Search, Product Recommendations, and Sales Concentration*

Andres Hervas-Drane†

June 7, 2012

Abstract

Personalization mechanisms such as recommender systems play an important role in online retail, and recent empirical contributions find that they reduce the concentration of sales and increase sales volume within product assortments. But the underlying drivers of these effects are not yet well understood. I present a model to explain the role of product recommendations on consumers' product discovery process and their implications for the firm. Consumers face a search problem within an assortment of horizontally differentiated products supplied by a monopolist, and may search for a product match by drawing products from the assortment or by seeking product recommendations from other consumers. I analyze the underlying consumer interactions that lead to the emergence of word of mouth, provide a rationale for the firm's adoption of personalization mechanisms such as recommender systems which generate personalized recommendations, and evaluate the impact on sales concentration, sales volume, firm prices, and firm profits. The model explains how personalization mechanisms contribute to lower the concentration of sales and is well 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, the CRES Strategy Conference at Olin Business School, Universitat Autònoma de Barcelona, the ZEW Online Markets Workshop, and the Max Planck's Strategy Theories Conference. The main body of this research was conducted at Harvard University. Financial support from the HBS Division of Research, the Real Colegio Complutense, the NET Institute, the Caja Madrid Foundation, and grant ECO2009-11307 of the Spanish Ministry of Science and Innovation are gratefully acknowledged.

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

"Today, online commerce saves customers money and precious time. Tomorrow, through personalization, online commerce will accelerate the very process of discovery." Jeffrey P. Bezos, 1997[1]

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. Electronic commerce is also reducing the concentration of sales within product assortments, increasing the market share of products catering to niche audiences. The main explanation for this phenomenon has focused on supply side factors, given 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 contributing to drive down sales concentration online. Brynjolfsson, Hu, and Simester [9] and Elberse and Oberholzer-Gee [22] 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 yet well understood.

This paper presents a formal model of consumer search that explains how demand side factors contribute to explain lower sales concentration in online retail. I examine how product recommendations impact the product choices of consumers and their participation in the market. The exchange of product recommendations is fundamental to the product discovery process of consumers, and online retailers are increasingly participating in the process with the adoption of personalization mechanisms such as recommender systems, which generate personalized product recommendations for consumers. Recent empirical contributions confirm that product recommendations play an important role in online retail.[2] Chevalier and Mayzlin [14] find that online consumer reviews of books increase relative sales at the retailer they are posted on. Feng and Zhang [41] find that online consumer reviews of videogames have a stronger impact on niche products. Brynjolfsson, Hu, and Simester [9] show with server log data that the recommender system is the major contributor to the lower online sales concentration observed in their study. Oestreicher-Singer and Sundararajan [33] find that sales concentration is lower among book categories on Amazon.com where personalization is expected to be more accurate. De, Hu, and Rahman [17] use server log data to show that a recommender system increases sales volume.

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

[2] Offline sales are also affected, with analysts estimating 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.

The modeling exercise presented here is based on the observation that online personalization mechanisms are substituting traditional word of mouth processes. The findings address several key questions. First, they explain consumers' demand (and supply) of product recommendations in the market. Second, they explain the impact of traditional word of mouth processes where consumers directly exchange product recommendations, which are found to increase sales concentration and sales volume within product assortments. And third, they explain how the introduction of personalization mechanisms such as recommender systems reduces sales concentration and further increases sales volume, rationalizing the findings reported above in the recent empirical literature. The model provides a novel explanation for several aspects of consumer product discovery and informs the design of marketing strategies that exploit it.

I proceed by considering a market where a monopolist supplies a large assortment of horizontally differentiated products and consumers face a search problem to identify products that match their taste prior to purchase. The monopolist supplies the assortment under a uniform pricing scheme and may be an online media store such as those of Apple or Amazon. The monopoly case is of interest because it allows for a direct evaluation of the value captured by the firm in the presence of consumer search, and which would otherwise be eroded by competition. I enrich the demand side of the market by allowing consumers to differ in both their product preferences and their search costs. Consumers arrive to the market uninformed about products 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. Consumers may search for a preferred product by directly sampling products from the assortment or by seeking product recommendations from others who have sampled products before them. The construction is well suited for experience goods such as music, films, books, or video games. These product categories exhibit large assortments with fairly homogeneous pricing, consumer preferences tend to be idiosyncratic, and consumers require direct exposure to products in order to identify their value.

The model reveals that word of mouth arises because of preference correlation across consumers, even if costly. Word of mouth increases sales because consumers that would otherwise not participate in the market now choose to do so by seeking product recommendations from others, as this reduces the expected number of products they need to sample in order to locate a preferred product. Word of mouth increases the concentration of sales within the assortment, and this is characterized by two different effects. First, word of mouth drives consumers to mass market products – those over which most consumers agree – as they are recommended more often in the word of mouth exchange and such recommendations carry a higher success rate. Second, word of mouth benefits mainstream consumers the most – those whose preferences are more widespread in the population – and the products that appeal to them benefit as a result.

When the firm introduces personalization mechanisms such as recommender systems, a distinctive feature of the online channel, consumers can obtain personalized product recommendations that account for their preferences. I provide insight on how the collaborative filtering algorithms that drive such systems can be interpreted as intermediating a word of mouth exchange, matching consumers with similar product preferences. And I argue that, to understand the impact of this technology, one should account for the search strategies (traditional word of mouth) that consumers employ in its absence. I show that this matching mechanism improves over the value of word of mouth, increasing sales by further reducing the average number of products consumers need to sample during their search, and reducing the concentration of sales by benefiting niche consumers with less prevalent preferences the most and thereby reverting the word of mouth effects on concentration discussed above. The model provides a rationale for the recent findings in the empirical literature, and explains the firm's investment in personalization mechanisms as well as consumers' willingness to embrace them in their product discovery process.

To derive these results, several simplifications are made in the analysis. On the supply side, assortment composition is assumed exogenous. The monopolist supplies all available product varieties and cannot restrict the length of the assortment, or is unwilling to do so to reduce consumer search costs. This could be due to threat of entry by a competitor, for instance. On the demand side, consumers sample with replacement over a continuum of products, so stopping rules do not depend on past search history or sunk search costs. I restrict consumer search strategies to sampling products from the assortment or seeking product recommendations from others. Search is meaningful because consumers cannot directly identify their preferred products; market shares and product popularity information are not available or readily observable.

The paper relates to the recent literature on consumer search and e-commerce. Kim, Albuquerque and Bronnenberg [27] estimate the consumer's search problem based on camcorder sales data retrieved from Amazon.com, and consistent with the results presented here, find that consumer search costs have a significant impact in the market and that Amazon's product recommendations lower them. Chen, Wang and Xie [12] study Amazon sales data to disentangle the effect on consumer demand of word of mouth from that of observational learning based on sales rankings. Sun [37] examines the informational role of consumer product ratings and shows that niche products are associated with higher variance of ratings. Choi and Bell [16] consider the benefits of e-commerce for preference minorities, consumers who are not well served by local brick and mortar stores due to the constraints of physical distribution. The findings reported here also suggest that consumers with niche preferences and the products that appeal to them benefit the most from online personalization.

To the best of my knowledge, no previous theoretical work has explored the links between word of mouth, personalization, and sales concentration. Bar-Isaac, Caruana and Cuñat [6] model how reductions in consumer search costs affect product design choices on the supply side of the market, which can lead to lower sales concentration by increasing the market shares of firms with rare designs. Fleder and Hosanagar [23] evaluate the impact of different recommender systems on sales concentration and volume by using simulations with consumers and products located on a

2-axis space. Other contributions have focused on improving the performance of personalization mechanisms by applying Bayesian learning or other methods, as in Ansari, Essegaier and Kohli [2], Ying, Feinberg and Wedel [40], Bodapati [8], and Atahan and Sarkar [4]. A large strand of literature in marketing has examined the impact of word of mouth. Recent contributions include that of Berger and Schwartz [7], who analyze the drivers of word of mouth across product categories and over time. Chen, Harper, Konstan and Li [11] study how social comparisons can boost word of mouth contributions in user communities. Aral and Walker [3] examine how the viral features of products can foster word of mouth. Chintagunta, Gopinath and Venkataraman [15] analyze how online user reviews impact the box office performances of movies through their geographical rollout. Manchanda, Xie and Youn [31] evaluate the comparative performance of word of mouth and marketing communication in the pharmaceutical industry. Cheema and Kaikati [10] evaluate how consumers' concerns for uniqueness impact word of mouth.

The paper is organized as follows. The next section introduces the building blocks of the search model. Section 3 solves the simplest instance of search when there is no word of mouth, which serves as a benchmark for the remaining of the analysis. Section 4 introduces word of mouth and allows consumers to exchange product recommendations. Section 5 considers the impact of personalization, when the firm supplies personalized product recommendations to consumers. I review the assumptions of the model and discuss some extensions in Section 6, and conclude in Section 7.

2 The model

Consider a market where a monopolist supplies a product assortment consisting of a continuum of products of measure one and quotes a common price p for all products. In the market there is a unit mass of consumers, and the preferences of consumers over products are simplified to a binary classification – a given 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 derives zero utility. Consumers exhibit unit demand, and may participate in the market to purchase and consume a preferred product or stay out.

Consumers differ on two separate dimensions: their preferred products and their costs of sampling products. The simplest instance of the model that yields the results is that where some product preferences are more prevalent than others in the population. In addition, the insights derived from the model are richer when consumers with different product preferences also agree on some subset of products.[3] To this end, I partition the product space into N equivalent product pools, which can be understood as product varieties, and consider $T = N - 1$ consumer types. Let

[3]The presence of products over which all consumers agree is not crucial to the results, but is useful to illustrate the concentration effects that arise when there is overlap in the preferred product space of different consumer types. Alternative preference specifications are discussed in Section 6.

consumers of type $t \in \{1, ..., T\}$ prefer products pertaining to product pools t and N. Thus all consumers derive utility u from any product in pool N but consumers of different types disagree over the remaining. I will assume throughout that $T \geq 3$ and refer to T as a measure of *taste diversity*, since the larger the value of T the more heterogeneous are consumer preferences and the more differentiated the product space.

To capture the fact that some product preferences are more prevalent in the population, denote the share of consumers of type t by s^t where $\sum s^t = 1$ and $s^t < s^{t+1}$. That is, 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. In particular, pool N is composed of *mass market* products that appeal to all consumers.

When entering the market, consumers are fully informed except that they cannot map individual products to product pools.[4] So consumers have some information about the assortment, akin to a notion of taste, but cannot readily identify their preferred products. All products are ex-ante identical and as a result consumers face a search problem in order to locate a preferred product. Consumers can become informed about products by sampling them. A product *match* is achieved when a preferred product is identified. Sampling products is costly, and consumers will form a rational expectation of the value of search. Sampling costs are uniformly distributed in the consumer population independently of product preferences, where the cost of consumer i is given by $c^i \sim U[0, \overline{c}]$. Thus sampling a product which does not yield a match incurs disutility c^i, and sampling and consuming a product match yields utility $u - c^i$. For experience goods, this can be interpreted as the time investment required to experience the good. I assume $\overline{c}$ is sufficiently high to ensure the market remains uncovered.[5]

To summarize the model:

- The monopolist supplies an assortment with N product pools and prices all products at p.
- There are $T = N - 1$ 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 $\sum s^t = 1$ and $s^t < s^{t+1}$.
- Consumers cannot map products to products pools and form a rational expectation of the value of search.

[4]This information structure can be further relaxed if consumers learn the value of search strategies based on past search experience. I.e., if consumers have previously observed their match probability with evaluations β (1) and with recommendations α^t (8), they do not require information about the preference structure of the consumer population to form correct expectations.

[5]This requires $\overline{c} > (u-r)/2$ throughout the analysis, where r is the cost of seeking a recommendation introduced in Section 4. The assumption simplifies the analysis by avoiding corner solutions in the pricing game, as a positive mass of consumers will not participate in the market in equilibrium.

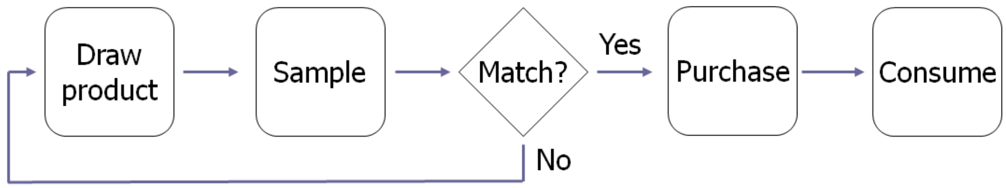

Figure 1: The sequential search process of consumers when searching with evaluations by sampling products directly from the assortment.

- Sampling costs are uniformly distributed in the consumer population, $c^i \sim U[0, \overline{c}]$.

3 Search benchmark

I start the analysis by considering the simplest instance of search in the model. This is the case where there is no word of mouth and consumers search only with their own product evaluations. I consider a two-stage game where the monopolist first chooses the price level in the market p and consumers search for a match in the second stage by sequentially drawing and sampling products from the assortment. Figure 1 depicts the sequential search process faced by consumers. On each draw, consumers incur sampling cost c^i and will execute a purchase at price p if they locate a match.[6]

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 isolates variations in the concentration of sales (or market share variations) from volume effects driven by shifts in the extent of consumer participation. In particular, it is of interest to analyze how product recommendations affect the concentration of sales. To compare concentration across sales distributions I 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.[7]

Consumer search strategy. I proceed by backwards induction and consider the search problem faced by consumers in the second stage given a price level p. The only available search strategy is to sequentially sample products until a match is located. A consumer of type t will only obtain a product match when drawing a product from pools t or N. It is useful to define

[6] The model also applies to a market where consumers can realize costless returns of undesired products.

[7] See Hall and Tideman [25] for an analysis of the desirable properties of a concentration index.

indicator function λ as $\lambda_n^t = 1$ if $n = t$ or $n = N$, and $\lambda_n^t = 0$ otherwise. Since consumers draw products randomly from the assortment, the probability of drawing from any given pool is $1/N$. The match probability on each draw, denoted by β, is then given by

$$\beta = \sum\nolimits_n (1/N)\lambda_n^t = \frac{2}{N}. \tag{1}$$

Search with replacement implies that each draw is a Bernoulli trial with success probability β, which is common for all consumers. The expected utility of a new product draw for an unmatched consumer with sampling cost c^i is

$$u_e^i = \beta(u - p) - c^i, \tag{2}$$

given that the consumer only purchases if a match is located but incurs 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 and is unaffected by past unsuccessful draws. Hence the consumer of each type which is strictly indifferent between evaluating products and not participating can be identified by equating u_e^i to zero. Denote the indifferent evaluator by c_e^i,

$$c_e^i = \beta(u - p). \tag{3}$$

Only consumers with a sampling cost $c^i \leq c_e^i$ choose to search, and participation is homogeneous across types. These consumers search by sequentially drawing products until they obtain a match, which on average requires $1/\beta$ draws. The search process finalizes once a match is located; searching for a second match cannot be optimal. Consumers with a higher sampling cost prefer not to participate in the market.

Firm pricing. I next turn to the first stage of the game and solve the firm's problem given the consumer participation constraint for all types (3). Firm profits are then given by

$$\pi_e = \frac{c_e^i}{\bar{c}} p = \frac{\beta(u - p)}{\bar{c}} p. \tag{4}$$

Solving for the firm's optimal price obtains

$$p_e = \frac{u}{2}. \tag{5}$$

Sales concentration. I next characterize the sales distribution with evaluations, denoted by $\sigma = \{\sigma_n\}_{n=1}^N$. Let s_e^t be the share of consumers of type t among the mass of consumers that searches with evaluations. I proceed by characterizing separately the sales distribution generated by each consumer type $\sigma^t = \{\sigma_n^t\}_{n=1}^N$, where $\sigma_n = \sum_t s_e^t \sigma_n^t$. To characterize σ^t note that the sales distribution generated by consumers of type t must equal their distribution of matches over products.

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 σ^t can be derived by characterizing the distribution of matches over products for a single evaluation of a consumer of type t. Using indicator function λ, 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

$$\sigma_n^t = \frac{(1/N)\lambda_n^t}{\beta} = \begin{cases} \frac{1}{2} & \text{if } n = t \text{ or } n = N \\ 0 & \text{otherwise} \end{cases}. \tag{6}$$

Proposition 1. *When consumers search with own product evaluations, consumers of all types with sampling cost $c^i \in [0, c_e^i]$ participate in the market, where $c_e^i = u/N$ and the share of participating consumers of each type is proportional to their prevalence in the population $s_e^t = s^t$. The firm prices products at $p_e = u/2$ and the sales distribution in the market is characterized by the following market shares:*

$$\sigma_n = \begin{cases} \frac{s_e^n}{2} & \text{if } n \in (1, N-1) \\ \frac{1}{2} & \text{if } n = N \end{cases}. \tag{7}$$

Consumers searching with own product evaluations decide to participate or not in the market based on product prices and their idiosyncratic sampling costs. As there is no word of mouth there is no interaction among consumers, so the prevalence of a consumer's preferences in the population does not affect her search strategy and participation in the market is homogeneous across consumer types. Participation pays off for consumers with lower sampling costs, who choose to search by repeatedly sampling products from the assortment until they locate a match. Consumer participation in the market increases with match utility u and decreases with sampling costs $\bar{c}$. A higher taste diversity T reduces consumer participation, as it implies that locating a match within the assortment requires on average more product draws. The sales distribution in the market is characterized by market shares which are a linear function of s^t across product pools and increasing in n, reflecting the product preferences of the population.

The firm's demand is downward sloping, as expected. The lower the firm's prices, the more consumers are willing to participate. The firm's profits increase with match utility u and decrease with taste diversity T and sampling costs $\bar{c}$, due to their impact on consumer participation. The latter implies that the firm has incentives to lower consumers' sampling costs 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

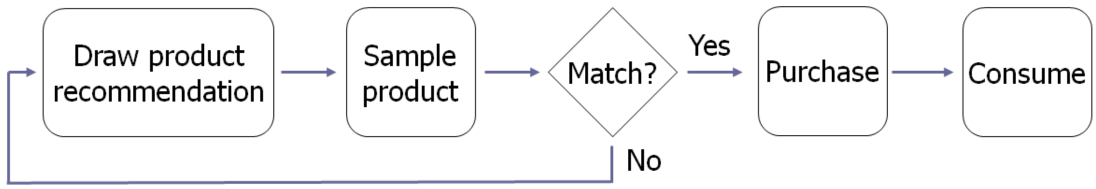

Figure 2: The sequential search process of consumers seeking product recommendations from others who previously completed their search by sampling products from the assortment.

from their product pages, becoming information gateways for products. The model suggests this increases the willingness of consumers to search within the assortment, allowing the firm to sustain higher prices and derive higher profits.

4 Word of mouth

Consumers stand to benefit from exchanging information about the products matches they uncover through search. I introduce word of mouth in the model by adding a third stage to the search benchmark. In the first stage, the monopolist chooses the price level in the market p. Consumers willing to participate in the market then choose between two mutually exclusive search strategies, and may either search during the second stage or the third stage. In the second stage, consumers may search for a match with own product evaluations as in the previous section, depicted in Figure 1. In the third stage, consumers may search for a match by seeking product recommendations from evaluating consumers who searched before them in the second stage, following the process depicted in Figure 2.

Recommendations are drawn randomly from the mass of evaluating consumers, who identify the product they matched with in the second stage when supplying a recommendation.[8] Since obtaining recommendations requires an additional step in the search process, demanding them is costly and each recommendation draw incurs a common cost r for all consumers. To ensure that recommendations hold in the market, I require $r < u/4$. For simplicity, I assume consumers supplying recommendations do so for free, and the analysis will reveal that they benefit from doing so. Nonetheless, the model is also robust to the presence of communication costs incurred by consumers supplying recommendations.[9]

[8] This communication can be understood to take place online or offline. In the first case, the firm provides a platform for evaluating 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. Also note that the reference to the product an evaluating consumer matched with is the core element of a product recommendation. References about which products *not* to sample, e.g. negative reviews or ratings, are not valuable in this setting. As there is a continuum of products in the assortment, discarding a finite number of products does not increase the probability of locating a match. The result carries over to discrete product spaces when the assortment is sufficiently large.

[9] The model is robust to communication costs as long as consumers demanding recommendations compensate those who supply them. That is, parameter r can be interpreted as a transfer from each consumer demanding a

I proceed by backwards induction and start by considering the search strategies of consumers in the third and second stages. The complexity of the problem increases with word of mouth because of the interdependency that arises between types, and search strategies are determined by a system of implicit equations. A closed form solution is not obtained but the properties of the solution are characterized below. I then to turn to the firm's pricing problem, and find that the demand function now exhibits kinks at the price points that trigger consumers to switch search strategies (it becomes convex). Optimal prices are characterized as a function of consumer search strategies and some threshold conditions on prices are identified. Finally, I turn to the impact of word of mouth on sales concentration. The precise sales distribution cannot be pinned down, but the characterization of search strategies is sufficient to analyze the direction of the market share shifts introduced by word of mouth (with respect to the sales distribution characterized in Proposition 1). I characterize the pattern of shifts across product pools and show that sales concentration always increases.

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 $\sigma = \{\sigma_n\}_{n=1}^{N}$ (7) carries over from the 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

$$\alpha^t = \sigma_t + \sigma_N = \frac{1 + s_e^t}{2}. \tag{8}$$

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. I 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 an unmatched consumer of type t with sampling cost c^i is

$$u_r^{t,i} = \alpha^t(u - p) - r - c^i, \tag{9}$$

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

recommendation to the consumer who supplies it when the latter incurs a communication cost r. I further discuss the factors surrounding the provision of recommendations in Section 6.

the consumer of type t which is strictly indifferent between seeking recommendations and not participating can be identified by equating $u_r^{t,i}$ to zero. Denote the indifferent recommendation seeker of type t by $c_r^{t,i}$, where

$$c_r^{t,i} = \alpha^t(u-p) - r. \tag{10}$$

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.

I 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 unsuccessful draws. This implies that no consumer that chooses to search with evaluations in the second stage will ever prefer to abort the search in order to search with recommendations in the third.

The indifferent evaluator of type t, denoted by $c_e^{t,i}$, is identified by equating the expected utility derived from both search strategies in order to locate a match, $u_r^{t,i} = u_e^i$. The expected utility of searching with evaluations u_e^i (2) carries over from the previous analysis and is type-independent. The expected number of draws required for a match with evaluations and with recommendations are given by $1/\beta$ and $1/\alpha^t$ respectively. Therefore,

$$u - p - \frac{r + c_e^{t,i}}{\alpha^t} = u - p - \frac{c_e^{t,i}}{\beta}, \tag{11}$$

and the indifferent evaluator of type t is identified by

$$c_e^{t,i} = \frac{\beta r}{\alpha^t - \beta}. \tag{12}$$

Consumers of type t with sampling cost $c^i \in [0, c_e^{t,i}]$ prefer to search with evaluations in the second stage over seeking recommendations. And consistency requires 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, a boundary price $\overline{p}^t$ can be identified by equating $c_e^{t,i} = c_r^{t,i}$,

$$\overline{p}^t = u - \frac{r}{\alpha^t - \beta}. \tag{13}$$

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 (3), following the previous analysis. Note that participation is homogeneous across types that search only with evaluations.

I 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 $c^i \in (c_e^{t,i}, c_r^{t,i}]$ search with 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.

I 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. I next argue that the solution to this system satisfies that $c_e^{t,i}$ and s_e^t are strictly decreasing and strictly increasing in t, respectively, for types that search with recommendations. I show this by contradiction. Assume recommendations hold for two types, t and $t+1$, and consider the case $c_e^{t,i} \leq c_e^{t+1,i}$. This requires that $\alpha^t \geq \alpha^{t+1}$ by (12), which then implies that $s_e^t \geq s_e^{t+1}$ by (8). 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 $c_e^{t,i} \leq c_e^{t+1,i}$ both imply $s_e^t < s_e^{t+1}$, which 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 $t+1$.

Some conclusions can now be drawn 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 strictly 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 strictly increasing in t, then α^t must also be strictly increasing in t. The latter implies that $c_r^{t,i}$ and $\overline{p}^t$ are strictly increasing in t, so s_r^t must also be strictly 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 $j > t$. Figure 3 depicts the utility of search strategies implied by the solution.

Firm pricing. I next turn to the first stage of the game and analyze the firm's pricing problem. Given a price level p in the market, I 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 $t_r(p)$ 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

$$\pi_r = [\sum_{t=1}^{t_r(p)-1} \frac{c_e^i}{\overline{c}} s^t + \sum_{t=t_r(p)}^{T} \frac{c_r^t}{\overline{c}} s^t] p. \tag{14}$$

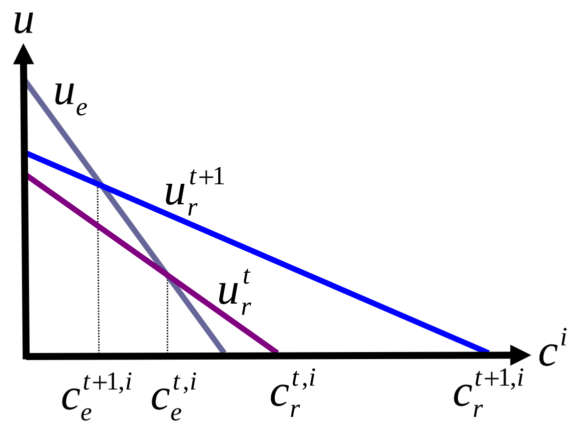

Figure 3: Utility of search strategies for consumers of types t and $t+1$ as a function of their sampling cost c^i when recommendations hold for both types. The utility of searching with evaluations u_e^i is common for all types, but the utility of searching with recommendations is higher for consumers with more prevalent preferences in the population, $u_r^{t+1,i} > u_r^{t,i}$.

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 search strategies across types when the marginal type seeking recommendations is t,

$$\widehat{\alpha}^t = \frac{\sum_{j=t}^{T} s^j}{\sum_{j=1}^{t-1} s^j \beta + \sum_{j=t}^{T} s^j \alpha^j}, \tag{15}$$

where $\widehat{\alpha}^t > 0$. For each component of demand such that $t_r \in \{1, ..., T\}$, the maximum of the corresponding profit curve can be derived from (14), and denoted by $\widehat{p}^t$, where

$$\widehat{p}^t = \frac{u - r\widehat{\alpha}^t}{2}. \tag{16}$$

For the component in which $t_r = T+1$, consumers search only with evaluations and $\widehat{p}^{T+1} = p_e$ as in (5).

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. The 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.[10] Therefore,

[10] This requires the maximum for the component without recommendations to not be well defined, $\widehat{p}^{T+1} < \overline{p}^T$, which implies $r < u(\alpha^T - \beta)/2$. 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 $r < u/4$ is sufficient to ensure recommendations hold in

the solution satisfies

$$p_r = \begin{array}{l} ArgMax_{\widehat{p}^t} \;\; \pi_r(\widehat{p}^t) \\ s.t.\; t_r(\widehat{p}^t) = t \text{ for some } t \in \{1, ..., T\} \end{array} . \tag{17}$$

Sales concentration. I next characterize the sales distribution with word of mouth, denoted by $\rho = \{\rho_n\}_{n=1}^N$. Let s_{er}^t 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). Consider the marginal type t_r that searches with recommendations, such that types $t \in \{1, ..., t_r - 1\}$ search only with evaluations and types $t \in \{t_r, ..., T\}$ search with both evaluations and recommendations. Consumers of type t searching with evaluations generate sales distribution $\sigma^t = \{\sigma_n^t\}_{n=1}^N$ (6). Denote by $\mu^t = \{\mu_n^t\}_{n=1}^N$the sales distribution generated by consumers of type $t \in \{t_r, ..., T\}$ searching with recommendations. And denote by S_e and S_r the shares of participating consumers that search with evaluations and recommendations, respectively. I can write ρ_n as

$$\rho_n = \begin{cases} S_e\, s_e^n\, \sigma_n^n & \text{for } n \in \{1, ..., t_r - 1\} \\ S_e\, s_e^n\, \sigma_n^n + S_r\, s_r^n\, \mu_n^n) & \text{for } n \in \{t_r, ..., T\} \\ \sum_{t=1}^T S_e\, s_e^t\, \sigma_N^t + \sum_{t=t_r}^T S_r\, s_r^t\, \mu_N^t & \text{for } n = N \end{cases} . \tag{18}$$

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 I 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 $\sigma_n \lambda_n^t$, where σ (7) identifies the distribution of matches of evaluating consumers over product pools, and the probability of a match over all products is given by α^t. This implies

$$\mu_n^t = \frac{\sigma_n \lambda_n^t}{\alpha^t} = \begin{cases} \frac{s_e^t}{1+s_e^t} & \text{if } n = t \\ \frac{1}{1+s_e^t} & \text{if } n = N \\ 0 & \text{otherwise} \end{cases} . \tag{19}$$

So $\mu_t^t < 1/2$ and $\mu_N^t > 1/2$. And I can now establish that ρ_n (18) is strictly increasing in n, given that both s_e^t and s_r^t are strictly increasing in t in equilibrium, and μ_t^t is also strictly increasing in t for $t \in \{t_r, ..., T\}$.

Unfortunately, because closed form solutions are not available for s_e^t and s_r^t, I cannot directly compare ρ and σ to determine the impact of word of mouth on sales concentration. Instead, I

equilibrium.

proceed by evaluating the market share shifts driven by the introduction of word of mouth and argue that concentration always increases. I show this in two steps. Starting from σ, I first account for the shift in consumer participation driven by recommendations (the *participation effect*) while keeping fixed the per-type sales distribution σ^t. To do so, I derive a participation-adjusted sales distribution $\overline{\rho} = \{\overline{\rho}_n\}_{n=1}^N$, where $\overline{\rho}_n = \sum_t s_{er}^t \sigma_n^t$, and evaluate how concentration differs from σ. In the second step, I account for the shift in the sales distribution generated by consumers seeking recommendations (the *mass market effect*), and evaluate how concentration differs between ρ and $\overline{\rho}$.

To account for the participation shift, I can directly write $\overline{\rho}$ as

$$\overline{\rho}_n = \sum\nolimits_t s_{er}^t \sigma_n^t = \begin{cases} \frac{s_{er}^n}{2} & \text{if } n < N \\ \frac{1}{2} & \text{if } n = N \end{cases}. \tag{20}$$

To see how $\overline{\rho}$ differs from σ, I need to compare how s_{er}^t with word of mouth compares with s_e^t in Proposition 1. Note that participation s_{er}^t with word of mouth for types $t \in \{1, ..., t_r - 1\}$ is given by c_e^i, as in Proposition 1, but is larger and increasing in t for types $t \in \{t_r, ..., T\}$. Inspection of $\overline{\rho}$ (20) and σ (7) reveals that this implies (1) a market share transfer from product pools $n \in \{1, ..., t_r - 1\}$ 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 rank product pools according to sales rank, the participation shift unambiguously increases concentration.

I next account for the mass market shift generated by recommendation seekers in order to evaluate how ρ differs from $\overline{\rho}$. Inspection of ρ_n (18) reveals that the mass market shift is driven exclusively by μ^t, as evaluating consumers continue to generate the same sales distribution unaffected by the presence of word of mouth. Inspection of μ^t (19) and σ^t (6) reveals that $\mu_N^t > \sigma_N^t$ and $\mu_t^t < \sigma_t^t$ for $t \in \{t_r, ..., T\}$ (the types for which μ^t is well defined). This implies a market share transfer from product pools $n \in \{t_r, ..., T\}$ to product pool N. Since this shifts market share from lower rank product pools to the highest rank product pool, the mass market shift unambiguously increases concentration.

I conclude that the introduction of word of mouth strictly increases the concentration of sales in the market. In addition, the participation shift implies that $\rho_n < \sigma_n$ for product pools $n \in \{1, ..., t_r - 1\}$, and the mass market shift implies that $\rho_N > \sigma_N$. Given that both ρ_n and σ_n are strictly increasing in n it must be the case, for some for some $\hat{n} \in \{t_r(p_r), ..., T\}$, that $\rho_n < \sigma_n$ for product pools $n \in \{1, ..., \hat{n} - 1\}$, $\rho_{\hat{n}} \geq \sigma_{\hat{n}}$, and $\rho_n > \sigma_n$ for product pools $n \in \{\hat{n} + 1, ..., N\}$.

Proposition 2. *When there is word of mouth in the market, consumers of types* $t \in \{0, ..., t_r(p_r) - 1\}$ *with sampling cost* $c^i \in [0, c_e^i]$ *and consumers of types* $t \in \{t_r(p_r), ..., T\}$ *with sampling costs* $c^i \in [0, c_e^{t,i}]$ *search with own product evaluations, and consumers of types* $t \in \{t_r(p_r), ..., T\}$ *with sampling cost* $c^i \in (c_e^{t,i}, c_r^{t,i}]$ *search with recommendations, where* $c_e^{t,i} < c_e^i < c_r^{t,i}$ *and* $c_e^{t,i}$ *is strictly* *decreasing and $c_r^{t,i}$ is strictly increasing, respectively, in t. Firm prices satisfy $p_r = (u - r\widehat{\alpha}^t)/2$, where $t = t_r(p_r)$ and $\widehat{\alpha}^t > 0$ is a weighted measure of the value of recommendations for the consumer population. The sales distribution is characterized by $\rho_n < \sigma_n$ for product pools $n \in \{1, ..., \hat{n} - 1\}$, $\rho_{\hat{n}} \geq \sigma_{\hat{n}}$, and $\rho_n > \sigma_n$ for product pools $n \in \{\hat{n} + 1, ..., N\}$, for some $\hat{n} \in \{t_r(p_r), ..., T\}$.*

This equilibrium compared to the one defined in Proposition 1 implies that word of mouth increases the concentration of sales. Consumer types with more prevalent preferences benefit the most from word of mouth and participate more in the market, increasing sales volume. The firm reduces prices and profits increase.

When there is word of mouth, consumers' participation in the market depends on the prevalence of their preferences in addition to their idiosyncratic sampling costs. Product recommendations allow consumers to benefit from those that searched before them, and consumers benefit the most from recommendations originating from others who share their product preferences (same type t) as those recommendations always yield a match. Recommendations originating from consumers with different product preferences (different type) only yield a match when referring to mass market products. Given that recommendations are drawn randomly from the mass of consumers that searched with evaluations, the value of recommendations for any given consumer (the probability they yield a match) then depends on the type-composition of evaluating consumers, and interdependencies across types influence the search strategies of consumers. Hence there is a trade-off among types, and the more valuable recommendations are for one type, the less valuable they are for the remaining.

In equilibrium, consumer search strategies are characterized across sampling costs and across types. Consumers with high sampling costs seek recommendations and consumers with low sampling costs prefer to search with evaluations. Recommendations may increase the probability of a match on every draw but are also costly, so consumers with low sampling costs who suffer less from failed draws are better off with evaluations. Consumers with more mainstream preferences (higher types t) are more inclined to search with recommendations than consumers with niche preferences (lower types t). This implies that in equilibrium the mass of evaluating consumers of each type must be increasing in t, as higher types t derive more value from recommendations. This is not surprising given that the share of consumers of each type is increasing in t; there are simply more consumers of the higher types in the population. But the analysis also reveals that consumers of lower types are over-represented among evaluating consumers, precisely because they benefit less from recommendations and search comparatively more with evaluations as a result. Some of the lower types may search exclusively with evaluations. This implies that, although the value of word of mouth is higher for mainstream consumers whose preferences are more prevalent in the population, the advantage is somewhat mitigated by the stronger incentives of niche consumers to search with evaluations.

The value of word of mouth increases with taste diversity T. This increases the share of participating consumers that seek recommendations, as a higher taste diversity lowers the match probability when sampling products more so than when searching with recommendations. The main factor contributing to this divergence is the presence of mass market products, those products over which consumers with different preferences agree on, as the probability of matching with these products is higher with recommendations. The share of participating consumers that search with recommendations also increases with consumption utility u and decreases with recommendation cost r and sampling costs $\bar{c}$.

Word of mouth expands the firm's demand in the low price range. Demand in the high price range is unaffected because high prices preclude word of mouth – no consumers seek recommendations when product prices are high, given that consumers seeking recommendations are those with high sampling costs and are less willing to participate in the market. As a result, the firm reduces prices for word of mouth to hold in equilibrium, and this ensures prices are lower and consumer participation and firm profits are higher than in its absence. The higher the value of recommendations for consumers the larger the reduction in prices by the firm, as the impact of the reduction is more than offset by the increase in participation.[11]

Similarly to lowering sampling costs for consumers, facilitating the exchange of product recommendations by lowering the cost r 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 pioneered mechanisms to facilitate consumer-to-consumer communication on their websites. Chevalier and Mayzlin [14] analyze the impact of online book reviews at Amazon and Barnes and Noble, and find that reviews increase the relative sales at the retailer they are posted on. Feng and Zhang [41] find that online consumer reviews of videogames have a stronger impact on niche products. The findings are consistent with the model, and suggest that part of the market growth spurred by electronic commerce may be attributable to facilitating consumer-to-consumer communication alone.

Word of mouth increases the concentration of sales within the assortment, as illustrated in Figure 4. The driver of this impact is the exchange of recommendations among consumers with different product preferences. The impact can be decomposed in two effects, a *mass market effect* and a *participation effect*, which are respectively characterized across product pools and across consumer types. The mass market effect follows from the fact that consumers of different types agree on some subset of 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 follows from the fact that consumers with widespread preferences exhibit higher participation, as they derive higher value from recommendations. This effect increases the market shares of product pools with widespread appeal and decreases that of

[11]Note that all consumers are better off in the presence of word of mouth, and this provides a rationale for evaluating consumers to willingly supply product recommendations in the market.

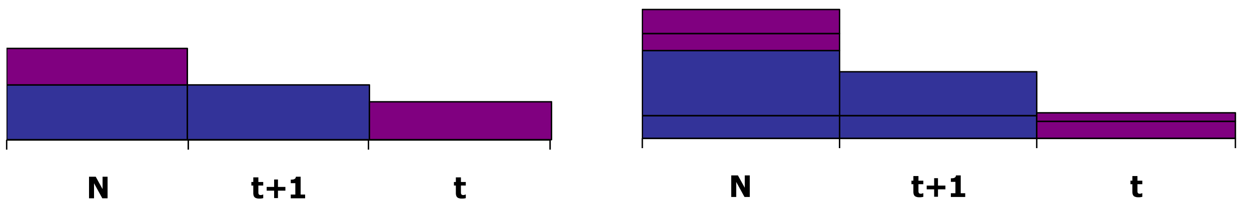

Figure 4: Sales distribution generated by consumers of types t (purple) and $t+1$ (blue) over product pools t, $t+1$, and N – considering only sales generated by these two types – in the search benchmark (left) and with word of mouth (right, with consumers searching with recommendations piled above those searching with evaluations). Word of mouth increases the concentration of sales benefiting mass market products N and consumers of type $t+1$ with more prevalent preferences the most, respectively the mass market effect and the participation effect.

product pools with low appeal. Both effects increase the concentration of sales, and the shift in concentration grows with the share of consumers searching with recommendations.[12]

The mechanisms predicted by the model that lead to an increase in sales concentration have been identified in the empirical literature. Leskovec, Adamic, and Huberman [30] 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. Related findings have been reported for popularity feedback mechanisms that inform consumers of the popularity ranking of products. Salganik, Dodds and Watts [35] study demand over a set of rare songs offered to test subjects on the Internet, and Tucker and Zhang [39] analyze the click-through rates of a webpage indexing marriage agencies, and in both cases popularity feedback increases concentration and consumer participation. The findings are reminiscent of the double jeopardy effect discussed by Ehrenberg, Goodhardt and Barwise [19], where small brands perform comparatively worse than large brands. The model suggests that the exchange of product recommendations, and perhaps word of mouth processes more generally, are an explanatory factor. The model shows that products enjoy increasing returns to appealing to a larger share of the consumer population, and this reinforces their market shares to the point that these overestimate the appeal of best-selling products and underestimate that of lesser performing products.

[12]Note that no product pool is worse off with word of mouth in terms of sales volume. Not all product pools benefit, however, and some that increase their sales volume with word of mouth may nonetheless exhibit a lower market share.

5 Personalization

Online retailers have implemented recommender systems to generate personalized product recommendations for consumers, and major players such as Amazon and Netflix have invested heavily in their development. These systems are one of the main novelties surrounding the online transition of word of mouth. Collaborative filtering algorithms underlie most commercial recommenders, identifying preference similarity among consumers in order to select which products to recommend (e.g. 'customers who bought this item also bought...'). Consider a simple instance of such an algorithm . The firm exploits a database containing a set C of consumers, a set N of products, and the ratings that consumers have provided for products. If consumer c_i has not rated product n_k, an expected value for that rating can be calculated by

$$E[Rating(c_i, n_k)] = \sum_{j \epsilon C} Similarity(c_i, c_j) * Rating(c_j, n_k), \tag{21}$$

where the *Similarity* function measures the taste proximity of any two consumers based on the correlation of their past product ratings. The algorithm will recommend to consumer c_i the unrated product which obtains a higher expected rating.[13]

I next analyze the impact of personalization on the market. Observe that a collaborative filter can be interpreted as a mechanism that matches consumers based on their preferences during the recommendation exchange. That is, it ensures that consumers always obtain recommendations from others that share their same product preferences. Therefore, if the firm introduces personalization, consumers seeking recommendations in the third stage will obtain them from others of their same type t (given that personalization outperforms traditional word of mouth, assuming the cost of obtaining a recommendation r remains constant)

Personalized recommendations always yield a match. The impact of introducing personalization in the market can be characterized by following the analysis in the previous section and taking into account that $\alpha^t = 1$ for all t, which homogenizes the value of recommendations across all consumer types. Inspection of $c_r^{t,i}$ (10) reveals that the indifferent recommendation seeker with personalization, denoted by c_{rp}^i, is given by

$$c_{rp}^i = u - p - r, \tag{22}$$

and is common across types. Inspection of $c_e^{t,i}$ (12) reveals that the indifferent evaluator with personalization c_{ep}^i is also common across types and given by

[13]See Adomavicius and Tuzhilin [1] for a taxonomy of recommender systems and an overview of the related computer science literature. For a brief discussion on the economics of recommender systems, see Resnick and Varian [34]. Murthi and Sarkar [32] review the general implications of personalization technologies in the context of the firm.

$$c^i_{ep} \;=\; \frac{\beta r}{1-\beta}. \tag{23}$$

Therefore participation is homogeneous across types $s^t_e = s^t_r = s^t_{er} = s^t$. The boundary price $\overline{p}^t$ also becomes homogeneous across types, and can be denoted simply by $\overline{p}$. The marginal type seeking recommendations $t_r(p)$ is now given by $t_r = 1$ if $p < \overline{p}$ or $t_r = T+1$ if $p \geq \overline{p}$. The firm's profit function π_r (14) then exhibits a unique non-differentiability at $\overline{p}$, so either all consumer types seek recommendations or none do. It follows from the assumption that $r < u/4$ that the firm's optimal price with personalization p_s always falls in the range $p_s < \overline{p}$ and recommendations continue to hold in equilibrium.[14]

I next consider the sales distribution with personalization and argue that the participation and mass market effects identified in the previous section are no longer present. Consider the participation shift given by $\overline{\rho} = \{\overline{\rho}_n\}^N_{n=1}$ (20) and the sales distribution with evaluations only $\sigma = \{\sigma_n\}^N_{n=1}$ (7). With personalization, given that participation is homogenous across types $s^t_{er} = s^t$, now $\overline{\rho} = \sigma$ and the participation shift is no longer present. Next, consider the mass market shift generated by recommendation seekers and given by $\mu^t = \{\mu^t_n\}^N_{n=1}$(19). With personalization, recommendations are drawn exclusively from evaluating consumers of the same type, so μ^t is now given by $\mu^t_n = \sigma^t_n\lambda^t_n/\alpha^t = \sigma^t_n$, and therefore $\mu = \sigma$ and the mass market shift is no longer present. I conclude that the sales distribution with personalization is equivalent to that derived in Proposition 1.

Proposition 3. *When there is personalization in the market, consumers of all types with sampling costs $c^i \in [0, c^i_{ep}]$ search with own product evaluations and consumers with sampling cost $c^i \in (c^i_{ep}, c^i_{rp}]$ search with recommendations, where $c^i_{ep} < c^{t,i}_e$ and $c^i_{rp} > c^{t,i}_r$. The firm prices products at $p_s = (u-r)/2$ and the sales distribution is equivalent to σ (7) derived in Proposition 1.*

This equilibrium compared to the one defined in Proposition 2 implies that personalization reduces the concentration of sales. All consumer types benefit from personalization and increase their participation, increasing sales volume and firm profits, but consumers with less prevalent preferences benefit the most.

Personalization increases the value of recommendations for consumers and ensures it no longer depends on how prevalent their preferences are in the population. Consumer participation as well as the share of consumers that seek recommendations increases across all types; some consumers who would otherwise stay of out of the market and some consumers who would otherwise search

[14] The maximum of the profit curve in the range $p \geq \overline{p}$ is given by p_e (5). For the solution to fall in the range $p < \overline{p}$ requires that $p_e < \overline{p}$, which implies $r < u(1-\beta)/2$. This always holds given assumption $r < u/4$. In addition, this equilibrium marks the highest consumer participation predicted in the model. An uncovered market then requires $c^i_{rp} < \overline{c}$, which given p_s implies $\overline{c} > (u-r)/2$. This lower boundary on $\overline{c}$ ensures the market is uncovered in all equilibria derived in the analysis.

with evaluations are now willing to seek recommendations instead. Both effects are stronger among consumers with less prevalent preferences (lower t), who benefit the most from personalization and are no longer in comparative disadvantage with respect to their peers.

The introduction of personalization expands the firm's demand in the lower price range. More consumers are now ready to participate by seeking recommendations at higher prices. The firm adjusts prices to account for the higher value of recommendations and these continue to hold in equilibrium. The firm may increase or reduce prices with the introduction of personalization, and the the sign of the adjustment depends on the value of word of mouth in the absence of personalization given by $\widehat{\alpha}^t$ (15) where $t = t_r(p_r)$. In general, the firm will lower prices when the introduction of personalization drives many consumer types to seek recommendations that would otherwise not do so. If personalization switches few consumer types to recommendations, the firm will increase prices.[15] Regardless of the direction of the price adjustment, firm profits increase and prices are always lower than in the equilibrium with evaluations only characterized in Proposition 1.

Personalization reduces the concentration of sales in the market, and this follows from the fact that it eliminates the exchange of recommendations among consumers with different product preferences. The mass market and participation effects driven by word of mouth are no longer present – mass market products no longer benefit comparatively more from recommendations and mainstream consumer types no longer participate comparatively more in the market. With respect to the sales distribution with word of mouth, this shifts market share from product pools with widespread appeal to those with narrow appeal, and the sales distribution with personalization matches that of evaluations derived in Proposition 1.

I have considered the most simple instance of personalization, that of a perfect taste matching mechanism. Although the precise impact on prices and across consumer types is difficult to pin down, three predictions stand out: personalization reduces the concentration of sales, increases consumer participation, and has a larger impact on consumers with niche preferences. Recent empirical work supports these predictions. Brynjolfsson, Hu and Simester [9] examine online and offline sales concentration for a clothing retailer and find that concentration is lower online, with the server log data showing that it is mainly due to the recommender system. Oestreicher-Singer and Sundararajan [33] examine sales concentration within book categories on Amazon.com and find that sales concentration is lower among categories with denser co-purchase networks, where personalization is expected to be more accurate. Using similar measures of concentration and co-purchase patterns, Ehrmann and Schmale [20] report the same findings on Amazon.de. De, Hu and Rahman [17] analyze server log data of an online clothing retailer and find that the recommender systems increases sales volume.

[15] Note that, given equilibrium prices p_r in Proposition 2, personalization increases prices whenever $\widehat{\alpha}^t > 1$ in its absence and lowers prices prices whenever $\widehat{\alpha}^t < 1$. Inspection of $\widehat{\alpha}^t$ (15) reveals that $\widehat{\alpha}^t \epsilon (\frac{1}{2}, \beta^{-1})$ and is decreasing in t, the more types seek recommendations the higher $\widehat{\alpha}^t$, given that $s^t > s^{t-1}$ and $\beta < \alpha^t < 1$ for types that seek recommendations.

The findings suggest additional considerations for the firm implementing personalization mechanisms. The introduction of these mechanisms provides a new and valuable search avenue to consumers, and the model predicts that the mass of consumers directly evaluating products across the assortment decreases as a result. The real-world performance of recommender systems, however, is reliant on consumer preference data, which implies that rewarding evaluating consumers for the input they provide becomes an important consideration in their implementation. Avery, Resnick and Zeckhauser [5] analyze several such reward schemes. From a mechanism design perspective, this search model contributes two additional insights. First, information on product matches, rather than on products that failed to yield a match, is more valuable for large assortments and should command a higher reward. Second, due to their lower prevalence in the population and the value generated from their input, product evaluations from consumers with less prevalent preferences should also command a higher reward.

6 Discussion

I next review the main assumptions underlying the preceding analysis. The review suggests the findings are qualitatively robust to changes on the demand side of the market, such as enriching consumer preferences and search strategies. Changes on the supply side of the market, such as non-uniform pricing or endogenizing assortment composition, present technical challenges and require additional modeling assumptions for the analysis to remain relevant. This suggests, again, that the analysis is most relevant to markets with extensive assortments and fairly uniform pricing schemes.

Mass market products. I have assumed consumers agree on some products but disagree on the remaining, and this is broadly consistent with the empirical evidence on consumer demand for media products. Goel, Broder, Gabrilovich and Pang [24] 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 [38] and Elberse [21] report similar findings. Consumer preferences could be enriched by considering partial overlapping in the product pools preferred by different consumer types, allowing some product pools to be preferred by a subset of types. This increases the complexity of the problem by increasing the interdependencies between types that arises with word of mouth, and would lead to a richer mass market effect dependent on the precise overlapping across types. Consumer preferences could be simplified by eliminating overlapping across consumer types, discarding mass market products. This simplifies the analysis by reducing the interdependencies between types with word of mouth, and ensures the mass market effect is no longer present. It is important to stress, however, that the qualitative results of the analysis do not hinge on the presence of mass market products – the participation effect is still present, and therefore the direction of the shifts in consumer participation and sales concentration in Propositions 2 and 3 continue to hold.

Utility levels. Consumers derive the same utility from all product matches, and this aspect of the model could 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. The empirical evidence is inconclusive on this point, however. Goel, Broder, Gabrilovich and Pang [24] analyze consumer product ratings for music and movies and find that niche products are rated higher than mass market products in music, but the trend is reversed for movies. Introducing heterogeneous match utilities alters the stopping rules of consumers during search. If different products yield different levels of utility, a consumer will be willing to resume search 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 (and this will depend on her idiosyncratic sampling cost). This increases the complexity of consumers' search strategies, but the general implications are clear: increasing the relative utility consumers obtain from a product pool will increase its market share in equilibrium.

Provision of recommendations. Consumers that search with evaluations are willing to freely supply product recommendations to others, and a large body of literature has documented several motivations for consumers to contribute to word of mouth processes. See Dellarocas [18] for a related discussion. In markets for media products, casual evidence suggests that recommendations are well provisioned, and consumers may enjoy the opportunity to discuss their preferred products with others. Although consumers derive no direct benefit (nor cost) from supplying recommendations in the model, they do benefit indirectly from lower prices – recall that prices with word of mouth and personalization (Propositions 2 and 3) are always lower than in their absence (Proposition 1). Nonetheless, and given that consumers seeking recommendations incur a sunk cost r on each draw, the model suggests that they are willing to reward those that provide them. Interpreting r as a transfer instead of a sunk cost ensures the model does not break down if consumers supplying recommendations incur a communication cost r. See Avery, Resnick and Zeckhauser [5] for an analysis of reward schemes for the optimal provision of recommendations. I have also assumed that recommendations enjoy no salience, as consumers do not place additional value on a match that results from a recommendation. Senecal and Nantel [36] report a series of experiments that suggest recommendations have an influential effect on consumers beyond awareness. Salient recommendations would increase the expected utility consumers derive from recommendations, increasing the share of consumers searching with recommendations and overall consumer participation.

Pricing. The monopolist quotes a single price and cannot price-discriminate across consumers or products. This simplifies the analysis and allows me to focus on demand-side effects. The assumption is restrictive, but then 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 major online retailers such as Amazon and Apple apply single price schemes across their digital content catalogs. If the monopolist could price-discriminate consumers based on their search strategies, it can be shown that she would charge a higher price to those searching with evaluations (price p_e (5) specifically) and would further discount prices for those seeking recommendations. The effect analyzed by Kuksov and Xie [29] of providing lower prices or unexpected frills to early customers in order to profit from later customers is not present, as the surplus of evaluating consumers does not impact the value of recommendations in the market. If the monopolist could price-discriminate across products, the complexity of the problem can increase substantially due to the informational content of prices. In a scenario where consumers perfectly observe all product prices, as in the preceding analysis, prices act as a signaling device and render search irrelevant.

Assortment composition. I have assumed that the composition of the assortment supplied by the monopolist is exogenous. The monopolist acts as an intermediary or gatekeeper that supplies all available product varieties on the market. If the monopolist were to choose the composition of the assortment to maximize profits, she would evaluate the marginal profitability of each stocked product by considering its impact on consumers' search costs and participation – the profit maximizing solution then entails reducing the assortment to only mass market products. The fact that smaller assortments may be more attractive for consumers has been explored in the literature. Iyengar and Lepper [26] and Chernev [13] report experiments suggesting that more choice is not always preferred by consumers. Kuksov and Villas-Boas [28] 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 my setup, and is ruled out by the corner case of mass market products which provide maximum utility to all consumers. Optimal assortment composition in the presence of consumer-to-consumer communication is a complex problem and is beyond the scope of this paper.

7 Conclusion

The exchange of product recommendations is valuable for consumers in markets characterized by large assortments of horizontally differentiated products such as those for music, cinema, literature, or video game entertainment. This contributes to explain the prevalence of word of mouth in those markets as well as the high concentration of sales they exhibit. Due to the mechanisms that drive the exchange of recommendations 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, and personalization mechanisms such as recommender systems are a prominent example of how firms can play an active role in the process.

Personalization mechanisms aid consumers in their product discovery process by reducing their search costs. This increases sales and benefits consumers with uncommon preferences the most. Hence firms supplying personalization generate value in the market, and firms with a large product assortment and customer base stand to profit the most. Superior personalization can sustain a competitive advantage if the firm supplying it can capture a share of the value it generates – that is, if it completes the transaction after helping consumers identify their preferred products. In the context of recommender systems and electronic commerce, the presence of switching costs and network effects suggests that firms can design strategies to achieve this. Recommender systems exhibit a learning curve to identify the preferences of new customers and benefit from large datasets of consumer activity to improve their accuracy. Consumers will receive less accurate recommendations when switching purchases across firms and, in general, when patronizing smaller firms. Both factors suggest a firm can benefit from rewarding consumers to join and grow its customer base, generating a lock-in effect to outperform competitors.

Personalization also reduces the concentration of sales and benefits niche products, and this drives other marketing considerations. By increasing the demand for products in the tail of the sales distribution, firms with low inventory costs stand to profit the most from personalization. These firms can increase the depth of their assortment beyond that of competitors, ensuring others cannot serve consumers demanding niche products in the tail. In the case of Netflix and Blockbuster, for example, a large share of the movies Netflix recommends to its customers are not available in Blockbuster stores. By generating personalized recommendations that help consumers navigate its assortment, Netflix is also increasing the value of stocking a deeper assortment than competitors. Thus the analysis provides a rationale for online retailers to pioneer the provision of personalization mechanisms in the marketplace, and suggests that their value-generating potential should not be underestimated.

References

[1] 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

[2] Ansari, Asim, Skander Essegaier and Rajeev Kohli (2000), 'Internet Recommendation Systems,' *Journal of Marketing Research* 27 363-375.

[3] Aral, Sinan and Dylan Walker (2011), 'Creating Social Contagion Through Viral Product Design: A Randomized Trial of Peer Influence in Networks,' *Management Science* (Forthcoming)

[4] Atahan, Pelin and Sumit Sarkar (2010), 'Accelerated Learning of User Profiles,' *Management Science* 57:2 215-239

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

[6] Bar-Isaac, Heski, Guillermo Caruana and Vicente Cuñat (2011), 'Search, Design, and Market Structure,' *American Economic Review* (Forthcoming)

[7] Berger, Jonah and Eric M. Schwartz (2011), 'What Drives Immediate and Ongoing Word of Mouth?,' *Journal of Marketing Research* 48 869-880

[8] Bodapati, Anand V. (2008), 'Recommendation Systems with Purchase Data,' *Journal of Marketing Research* 45 77-93

[9] Brynjolfsson, Erik, Yu Jeffrey Hu and Duncan Simester (2011), 'Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales,' *Management Science* (Forthcoming)

[10] Cheema, Amar and Andrew M. Kaikati (2010), 'The Effect of Need for Uniqueness on Word of Mouth,' *Journal of Marketing Research* 47 553-563

[11] Chen, Yan, F. Maxwell Harper, Joseph Konstan and Sherry Xin Li (2010), 'Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens,' *American Economic Review* 100 1358-1398

[12] Chen, Yubo, Qi Wang and Jinhong Xie (2011), 'Online Social Interactions: A Natural Experiment on Word of Mouth Versus Observational Learning,' *Journal of Marketing Research* 48 238-254

[13] 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

[14] 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

[15] Chintagunta, Pradeep K., Shyam Gopinath and Sriram Venkataraman (2010), 'The Effects of Online User Reviews on Movie Box Office Performance: Accounting for Sequential Rollout and Aggregation Across Local Markets,' *Marketing Science* 29:5 944-957

[16] Choi, Jeonghye and David R. Bell (2011), 'Preference Minorities and the Internet,' *Journal of Marketing Research* 48 670-682

[17] De, Prabuddha, Yu Jeffrey Hu, and Mohammad S. Rahman (2010), 'Technology Usage and Online Sales: An Empirical Study,' *Management Science* 56:11 1930-1945

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

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

[20] 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*

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

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

[23] 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

[24] 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*

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

[26] 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

[27] Kim, Jun B., Paulo Albuquerque and Bart J. Bronnenberg (2010), 'Online Demand Under Limited Consumer Search,' *Marketing Science* 29:6 1001-1023

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

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

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

[31] Manchanda, Puneet, Ying Xie, Nara Youn (2008), 'The Role of Targeted Communication and Contagion in Product Adoption,' *Marketing Science* 27:6

[32] Murthi, B. P. S. and Sumit Sarkar (2003), 'The Role of the Management Sciences in Research on Personalization,' *Management Science* 49:10 1344-1362

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

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

[35] 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

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

[37] Sun, Monic (2011), 'How Does Variance of Product Ratings Matter?,' *Management Science* (Forthcoming)

[38] 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*

[39] Tucker, Katherine and Juanjuan Zhang (2011), 'How Does Popularity Information Affect Choices? A Field Experiment,' *Management Science* 57:5 828-842

[40] Ying, Yuanping, Fred Feinberg and Michel Wedel (2006), 'Leveraging Missing Ratings to Improve Online Recommendation Systems,' *Journal of Marketing Research* 43 355-365

[41] 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