type of price discrimination would be profitable for the monopolist and would increase the share of consumers searching with recommendations, further increasing sales concentration when there is no personalization. Kuksov and Xie [24] analyze a firm's incentives to alter price or provide unexpected frills to early customers in order to obtain better product ratings and profit from later customers. This effect is not present in our model because the surplus and the mass of consumers searching with evaluations do not impact product recommendations. Otherwise, the monopolist would lower the price for evaluators to internalize their impact on recommendation seekers.

A more complex type of discrimination would entail pricing product pools independently. Such a pricing scheme increases the complexity of the problem significantly due to the informational content of prices, and additional assumptions on how consumers observe prices would be required. For instance, if consumers perfectly observe the prices of all products, then prices could act as a signaling device and render search irrelevant. If consumers observe only the distribution of prices, then price dispersion across product pools would imply that consumers update their stopping rules with the information acquired during search. The complexity of this case is beyond the scope of our analysis.

Provision of recommendations. We have assumed that informed consumers in the market freely supply product recommendations, and casual evidence suggests that recommendations are well provisioned in the markets we have considered. A large body of literature has documented several motivations for consumers to contribute to word of mouth processes, see Dellarocas [12] for a related discussion. Consumers may enjoy the opportunity to discuss their preferred entertainment products with others. The existence of such positive network effects on the demand side of artistic markets was proposed by Adler [1]. In our model, consumers providing recommendations derive no direct benefit (nor cost) in the process, but benefit indirectly from lower prices. Since consumers seeking recommendations incur a sunk cost r on each draw, they would be willing to reward those that provide them instead. Avery et al. [6] explore reward mechanisms for the optimal provision of recommendations. But assuming r is a sunk cost instead of a transfer allows us to ignore the bargaining problem that could arise between consumers.

We have assumed recommendations enjoy no salience, as consumers do not place additional value on a match that results from a recommendation. Senecal and Nantel [31] report a series of experiments that suggest recommendations have an influential effect on consumers beyond awareness. In our framework, salient recommendations would increase the expected utility consumers derive from recommendations $u_r^{t,i}$, increasing consumer participation and the share of consumers searching with recommendations. Hence salience would reinforce our findings on sales concentration.

7 Conclusion

We have provided a theoretical framework to understand the impact of information provision on consumer search. We have examined the firm's incentives to supply product information that allows consumers to better evaluate products, to provide a platform for consumers to exchange product recommendations, and to develop mechanisms to generate personalized product recommendations. We have shown that all the above reduce consumer search costs in the market, enabling the firm to extract more consumers surplus. The findings illustrate why it is profitable for the firm to become an information gateway for products, a role adopted by successful online retailers such as Amazon for books or NewEgg for electronics.

Our model explains the value of consumer-to-consumer communication in markets characterized by large assortments of horizontally differentiated products. This contributes to its prevalence in the markets considered, such as music, cinema, literature or video game entertainment. These markets are also characterized by a high concentration of sales, and our model explains why consumer-to-consumer communication tends to increase sales concentration. Another finding is that, due to the mechanisms that drive the exchange of information among consumers, those with uncommon preferences in the population and the products that appeal to them are underserved in the market. Thus there is an opportunity for firms that can help connect these consumers and products.

Personalization mechanisms such as recommender systems, that generate personalized products recommendations for consumers, are a prominent example of how firms can play an active role to reduce search costs in the market. Major online retailers have pioneered their implementation in their storefronts, and personalization will benefit most firms that have large product assortments and access to rich datasets on consumer preferences. Our model explains why personalization can reduce the concentration of sales in the market, having a stronger impact on consumers with uncommon preferences in the population, and thus contributing to generate a long tail effect. To understand lower sales concentration online, we need not only consider the depth of product assortments but also the personalization features that help consumers navigate them and update their product search strategies.

Personalization mechanisms can sustain a competitive advantage in the market. This will be the case if firms offering better personalization can capture a share of the value they generate. In real world applications, the existence of switching costs and network effects suggests that firms can design strategies to achieve this. For example, recommender systems exhibit a learning curve to identify the preferences of a consumer, and benefit from large datasets of consumer activity to improve their accuracy. So consumers will receive less accurate recommendations when patronizing a new firm and, in general, when patronizing smaller firms. Both factors suggest a firm can benefit from rewarding consumers to join the system and grow its userbase, generating a lock-in effect to outperform competitors.

The potential of personalization schemes to reduce the concentration of sales drives other strategic considerations. Since personalization increases the demand for products in the tail of the sales distribution, firms with low inventory costs stand to benefit the most. These firms can increase the depth of their assortment beyond that of competitors, ensuring competitors cannot serve consumers demanding products in the tail. For instance, in the case of Netflix and Blockbuster, a large share of the movies Netflix recommends to customers are not available in Blockbuster stores. By generating personalized recommendations that help consumers navigate its assortment, Netflix is also maximizing the value of stocking a deeper assortment than competitors.

The value of personalization mechanisms in the marketplace cannot be underestimated. Our results provide a rationale for the provision of unbiased product recommendations, and new business models have emerged in recent years with the promise to deliver them. This is the case of recommendation services such as Pandora for music or FilmAffinity for movies, or that of third-party providers of personalization tools such as Strands. Their potential to generate value will grow with the amount of information consumers disclose online, and may soon be limited only by the privacy concerns of consumers themselves.

References

[1] Adler, Moshe (1985), 'Stardom and Talent,' *American Economic Review* 75:1 208-212

[2] Adomavicius, Gediminas and Alexander Tuzhilin (2005), 'Towards the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions,' *IEEE Transactions on Knowledge and Data Engineering* 17:6 734-749

[3] Anderson, Chris (2004), 'The Long Tail,' *Wired*.

[4] Anderson, Simon P. and Régis Renault (1999), 'Pricing, product diversity, and search costs: a Bertrand-Chamberlin-Diamond model,' *RAND Journal of Economics* 30:4 719-735

[5] Anderson, Simon P. and Régis Renault (2006), 'Advertising Content,' *American Economic Review* 96:1 93-113

[6] Avery, Chris, Paul Resnick and Richard Zeckhauser (1999), 'The Market for Evaluations,' *American Economic Review* 89:3 564-584

[7] Bakos, Yannis J. (1997), 'Reducing Buyer Search Costs: Implications for Electronic Marketplaces,' *Management Science* 43:12 1676-1692

[8] Brynjolfsson, Erik, Yu Jeffrey Hu and Duncan Simester (2007), 'Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales,' *Working Paper*

[9] Brynjolfsson, Erik, Yu Jeffrey Hu and Michael D. Smith (2006), 'From Niches to Riches: Anatomy of the Long Tail,' *MIT Sloan Management Review* 47:4 67-71

[10] Chernev, Alexander (2003), 'When More Is Less and Less Is More: The Role of Ideal Point Availability and Assortment in Consumer Choice,' *Journal of Consumer Research* 30 170-183

[11] Chevalier, Judith A. and Dina Mayzlin (2006), 'The Effect of Word of Mouth on Sales: Online Book Reviews,' *Journal of Marketing Research* 43 345-354.

[12] Dellarocas, Chrysanthos (2006), 'Strategic Manipulation of Internet Opinion Forums: Implications for Consumers and Firms,' *Management Science* 52:10 1577-1593

[13] Ehrenberg, Andrew S. C., Gerald J. Goodhardt and T. Patrick Barwise (1990), 'Double Jeopardy Revisited,' *Journal of Marketing* 54:3 82-91

[14] Ehrmann, Thomas and Hendrik Schmale (2008), 'The Hitchhiker's Guide to the Long Tail: The Influence of Online-Reviews and Product Recommendations on Book Sales – Evidence from German Online Retailing,' *ICIS 2008 Proceedings*

[15] Elberse, Anita (2008), 'Should you invest in the long tail?' *Harvard Business Review* 86:7/8 88-96

[16] Elberse, Anita and Felix Oberholzer-Gee (2008), 'Superstars and Underdogs: An examination of the Long Tail Phenomenon in Video Sales,' *Working Paper*

[17] Fleder, Daniel and Kartik Hosanagar (2007), 'Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity,' *Management Science* 55:5 697-712

[18] Goel, Sharad, Andrei Broder, Evgeniy Gabrilovich and Bo Pang (2010), 'Anatomy of the Long Tail: Ordinary People with Extraordinary Tastes,' *Third ACM International Conference on Web Search and Data Mining*

[19] Hagiu, Andrei and Bruno Jullien (2010), 'Why Do Intermediaries Divert Search?' *Working Paper*

[20] Hall, Marshall and Nicolaus Tideman (1967), 'Measures of Concentration,' *Journal of the American Statistical Association* 62:317 162-168

[21] Iyengar, Sheena S. and Mark R. Lepper (2000), 'When Choice is Demotivating: Can One Desire Too Much of a Good Thing?' *Journal of Personality and Social Psychology* 79:6 995-1006

[22] Johnson, Justin P. and David P. Myatt (2006), 'On the Simple Economics of Advertising, Marketing, and Product Design,' *American Economic Review* 96:3 756-784

[23] Kuksov, Dmitri and J. Miguel Villas-Boas (2010), 'When More Alternatives Lead to Less Choice,' *Marketing Science* 29:3 507-524

[24] Kuksov, Dmitri and Ying Xie (2010), 'Pricing, Frills, and Customer Ratings,' *Marketing Science* forthcoming

[25] Leskovec, Jure, Lada A. Adamic, and Bernardo A. Huberman (2007), 'The Dynamics of Viral Marketing,' *ACM Transactions on the Web* 1:1

[26] MacDonald, Glenn M. (1988), 'The Economics of Rising Stars,' *American Economic Review* 78:1 155-166

[27] Oestreicher-Singer, Gal and Arun Sundararajan (2010), 'Recommendation Networks and the Long Tail of Electronic Commerce,' *Working Paper*

[28] Resnick, Paul and Hal R. Varian (1997), 'Recommender Systems,' *Communications of the ACM* 40:3 56-58.

[29] Rosen, Sherwin (1981), 'The Economics of Superstars,' *American Economic Review* 71:5 845-858

[30] Salganik, Matthew J., Peter Sheridan Dodds and Duncan J. Watts (2006), 'Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market,' *Science* 311 854-856

[31] Senecal, Sylvain and Jacques Nantel (2004), 'The influence of online product recommendations on consumers' online choices,' *Journal of Retailing* 80 159-169

[32] Tan, Tom F. and Serguei Netessine (2009), 'Is Tom Cruise Threatened? Using Netflix Prize Data to Examine the Long Tail of Electronic Commerce,' *Working Paper*

[33] Tucker, Katherine and Juanjuan Zhang (2009), 'How Does Popularity Information Affect Choices? Theory and A Field Experiment,' *Working Paper*

[34] Zhu, Feng and Xiaoquan Zhang (2010), 'Impact of Online Consumer Reviews on Sales: The Moderating Role of Product and Consumer Characteristics,' *Journal of Marketing* 74 133-148