

Word of Mouth and Recommender Systems: A Theory of the Long Tail*

Andres Hervas-Drane[†]

Job Market Paper

November 12, 2007

Abstract

I present a model to assess the extent to which recommender systems can account for the ‘long tail’, an increase in the tail of the sales distribution. Consumers face a search problem within a pool of horizontally differentiated products supplied by a monopolist. They are endowed with a taste profile that determines their probability of matching with any given product, but arrive to the market uninformed and cannot identify which products are more likely to yield a match. Consumers may search for a match by drawing products from the assortment or by seeking word of mouth recommendations from other consumers. Product evaluations prior to purchase and the exchange of recommendations are both shown to arise endogenously, increasing firm profits and the concentration of sales. Introducing a recommender system to act as an intermediary in the recommendations exchange further increases firm profits and affects sales concentration. Insights are derived on the mechanisms driving concentration in artistic markets and their implications for the long tail debate. The model is suited for experience good markets such as music, cinema, literature and video game entertainment.

Keywords: Search, Product Recommendations, Sales Concentration, 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 also thank seminar participants at the Harvard Economics of Organizations Brown Bag Lunch. Financial support from the HBS Division of Research, the Real Colegio Complutense and the NET Institute are gratefully acknowledged.

[†]Harvard Business School and Universitat Autònoma de Barcelona (ahervas@fas.harvard.edu)

1 Introduction

The expansion and development of electronic commerce in recent years has brought radical change to the distribution landscape. Products previously limited to specialized stores are now only clicks away from delivery, offering consumers access to a larger variety of goods than ever before. This evolution has been most noticeable in product categories such as books, music and films, where assortment sizes have increased dramatically. For example, Amazon sells over 3 million book titles compared to the 100,000 stocked by an average Barnes & Noble store.¹ The digitalization of content paired with the advent of digital distribution is further fueling this trend. Observers and industry analysts 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 the traditional distribution channel. This phenomenon was coined by Anderson [3] as the *long tail*, referring to the increase in the tail of the sales distribution. As empirical studies turn their attention to the available data and the mechanisms driving these changes are discussed, the long tail has become an object of academic debate.

Increased availability of products is understood to be the explanatory factor for this phenomenon, given that more niche consumers can now access their preferred products through the online channel. Some of these transactions were previously excluded from the market due to the logistical constraints of traditional distribution, which limited the availability of products with a low market share. However, recent studies suggest that factors beyond availability seem to be driving down sales concentration. Brynjolfsson et al. [7] analyze the sales distribution of a clothing retailer offering the same product selection across two separate channels: catalog and online. Both channels offer equal prices and conditions. Considering consumers that purchase through both channels, they find that sales concentration is lower online. In another study, Elberse and Oberholzer-Gee [12] report decreasing sales concentration within a sample of video titles over a five year period. Their data source covers both online and offline retail channels. By controlling for the introduction of new titles in the market, they conclude the changes observed are driven by demand side effects and online retailing. Both studies suggest that online distribution is triggering changes in

¹See Brynjolfsson et al. [8].

consumption patterns, but the drivers of these changes are not well understood.

This paper presents a model that can rationalize these facts. Our approach is motivated by the impact that recommender systems, implemented by major online retailers, have on consumer choice. Our model explains how personalized recommendations, such as those generated by recommender systems, can lower the concentration of online sales. Our results stem from the improvement that these systems provide over word of mouth by means of personalized filtering. We define word of mouth as the direct exchange of product recommendations among consumers. We show that word of mouth recommendations benefit mostly *mainstream consumers*, those whose product preferences are more widespread in the population. This asymmetry is reduced when recommender systems are introduced as an intermediary in the exchange.

Our model explains how recommender systems reduce search costs in the market, thereby increasing consumer participation and firm profits. They achieve this by processing data on consumer preferences which can be retrieved, for example, from product purchase and browsing history, product ratings and consumer demographics. While the development of these technologies has been pioneered by online retailers, traditional retailers are increasingly implementing them. By better exploiting consumer information to improve the quality of product recommendations, firms can sustain a competitive advantage. In addition to this, the model suggests that recommender systems can increase product variety in the long term. In artistic markets, increased demand for products that appeal to *niche consumers*, those with a rare taste in the population, provide incentives for emerging artists to participate. These effects showcase how innovation fostered by electronic commerce can extend beyond online distribution channels.

To the best of our knowledge, no previous theoretical work has explored the link between product recommendations and sales concentration. We consider a market of horizontally differentiated products supplied by a monopolist at a common price. The monopolist may be an electronic retailer or content provider offering a large product assortment. Each consumer's preferences are simplified to a partition of the product space into preferred and non-preferred products, as determined by her *taste*. A consumer derives positive utility from the consumption of a product which belongs to her preferred set and zero utility otherwise. Consumers arrive to the market uninformed and cannot identify their preferred products. As all products are ex-ante identical, the value of each product

can only be determined by sampling it. A product *match* is achieved when a consumer locates a product which belongs to her preferred set. But sampling products is costly, as it requires time and attention, and thus consumers face a search problem driven by taste.

We let consumers in the population differ in their taste and sampling costs. Consumers search for a match by sampling products, and may either draw products randomly from the assortment or seek recommendations from others. Recommendations exchanged by word of mouth are drawn randomly from the population of consumers that located a match. We find that consumers choose to seek and follow recommendations because they increase their probability of locating a product match. Mainstream consumers benefit more from word of mouth than niche consumers because recommendations are more likely to originate from others that share their taste, thus enjoying a larger probability of locating a match. We introduce a simple recommender system that filters recommendations for consumers based on their taste. We show that such a system can equate the probability of locating a match for all consumers, yielding a larger improvement for niche consumers.

The construction is well suited for experience goods such as music, films, books 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 objective characteristics of a product, such as genre, characteristics or plot, personal judgment requires direct exposure. Furthermore, due to exogenous factors beyond those explored here, price dispersion across titles is generally low in these markets. Hence product preferences have a larger impact than price on the concentration of sales.

1.1 Literature

Little theoretical work has focused on the mechanisms driving sales concentration within product assortments. Product differentiation models, for example, cannot readily explain how changes in the distribution channel affect the composition of sales. The search literature has mainly focused on price dispersion, by considering homogeneous goods offered by different sellers. These models are suited for settings where price dominates the search, but provide no insights on sales concentration across heterogeneous products. Some instances have explored heterogeneous consumer preferences with location models, such as Bakos [6]. But in this case the equilibrium is symmetric for all

consumer types and sellers, and no sales concentration is predicted by the model.

Recent work related to the long tail debate has proposed several factors that may explain sales concentration. Brynjolfsson et al. [7] present a search model with advertising. Consumers arrive to the market informed about advertised products, but incur search costs to learn about the remaining products. Sales concentration depends on how the size of the advertised and non-advertised product pools compare. Product popularity information is analyzed in an experiment by Salganik et al. [17]. They study demand concentration over a set of rare songs offered to test subjects on the Internet, with some treatments including popularity feedback and others not. They find that popularity information increases both concentration and the unpredictability of popularity in the outcome. Tucker and Zhang [18] 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 when popularity information is provided. However, it is unclear to what extent advertising and product popularity information can explain lower sales concentration online. As these factors have been shown to increase concentration, this would require online shoppers to be less exposed to both.

More closely related to the mechanisms explored here, Fleder and Hosanagar [13] analyze the impact of recommender systems on sales concentration. In their analytical model, they consider consumer purchases that follow product recommendations given an exogenous probability. The recommender system follows a popularity rule, recommending the bestselling product to all consumers, and they show the process tends to increase the concentration of sales. As a result, the treatment is somewhat akin to providing product popularity information. It does not account for consumer preferences and their incentives to follow recommendations or not.²

Our approach is focused on the demand for recommendations and their impact on sales; we simplify the problem by assuming the provision of recommendations as given. A large body of literature

²Simulation results are presented where consumers and products are located on a 2-axis space. In this setting the recommender model is richer and consumer preferences are well defined, but the incentives to follow recommendations are still exogenous. In the base scenario consumers are assumed to derive increased utility when following recommendations. It is not clear why products with a worse fit become more valuable after being recommended; in absence of this effect, consumers would strictly prefer to ignore recommendations. An interesting extension considers the case where consumers are not fully informed about products. The recommender system may then help consumers discover relevant products. But as awareness is randomized, the treatment does not account for word of mouth or advertising, and consumers' awareness of potentially popular products is equivalent to that of rare alternatives.

has documented several motivations for consumers to contribute to word of mouth processes, see Dellarocas [10] for a related discussion. Avery et al. [5] explore reward mechanisms for the optimal provision of recommendations. In our model, we assume consumers providing recommendations derive no immediate benefit (nor cost) in the process. Although an opportunity for profit may exist from a bargaining perspective, and our model does suggest that consumers are willing to reward others for recommendations, we do not further explore this aspect of the problem. Casual evidence suggests that recommendations are well provisioned in the markets considered here. 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] and may well offset any bargaining opportunity.

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 [16] famous superstars model as well as later contributions, such as MacDonald [14]. 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 frictions 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. In section 3 we derive the equilibrium with word of mouth. We enrich the model by introducing taste heterogeneity in section 4. In section 5 we introduce a recommender system that acts as an intermediary in the exchange of recommendations and analyze its impact on the market. Section 6 concludes. The Appendix proves Herfindahl index properties used in the analysis and contains the proofs not included in the text.

2 The model

Consider a market where a monopolist supplies $N \geq 3$ horizontally differentiated products. The monopolist quotes a common price for all products, denoted by p , and incurs a transaction cost t per unit sold. The number of products, N , is determined by outside factors and considered exogenous throughout the analysis.³

In this market there is a continuum of consumers with measure one. A consumer may derive positive utility from a product or not. In the first case, the consumer *matches* with the product and derives utility u from its consumption. In the second case, the consumer derives zero utility from the product. Consumers exhibit unit demand; they may match with several products but will only consume one. Consumers participate in the market to consume one product they match with or remain out. Utility of the outside option is normalized to zero.

Consumers may only learn if they match with a product by sampling it. Sampling products is costly for consumers, and we let the cost differ across consumers in the population. Sampling costs are drawn from a uniform cost distribution, where the cost of consumer i is given by $c^i \sim U[0, \bar{c}]$. Thus sampling a product which does not result in a match incurs disutility c^i , and sampling and consuming a product match yields utility $u - c^i - p$.

The outcome of sampling a product is modeled as a Bernoulli trial, with success probability dependent on the consumer's taste for the specific product. We refer to this probability as *match probability*. A *taste profile* specifies the match probability over all products, fully characterizing the taste of any given consumer in the market. We start by considering the homogeneous case, where all consumers share a common taste profile λ defined over the N products, $\lambda = \{\lambda_1, \lambda_2, \dots, \lambda_N\}$ where $\lambda_n \in [0, 1]$. Consumers sharing a common taste profile are ex-ante identical, as they have an equal probability of matching with any given product. But given that matches are random outcomes, they will generally differ ex-post in their realized product matches.⁴

We will refer to products with positive match probability as *relevant products* and to products with zero match probability as *irrelevant products*. A special case for the analysis is that of *quasi-*

³The single price assumption simplifies consumer search strategies while ensuring prices are not signaling.

⁴Thus consumers with the same taste profile may end up purchasing different products. The random element captures different consumption settings or other idiosyncratic factors influencing product choices. In general, consumers with the same taste profile will exhibit correlated preferences over products but need not agree on each one.

niform taste. A taste profile is quasiuniform if the match probability is common across all relevant products, that is $\lambda_n \in \{0, a\}$ for all n and $a \in (0, 1]$.

It is useful for the analysis to decompose the taste profile as follows. We construct *taste distribution* τ , a discrete distribution with support over the N products, by

$$\tau_n = \frac{\lambda_n}{\sum \lambda_n}. \quad (1)$$

Where, unless otherwise stated in notation, all summations extend over the N products. The taste distribution captures the comparative market appeal of products by providing a measure of each product's match probability with respect to others. If the probability of a match across products is constant and all products have an equal match probability, τ is uniform. If only one relevant product exists, that is one product has positive match probability and $N - 1$ products never yield a match, τ is a degenerate distribution.

For our analysis, we assume that products within the assortment vary in their appeal for consumers and $\lambda_i \neq \lambda_j$ for at least some i and j . This implies that at least some product has a positive probability of resulting in a match and τ is a non-uniform distribution. In the quasiuniform taste case, τ is uniform over the support with positive probability (over relevant products).

The taste distribution also serves as a benchmark for the concentration of sales in the market. If the distribution of sales coincides with taste distribution τ , product sales are representative of the underlying preferences of consumers. To see this, consider the case in which all consumers are informed about products. In our framework, a unit purchase scenario, an informed consumer is a consumer that has located a match. If consumer matches in the population are distributed over products according to their comparative appeal, the proportion of consumers matched to product n is exactly τ_n . If all consumers purchase their match, the market share of product n is given by τ_n and the distribution of sales coincides with τ .

We define *taste intensity* β as

$$\beta = \frac{\sum \lambda_n}{N}. \quad (2)$$

Taste intensity is the average probability of a match across the N products. The higher the

taste intensity, the more probable it is for consumers to locate a match within the product space. Thus β measures the value of the assortment for consumers. The original taste profile λ is fully recoverable from β and τ .

Taste concentration is defined as the Herfindahl index of τ , denoted by H ,

$$H = \sum \tau_n^2. \tag{3}$$

The higher the taste concentration, the larger the variations in match probabilities across products. In the Appendix we show that the Herfindahl index can be expressed as a function of N and the sample variance of τ .

When arriving to the market consumers observe only the size of the assortment N , the price level set by the monopolist p and their taste intensity β and taste concentration H . Consumers cannot map their taste profile to the set of available products, so they cannot identify which products are more likely to yield a match. All products are ex-ante identical, and as a result consumers face a search problem in order to locate a match.⁵

The search process is modeled in three stages. In the first stage, the monopolist chooses the price level in the market, p . Note that by construction the monopolist cannot price discriminate consumers. In the second stage, consumers may search for a match by sequentially drawing and sampling products from the assortment. Consumers incur sampling cost c^i on each draw to learn if they match with the product. In the third stage, consumers may search by sequentially seeking recommendations from those that searched in the second stage. Each recommendation draw obtains a product reference from a consumer that searched and located a match in the second stage. The consumer providing the recommendation identifies the product she matched with. To draw a recommendation incurs fixed cost r , as information must be obtained or requested from others. After receiving a recommendation, the searching consumer may then draw and sample the product at cost c^i .⁶

⁵An alternative interpretation of the model is the following: nature moves first in the game and assigns each consumer's matches. Each consumer is matched with product n with probability λ_n . Under this interpretation, λ characterizes the proportion of the population matched with each product. Match realizations are not observed by consumers; each consumer arrives to the market ex-ante matched with a set of products but cannot identify them among the assortment.

⁶The recommendation exchange can be understood to take place either online or offline. In the first case, sampling consumers actively publish their recommendations and consumers seeking recommendations browse them. In the

Consumers form a rational expectation of their participation costs in the market. Participation costs have two components: the search costs to locate a match and the price to be paid for the desired product. Search costs include both sampling and recommendation costs. We define *search frictions* as the total search costs required for all consumers to locate a match in the market. In order to compute search frictions, we let each consumer choose the search strategy that minimizes her search costs. While this may include consumers that prefer not to participate in the market, it provides a consistent measure for the whole consumer population.

The search problem is solved assuming uniform sampling with replacement. Uniform draws from the product space are consistent with the fact that products are ex-ante identical to consumers. Recommendations are also drawn uniformly from the mass of consumers that searched in the second stage. This implies that recommendations are representative of the population's preferences. Sampling the product space with replacement simplifies the problem and ensures tractability. The approach approximates the model without replacement as long as the number of draws in equilibrium is small with respect to N and the match probabilities across products ensure most consumers will locate a match within the assortment.

We assume the cost of recommendations is low $r \in (0, \frac{\beta(NH-1)(u-t)}{2})$ and sampling costs in the consumer population are high $\bar{c} \geq \frac{\beta NH(u-t)-r}{2}$. These conditions simplify the analysis by avoiding corner solutions in the pricing game, and allow us to restrict attention to the interesting cases. The first condition bounds the cost of seeking recommendations and ensures they are exchanged in equilibrium. For higher values of r , recommendations play no role in the market. The second condition ensures the market is uncovered in equilibrium and consumers with high sampling costs prefer not to participate. So we need not consider corner pricing solutions where the market is covered. These restrictions are further discussed when analyzing the firm's strategy in the proof of Proposition 2. All games are solved by backwards induction.

second case, consumers seeking recommendations observe which consumers have already matched and request product references from them.

3 Search with word of mouth

We start by solving a simplified game without evaluations or recommendations. We define a product evaluation as the opportunity to sample a product before purchase. In the next subsection, we consider the two stage game where consumers cannot sample products prior to purchase nor search with recommendations. As product evaluations are not possible, the only feasible strategy to locate a match is to purchase products from the assortment in order to sample them. Starting in subsection 3.2 we introduce evaluations in the second stage. Consumers can now sample products prior to purchase. In subsection 3.3 we introduce recommendations and solve the full three stage game where consumers can seek product references from others. This approach simplifies the exposition and allows us to isolate the impact of evaluations and recommendations on the market.

3.1 No evaluations or recommendations

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 expected probability of a match on each purchase when consumers sample from the product space. As products are drawn uniformly with replacement, the probability of drawing any given product at any step in the search is equal to $1/N$. The probability of a match is then

$$\sum \frac{1}{N} \lambda_n = \sum \beta \tau_n = \beta. \quad (4)$$

Each purchase is a Bernoulli trial with success probability equal to taste intensity β . The expected utility of a new purchase for consumer i is

$$u_s^i = \beta u - c^i - p, \quad (5)$$

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 will also vary across consumers, as they differ in c^i . The utility of a successive draw, however, is constant throughout the search for any given consumer. Hence consumers either search until a match is obtained or don't participate in the

market. We can identify the consumer which is strictly indifferent between both alternatives by equating u_s^i to zero. Denote this indifferent consumer by c_s^i ,

$$c_s^i = \beta u - p. \quad (6)$$

Only consumers with a sampling cost $c^i \leq c_s^i$ choose to search. 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.

Given the consumer participation constraint (6), which is a function of price level p , we can now solve the first stage of the game. 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 β is low, $t \geq \beta u$, no feasible transaction is profitable and the market breaks down. So 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 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 over all products are

$$\pi_s = \frac{c_s^i}{\bar{c}} \beta^{-1} (p - t) = \frac{(u\beta - p)(p - t)}{\bar{c}\beta}. \quad (7)$$

Solving for the firm's optimal price we obtain

$$p_s = \frac{u\beta + t}{2}. \quad (8)$$

As shown by this first result, all products sell in positive volume when consumers are uninformed about products and no evaluations or recommendations are available. Consumers draw products uniformly from the assortment, so irrelevant products that appeal to no consumer also enjoy positive market shares due to unsuccessful purchases.

Lowering sampling costs in the population increases participation, as consumers anticipate the costs of locating a match and do not participate in the market if it does not pay off. When

transaction costs are high or taste intensity is low, the market breaks down. In these cases, no profitable price for the monopolist faces positive demand in the market due to consumers anticipating costly unsuccessful purchases. Similar findings were reported by Bakos [6] in a search model with horizontally differentiated products.

3.2 Evaluations prior to purchase

We next introduce consumer evaluations before purchase. Consider the consumer's problem in the second stage given a price level p . The probability of a match when 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. Note that consumers strictly prefer to search with evaluations by sampling products before purchase rather than after, as this avoids unsuccessful purchases. The expected utility will also vary across consumers, as they differ in 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 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. 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.

Next we characterize the sales distribution. As consumers only purchase when they locate a product match, the sales distribution must equal the distribution of matches over products. Denote this distribution by σ . Note that all consumers are identically and independently distributed in the sampling outcome, as every product evaluation is independent of past evaluations and those of other consumers. The sales distribution must also be independent of market participation. We can

derive σ by characterizing the distribution of matches over products for a single evaluation. The probability that a consumer matches product n is $(1/N)\lambda_n$ and the probability of a match over all products is given by β . This implies

$$\sigma = \frac{(1/N)\lambda_n}{\beta} = \frac{\beta\tau_n}{\beta} = \tau_n, \quad (11)$$

and the sales distribution coincides with the taste distribution. As a result, irrelevant products enjoy no sales in the market.

We can now turn to the firm's problem given the consumer participation constraint (10). Firm profits are

$$\pi_e = \frac{c_e^i}{\bar{c}}(p - t) = \frac{\beta(u - p)(p - t)}{\bar{c}}. \quad (12)$$

Solving for the firm's optimal price we obtain

$$p_e = \frac{u + t}{2}. \quad (13)$$

We next derive social welfare with evaluations SW_e , defined as the sum of consumer surplus and firm surplus. Every product sale generates social surplus u net of transaction cost t and sampling costs, and every consumer samples on average β^{-1} products to locate a match,

$$SW_e = \frac{c_e^i}{\bar{c}}(u - t) - \int_0^{c_e^i} \beta^{-1} c^i dc^i. \quad (14)$$

Social welfare without evaluations SW_s , differs by the fact that transaction costs are socially incurred each time a product is sampled. As every participating consumer samples β^{-1} products on average,

$$SW_s = \frac{c_s^i}{\bar{c}}(u - \beta^{-1}t) - \int_0^{c_s^i} \beta^{-1} c^i dc^i. \quad (15)$$

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 $\bar{c} \leq 4$.

Proposition 1 *Evaluations decrease search frictions and increase the concentration of sales. Firm profits increase due to higher participation and higher product prices, but the effect on consumer surplus depends on the population's sampling costs. Lowering sampling costs increases both firm profits and consumer surplus.*

Evaluations allow consumers to purchase only products they match with, ensuring the distribution of sales coincides with consumers' taste distribution over products. This increases the concentration of sales over the case where consumers cannot sample products before purchase, as irrelevant products no longer sell in the market. Although evaluations have generally been proposed to decrease sales concentration due to increased product exploration, our analysis suggests otherwise. The mechanism is simple: evaluations ensure consumers do not purchase products they do not match with. This shifts market share from irrelevant to relevant products, as consumers now only purchase the latter.

Every consumer participating in the market realizes a unique purchase once a match is located. This rotates the firm's demand curve contracting it over the range with positive demand, as less product sales per consumer are realized for any level of market participation. Evaluations also reduce search frictions as consumers need not incur product purchases in their search for a match. This increases market participation, rotating the demand curve and expanding it for all positive prices. In equilibrium, evaluations ensure the firm is better off sustaining higher prices and consumer participation increases. Although the overall volume of transactions is lower, firm profits are strictly higher with evaluations. Consumer surplus is only higher when sampling costs are low, $\bar{c} \leq 2$. When sampling costs are higher, consumers are better off in the absence of evaluations. The net effect on social welfare is only positive when sampling costs are below a certain threshold, $\bar{c} \leq 4$.

Evaluations may be costly for the firm if additional resources or infrastructure are required. When transaction costs are high or taste intensity is low, $t < \beta u$, evaluations can enable markets that would otherwise break down. In these cases, the firm has strong incentives to implement evaluations. Additional profits gained from evaluations also decrease quickly for larger values of taste intensity, $\beta \rightarrow 1$, as consumers incur few unsuccessful purchases without evaluations in the first place. Hence we should expect evaluations to be implemented for lower values of taste intensity. 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.

The firm also has incentives to lower sampling costs for consumers. Casual evidence suggests that firms invest in doing so. Many bookstores, for example, provide a comfortable environment with cafeteria services for their customers to browse books. Online retailers invest in the infrastructure required to directly stream book excerpts, song clips and movie trailers to their customers. This in turn provides incentives for more consumers to search the assortment and participate in the market.

3.3 Evaluations and recommendations

We next analyze the full three stage game with recommendations. Note that our previous results on the utility of an evaluation draw u_e^i (9) and the distribution of evaluating consumers σ (11) must carry over unaffected from our previous analysis.

We next characterize search with recommendations. 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. We start by assuming that a positive mass of evaluating consumers exists. The expected probability of a match for a consumer seeking recommendations, denoted by α , is given by

$$\alpha = \sum \sigma_n \lambda_n = \sum \tau_n (\beta N \tau_n) = \beta N \sum \tau_n^2, \quad (16)$$

where the sum of squares is equivalent to the Herfindahl index H ,

$$\alpha = \beta N H. \quad (17)$$

This expression provides a measure of the concentration of τ_n . Note that $\alpha > \beta$ as $H > 1/N$ (this result is shown in the Appendix). Hence the probability of a match is higher when seeking recommendations than when searching with evaluations.

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

$$u_r^i = \alpha(u - p) - r - c^i. \quad (18)$$

As every recommendation draw incurs cost r in addition to sampling cost 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 expected utility of searching with recommendations will also vary across consumers as they differ in their sampling cost c^i . The utility of a successive draw, however, is constant throughout the search for any given consumer. Hence we can identify the consumer which is strictly indifferent between seeking recommendations and not participating by equating u_r^i to zero. We denote the indifferent recommendation seeker by c_r^i ,

$$c_r^i = \alpha(u - p) - r. \quad (19)$$

Unmatched consumers with a sampling cost $c^i \leq c_r^i$ choose to search with recommendations in the third stage and consumers with an evaluation cost $c^i > c_r^i$ prefer to stay out of the market.

We next turn to the second stage of the game. Consumers decide to search with evaluations or not. 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 them as $\alpha > \beta$, consumers will 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 draws. 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 c_e^i we equate the expected utility derived from both search strategies in order to locate a match. The expected number of draws required for a match with evaluations and recommendations are given by β^{-1} and α^{-1} respectively. The indifferent evaluator is then

$$\begin{aligned} u - p - \frac{r + c_e^i}{\alpha} &= u - p - \frac{c_e^i}{\beta} \\ c_e^i &= \frac{\beta r}{\alpha - \beta}. \end{aligned} \quad (20)$$

Thus consumers with an evaluation cost $c^i \in [0, c_e^i)$ prefer to search with evaluations in the second stage over seeking recommendations.

We can now characterize consumer's search strategy. Two cases exist, $c_e^i \leq c_r^i$ and $c_e^i > c_r^i$. As c_r^i is a decreasing function of price level p , we can identify the boundary price \bar{p} that separates both solutions by equating $c_e^i = c_r^i$,

$$\bar{p} = u - \frac{r}{\alpha - \beta}.$$

If the firm's price is low, $p < \bar{p}$, then $c_e^i < c_r^i$ and some consumers seek recommendations. That is, consumers with sampling cost $c^i \in [0, c_e^i)$ search with evaluations, consumers such that $c^i \in [c_e^i, c_r^i)$ seek recommendations in the third stage and consumers with an evaluation cost $c^i > c_r^i$ choose to stay out of the market. If the firm's price is high, $p \geq \bar{p}$, consumers only search with evaluations. The market configuration in the high price range is characterized by our previous analysis where consumers only search with evaluations.

We next derive the sales distribution generated by consumers seeking recommendations, ρ . As every recommendation draw is independent from past draws, all consumers seeking recommendations are identically and independently distributed. Hence we need only characterize the distribution of matches for a single recommendation draw. The probability that a consumer matches product n is given by $\sigma_n \lambda_n$ and the probability of a match over all products is given by α . This implies

$$\rho_n = \frac{\sigma_n \lambda_n}{\alpha} = \frac{\tau_n (N \beta \tau_n)}{\beta N H} = \frac{\tau_n^2}{H}. \quad (21)$$

The sales distribution is skewed with respect to the taste distribution whenever consumers' taste is non-quasiuniform. This follows from the fact that, when relevant products differ in their match probability, those with higher match probability enjoy a proportionally larger share of sales when consumers search with recommendations instead of evaluations. There is a market share transfer from low to high appeal products. Technically, the shift is driven by the convexity of the square operator in (21).

We next turn to the first stage of the game and analyze the firm's pricing problem. We have

established that the firm's demand curve has two separate ranges; when $p < \bar{p}$ recommendations hold in the market, and when $p \geq \bar{p}$ they do not. The demand curve is composed of two linear components, it is continuous, (non-strictly) convex, and non-differentiable at \bar{p} . For the firm choosing the optimal price to maximize profits, the two separate ranges describe two concave profit curves, each of which lies above the other in its own range given the convexity of the demand curve and both intersect at the boundary price \bar{p} .

Consider the firm's solution in price range $p < \bar{p}$. As consumer participation is determined by c_r^i (19), firm profits are

$$\pi_r = \frac{c_r^i}{\bar{c}}(p - t) = \frac{\alpha(u - p) - r}{\bar{c}}(p - t). \quad (22)$$

Solving for the price level in equilibrium yields

$$p_r = \frac{u + t}{2} - \frac{r}{2\alpha}. \quad (23)$$

Our restriction on r ensures this solution is the global maximum over the full demand curve.⁷ Hence recommendations are exchanged in equilibrium, thereby increasing consumer participation and inducing the monopolist to reduce prices, which in turn implies that both firm profits and social welfare must increase.

Proposition 2 *Recommendations decrease search frictions and, whenever consumers' taste is non-quasiuniform, increase the concentration of sales. Firm profits and consumer surplus increase due to higher participation and lower product prices. Lowering the cost of recommendations intensifies the previous effects.*

Recommendations allow consumers to benefit from those that searched before them. The prob-

⁷For p_r to be the profit maximizing price, two conditions must be met. First, p_r must be well defined, $p_r < \bar{p}$. This condition reduces to $r < (u - t)(\alpha - \beta)\alpha/\beta + \alpha$ (I). Second, maximum profits in the price range $p \geq \bar{p}$ must be lower than those obtained with p_r . The maximum on the profit curve with evaluations is given by p_e in (13). A sufficient condition given the properties of both profit curves is $p^e < \bar{p}$, which reduces to $r < \frac{1}{2}(u - t)(\alpha - \beta)$ (II). As the bound imposed by (II) on r is lower than that of (I), condition (II) is sufficient for p_r to be the firm's solution. Substituting α in (II) as a function of H obtains $r \leq \beta(NH - 1)(u - t)/2$.

This equilibrium also marks the highest consumer participation level predicted in the model. For the market to be uncovered in equilibrium, $c_r^i < \bar{c}$ must hold, which given p_r reduces to $\bar{c} \geq (\beta NH(u - t) - r)/2$. This lower boundary on \bar{c} ensures the market is uncovered in all equilibria derived in our analysis.

ability of a match is higher when following recommendations, as information is gathered about which products to sample. The degree of this effect increases with the concentration of taste. In the case of minimum concentration, where all products were equally likely to yield a match, the match probability when following a recommendation would be no higher than randomly drawing a product from the assortment. When only one relevant product is present in the assortment, taste concentration is high and recommendations are highly valuable. In this case, the number of product draws required to locate a match searching with evaluations grows with the number of irrelevant products. But only one recommendation draw is required to locate the relevant product, as any consumer searching with evaluations must have previously matched with this product. This suggests that recommendations are particularly valuable for consumers in navigating large assortments with heterogeneous appeal.

Whenever consumers exhibit non-quasiuniform taste and relevant products differ in their match probability, recommendations increase the concentration of sales in the market. The effect can be explained as follows; products with a large appeal enjoy both a higher share of recommendations and a higher probability of resulting in a match than those with low appeal. Recommendations shift market share from low to high appeal products as consumers purchase proportionally more of the latter. The aggregate sales distribution is skewed *with respect* to the taste distribution. And the higher the proportion of consumers searching with recommendations, the higher the concentration of sales. Thus sales overestimate the appeal of bestsellers and underestimate that of products in the tail of the sales distribution.⁸

Recommendations also reduce search frictions, as consumers that choose to search with recommendations incur lower search costs to locate a match than otherwise. These consumers enjoy a higher match probability on each product draw at additional cost r , a trade-off they benefit from due to their high sampling costs. Consumers with low sampling costs prefer to search with evaluations as failed product draws are less costly for them. All consumers are strategic when choosing

⁸The result is robust. If we considered a dynamic model of consumer arrival where recommendations originated from all earlier consumers (not only those that searched with evaluations), the concentration of sales would increase. In this case, recommendations originating from consumers that previously matched seeking recommendations themselves would increase the skew of the recommendation source. This effect would increase both the match probability with recommendations and the concentration of sales they generate, ensuring a growing proportion of new consumers search with recommendations. Such a dynamic model could approximate the findings on popularity feedback reported by Salganik et al. [17] and Tucker and Zhang [18].

their search strategy; many that would otherwise search with evaluations now choose to seek recommendations. This follows from the fact that the indifferent evaluator c_e^i has lower sampling costs in equilibrium with recommendations than without. The proportion of participating consumers that seek recommendations increases with the concentration of taste H and consumption utility u , and decreases with taste intensity β , the cost of recommendations r and price level p . As long as some consumers search with recommendations, marginal changes in price do not affect consumers searching with evaluations.

As a consequence of reduced search frictions, recommendations also increase consumer participation in the market. This effect rotates the demand curve expanding demand in the lower price range, where recommendations are exchanged in the market. Given the upper threshold we impose on the cost of recommendations r , recommendations are exchanged in equilibrium and the monopolist sustains lower prices as a result. This increases consumer participation, and implies that consumers searching with evaluations also enjoy reduced prices. As a result, recommendations increase both consumer surplus and firm profits, unambiguously increasing social welfare.

Just as lowering sampling costs for consumers, improving the quality or lowering the cost or recommendations has the potential to expand markets. This provides incentives for the firm to increase the value of recommendations and facilitate their provision. The advent of electronic commerce is increasingly allowing firms to become active players in this process. By offering a platform for recommendations, online retailers such as Amazon have become valuable resources for consumers. Chevalier and Mayzlin [9] analyze the impact of online book reviews at two major online retailers. They find that most reviews are overwhelmingly positive and increase the relative sales at the retailer they are posted on. This is consistent with our model, as consumers recommend the products they have matched with and these recommendations generate additional sales. These findings suggest that part of the market growth generated by electronic commerce in these product categories may be attributable to these effects alone.

4 Taste heterogeneity

In artistic markets, consumers differ in their product taste. As a result, recommendations are frequently exchanged between consumers with different taste. Due to these interactions, consumers are generally aware of how prevalent their taste is in the population and how similar it is to that of others. In most cases, however, consumers cannot discriminate the taste of others they interact with. They may lack the necessary information, or their social networks may be connected based on other factors. But consumers can generally anticipate how valuable recommendations obtained from others in the population will be. As interactions between consumers with different taste are relevant for word of mouth processes, we next introduce taste heterogeneity in the model to analyze their impact.

We consider the case of two consumer types that differ in their taste profile. To impose structure on the problem, we take the view that the most significant difference in taste across consumers is the selection of products they prefer. Thus we restrict our analysis to the case in which both taste profiles differ by a permutation of their elements. This ensures that consumers differ only in their match probability across products, but not in their taste intensity or taste concentration. While the analysis could be extended to unrestricted taste profiles or additional consumer types, at the cost of added complexity, we model the simplest case for taste heterogeneity.

Let consumers of type A and B have taste profiles λ^a and λ^b respectively. Both taste profiles differ by a permutation of their elements such that $\lambda_n^a \neq \lambda_n^b$ for at least some product n . We refer to any product n such that $\lambda_n^a > \lambda_n^b$ as products *preferred* by type A , conversely $\lambda_n^a < \lambda_n^b$ are products preferred by type B . *Jointly relevant products* are those that may yield a match for both consumer types, $\lambda_n^a > 0$ and $\lambda_n^b > 0$. Denote the taste distributions of both types by τ^a and τ^b . The construction ensures consumer types differ in their taste distribution but not in their taste intensity β and taste concentration H , which is common for both types.

It is useful for the analysis to define *taste similarity* as the sum of the pairwise products of both distributions, which we denote as the Joint Herfindahl index JH ,

$$JH = \sum \tau_n^a \tau_n^b. \tag{24}$$

This index measures the joint concentration of both taste distributions across products. In the Appendix we show that it can be decomposed as a function of N and the covariance of both distributions.

Denote by s^a and s^b the share of consumers of each type in the population. The analysis is of interest when consumer types differ in their prevalence. Without loss of generality we let consumers of type A be more prevalent, $s^a > s^b$. Consumers of type A are *mainstream consumers*, as their taste is more widespread in the population. Consumers of type B are *niche consumers*. We assume search costs are distributed independently of taste, so \bar{c} and r are common across both consumer types.

On arriving to the market consumers observe their type, the share of types in the population and the taste similarity across types. Consumers cannot discriminate the taste of others when seeking a recommendation; they cannot observe the type of a consumer that provides a recommendation. The timing of the game carries over from our previous analysis. To ensure that recommendations play a role in the market we now require $r \in (0, \frac{\beta(H+JH)N-2)(u-t)}{4})$. We can then prove the following Proposition.

Proposition 3 *Mainstream consumers benefit from a higher match probability than niche consumers when seeking recommendations, incurring lower search frictions and exhibiting higher market participation.*

Proof. *Included in the Appendix.* ■

Mainstream consumers derive a larger benefit from word of mouth in the market. The intuition is the following: most recommendations originate from mainstream consumers, and as a result are more likely to yield a match for mainstream consumers than niche consumers. For this reason, mainstream consumers incur lower search frictions than niche consumers. A larger proportion of mainstream consumers participate in the market, and a larger proportion of mainstream consumers search with recommendations. The degree of this asymmetry increases with the share of mainstream consumers in the population and decreases with the taste similarity between both types.

In equilibrium, the indifferent evaluator with mainstream taste always has a lower sampling cost than the one with niche taste, $c_e^{a,i} < c_e^{b,i}$. For recommendations to be more valuable for main-

stream consumers, however, a larger share of all consumers that search with evaluations must have mainstream taste, $s_e^a > s_e^b$. The model predicts that in equilibrium both conditions are satisfied, and this result arises endogenously from the search strategies chosen by consumers. Intuitively, a marginal increase in the mass of niche consumers searching with evaluations lowers the value of recommendations for mainstream consumers. As a result, any such marginal increase is met with a larger increase in the mass of mainstream evaluators, as there is a larger mass of mainstream consumers within any given range of sampling costs.

As recommendations are more valuable for mainstream consumers in equilibrium, $\alpha^a > \alpha^b$ and mainstream consumers participate proportionally more in the market. That is, consumers with high sampling costs and mainstream taste choose to participate that would not do so if they had niche taste. In equilibrium, not only a higher proportion of mainstream consumers participate in the market, but a higher proportion of those that participate search with recommendations. The asymmetry between types is driven by difference in the value of recommendations, $\alpha^a - \alpha^b$, which in turn increases with the prevalence of mainstream taste in the population and decreases with taste similarity. When the prevalence of mainstream taste is high and taste similarity is low, most product recommendations in the market originate from mainstream consumers and are unlikely to yield a match for niche consumers. In these cases, niche consumers will search only with evaluations. Participation thresholds in equilibrium satisfy $c_r^{a,i} > c_r^{b,i}$ if niche consumers search with recommendations, and $c_r^{a,i} > c_e^{b,i}$ if they only search with evaluations.⁹

Recommendations are still informative signals for niche consumers when they are unlikely to yield a match. They identify products that may not be evaluated. Niche consumers could seek recommendations in order to discard the recommended products from the set of products to evaluate, reducing the expected number of draws required for a match. We show, however, that the informativeness of recommendations quickly decreases with N in this case. Given that recommendations are costly, this search strategy is unlikely to pay off with large assortments. As an example, consider the case where $N = 2$, $\lambda_1 = 1$ and $\lambda_2 = 0$, $\beta = 1/2$. Informativeness can be expressed as the improvement gained in the match probability when sampling with and without a recommenda-

⁹Niche consumers provide recommendations even if no other consumers with niche taste search with recommendations. This increases the probability that other consumers match with the products they have matched with.

tion. The informativeness of following a good recommendation (product 1) is equal to $1 - \beta = 1/2$. The informativeness of discarding the product identified by a bad recommendation (product 2) is equivalent, as $\beta_{-2} - \beta = 1 - \beta = 1/2$. When only two alternatives exist, both signals are equally informative. Next consider the case where N is large, $\lambda_1 = 1$ and $\lambda_2 = \lambda_3 = \dots = \lambda_N = 0$. Now $\beta \approx 0$ and a good recommendation is highly informative, $1 - \beta \approx 1$. A bad recommendation is not, however, as discarding an irrelevant product helps little in identifying the relevant one, $\beta_{-j} - \beta \approx 0$. Presumably for this reason, we do not observe consumers seeking recommendations on what to dislike within large assortments.

Our model shows that products that appeal to a large audience suffer comparatively less search frictions than those that appeal to a small one. This result is reminiscent of the double jeopardy effect discussed by Ehrenberg et al. [11]. They observe that small brands perform comparatively worse than large brands on metrics such as consumer loyalty or purchase frequency, a pattern originally identified for comic strips and radio presenters. Our model suggests that word of mouth could be an explanatory factor for such effects.

5 Recommender systems

Recommender systems generate product recommendations by exploiting databases of product information and consumer taste. Two basic techniques have been developed to generate recommendations for a given consumer: content-based and collaborative filtering. The content-based technique selects which products to recommend based on product similarity, comparing the characteristics of available products and those that were preferred by the consumer in the past. The collaborative filtering technique selects products that were preferred by other consumers with a similar taste. Hybrid techniques attempt to combine the two. The algorithms driving recommender systems, from here on also referred to as recommenders, are an active area of research. Most specialists believe that scope for improvement exists, and commercially deployed algorithms continue to improve their performance. For a brief discussion on the economics of recommender systems, see Resnick and Varian [15]. A taxonomy of recommender systems and an overview of the computer science literature are presented by Adomavicius and Tuzhilin [2].

Here we consider a simple collaborative filter that is suitable for our framework. The recommender takes input from consumers evaluating products and generates output for consumers seeking recommendations. Each time a consumer requests a product recommendation, the recommender selects the product to recommend by estimating the consumer's taste. We model this process by considering a quality level q that characterizes the performance of the recommender. With probability q on each request, the recommender identifies the consumer's type. In this case the recommender provides a randomly picked recommendation from an evaluating consumer of the same type. If no input is available from such consumers or if the recommender cannot identify the consumer's type, a recommendation is drawn from the whole population of evaluating consumers. Thus our recommender acts as a taste profile filter for consumers, with filtering success q as long as the required input is available. We assume that consumers observe the quality of the recommender q , and seeking a recommendation from the recommender incurs cost r as in our previous analysis.

This approach to modeling a recommender system has several good properties. There is no free lunch, as the recommender learns from consumers searching with evaluations. Hence the capability of the recommender to generate personalized recommendations is constrained by the presence of evaluating consumers with the required taste. Furthermore, for $q < 1$ the recommender performs worse for the consumer type with a lower share of evaluating consumers. This captures the impact of information sparsity; the less information is available for a given consumer type, the lower the comparative performance of the system for these consumers. The bounds on performance imply that a recommender of minimum quality, $q = 0$, does not improve upon word of mouth. A recommender that performed worse would be ignored by consumers (at the same cost r). A maximum quality recommender with $q = 1$ always identifies the consumer's type. As long as the required information is available to the system, it positions the consumer in a world where everyone shares her taste. This seems a plausible upper bound on the performance of a collaborative filter.

We next analyze the impact of the recommender on firm profits and consumer participation. The next Proposition summarizes our findings.

Proposition 4 *Firm profits, consumer participation and the surplus enjoyed by consumers searching with recommendations are all strictly increasing in the quality of the recommender system. The latter effects are more pronounced for niche consumers.*

***Proof.** Included in the Appendix. ■*

The recommender system improves upon word of mouth by increasing the match probability of recommendations for all consumers. The improvement is achieved by ensuring that more product recommendations match the taste of the consumers that receive them. Increasing the match probability of recommendations reduces the search frictions in the market, and more consumers decide to participate and search with recommendations or change their search strategy to do so. Higher consumer participation implies that firm profits are strictly increasing in the quality of the recommender.

The recommender rotates and expands the firm's demand curve within the price range where consumers search with recommendations. Although in equilibrium the firm increases prices as a result, this increase does not offset the higher surplus enjoyed by consumers searching with recommendations. Consumers that search with evaluations, on the other hand, do not benefit from the recommender and are worse off after the price change. The overall effect on consumer surplus is ambiguous.

The impact of the recommender system varies across consumer types. When the share of mainstream consumers in the population is high and taste similarity low, a positive quality threshold may exist for the recommender to have an impact on niche consumers. Below this threshold, niche consumers search only with evaluations and do not use the recommender. Once this quality threshold is met and both consumer types use the recommender, however, any quality improvement has a larger impact on niche consumers. A marginal increase in q results in a larger increase in the participation and the match probability of recommendations for niche consumers. A perfect recommender, $q = 1$, ensures recommendations are equally valuable for both types and the same proportion of mainstream and niche consumers participate in the market.

Recommender systems have been proposed to suffer from incentive problems if less consumers choose to evaluate products when they are available. See for instance Resnick and Varian [15]. In equilibrium, the recommender learns from both consumer types as a positive mass of evaluating consumers always exists. The model predicts, however, that the proportion of consumers searching with evaluations decreases with the quality of the recommender. This aspect merits further attention as recommenders continue to improve and become more widespread. Rewarding evaluating

consumers for the information they provide may become an important strategic consideration for recommender systems. Avery et al. [5] explore reward schemes applicable to this scenario. Our model suggests that information on matches, rather than on products that did not yield a match, is more valuable to the system. Similarly, reward schemes may also benefit by accounting for the taste prevalence of consumers providing feedback.

Our results also suggest that niche consumers have a higher willingness to pay for recommender systems. Consumer feedback seems to suggest that recommenders are indeed more valuable for niche audiences, with frequent reports of product discoveries that have a narrow appeal in the market. This prediction of the theory could be empirically tested and serve as an avenue for further research. The emergence of business models based primarily on recommender systems calls for a better understanding of their target audience and potential revenue streams.

5.1 Sales concentration

We now turn to the impact of the recommender system on sales concentration. We proceed by comparing the concentration of the sales distribution in the market before and after the recommender is introduced. For simplicity we consider only the benchmark case of a perfect recommender, $q = 1$. We present analytical results for the family of quasiuniform taste profiles, and discuss other taste profiles below. The result requires the following property on the sales concentration measure: when products are ranked in decreasing market share order, a market share transfer from a low rank product to a higher rank product that preserves the rank of the first must lower sales concentration. This is a standard concentration index property, satisfied for example by the Gini index. The following Proposition states the result.

Proposition 5 *When consumers exhibit quasiuniform taste, the recommender system strictly decreases the concentration of the sales distribution in the market.*

Proof. *Included in the Appendix.* ■

Two separate effects drive the result. The first effect is a change in the distribution of sales generated by consumers when the recommender is introduced, the *sales distribution effect*. As the recommender filters recommendations by consumer type, jointly relevant products no longer

benefit of recommendations from evaluating consumers of both types. This effect shifts market share from jointly relevant products to all other relevant products. Second, the recommender has a larger impact on niche consumers, increasing the share of participating consumers with niche taste in the market. This *participation effect* shifts market share from products relevant only to mainstream consumers to those relevant only to niche consumers. The market share of jointly relevant products is unaffected given that they are equally demanded by both types (as taste profiles are quasiuniform). Both effects reduce the concentration of sales in the market.

The reduction in sales concentration achieved by the recommender varies with taste similarity. When taste is quasiuniform, taste similarity is determined by the proportion of relevant products that are jointly relevant. The intensity of the sales distribution effect is concave with respect to taste similarity, and peaks at an intermediary value (when approximately half of all relevant products are jointly relevant). The intensity of the participation effect is decreasing in taste similarity beyond a threshold value. This threshold value determines when niche consumers search with recommendations before the recommender is introduced in the market. In this range, the recommender has lower impact on the participation of niche consumers in the market. Maximum taste similarity would imply that all relevant products are also jointly relevant, rendering the problem equivalent to the basic model with homogeneous taste. In this case, the recommender has no impact on the market. We conclude that the reduction in sales concentration is generally higher when taste similarity is low. That is, when most products that appeal to mainstream consumers do not appeal to niche consumers and conversely, as niche consumers are comparatively worse off without the recommender.¹⁰

When consumers do not exhibit quasiuniform taste, the sales distribution and participation effects increase complexity and may change sign. On the one hand, the participation effect varies if jointly relevant products differ in their match probability for both types. On the other hand, the sales distribution effect may increase concentration, and this is illustrated in the following example. Consider the case where relevant products for each consumer type have heterogeneous match probabilities and no jointly relevant products exist. The latter assumption simplifies the exposition by

¹⁰Most commercial recommender systems have a popularity discounting rule to avoid recommending popular products (based on the assumption that most consumers are informed about them). Such a rule would further intensify the sales distribution effect when taste similarity is positive.

ensuring that the sales distribution generated by each consumer type is independent, and equivalent to the homogeneous taste case when considered in isolation. We established in Proposition 2 that recommendations generate a market share shift from low to high match probability products when taste is non-quasiuniform. We also established in Proposition 4 that introducing a recommender system increases the proportion of consumers of both types that search with recommendations. Due to both effects, the recommender must increase the concentration of the sales distribution generated by both consumer types. In this example, the sales distribution effect increases sales concentration in the market.

When taste is non-quasiuniform, simulations confirm that the recommender may increase or decrease sales concentration in the market. As suggested by the example above, taste concentration pushes the recommender to increase concentration. This follows from recommendations increasing the concentration of sales *within* the set of relevant products for each consumer type. Hence the recommender can contribute to create bestsellers for niche audiences just as traditional word of mouth favors mainstream bestsellers.

6 Concluding remarks

We have provided a framework for understanding the impact of product recommendations on the market. Recommendations create value by reducing the search costs incurred by consumers. We have analyzed a simple recommender system based on a collaborative filter and found that it can improve the quality of recommendations over word of mouth, thereby increasing consumer participation and firm profits. Thus our model can explain the adoption and development of recommender systems in recent years. If firms offering a high quality recommender system capture a share of the value they generate, this technology can sustain a competitive advantage in an industry. Our results suggest the incentives of the firm and those of consumers are aligned with respect to recommender systems, and may outweigh strategic opportunities for manipulation by these parties. Accounting for consumer trust and competition between recommender systems would only intensify this result.

Recommender systems have a larger impact on niche consumers, reducing asymmetries in the

market that originate from word of mouth. They reduce sales concentration by increasing the participation of niche consumers in the market and filtering recommendation exchanges between consumers with different taste. When niche consumers strongly favor some of their preferred products, recommender systems can fuel new bestsellers for niche audiences. While this effect can increase sales concentration, the limited empirical evidence available so far suggests it does not prevail. Our results call for further evidence on the sales impact of recommender systems.

While the long tail debate has focused on the concentration of sales, our results show that recommender systems unambiguously increase the sales volume of products that appeal to niche consumers. Hence recommender systems can increase product variety in the long term by driving demand for products in the tail of the sales distribution. This effect has implications for artistic markets, increasing the incentives to participate for artists that appeal to niche audiences. Lower sales concentration may only be one of the shorter term implications of personalized recommendations in markets.

7 Appendix

7.1 Herfindahl index decomposition

The Herfindahl index can be decomposed as follows,

$$\begin{aligned}
 H &= \sum \tau_n^2 \\
 &= \sum \left[\frac{1}{N} + \left(\tau_n - \frac{1}{N} \right) \right]^2 \\
 &= \sum \frac{1}{N^2} + \sum \frac{2}{N} \left(\tau_n - \frac{1}{N} \right) + \sum \left(\tau_n - \frac{1}{N} \right)^2.
 \end{aligned} \tag{25}$$

Where the first summation adds up to $1/N$ and the second up to zero. The third summation equals N times the sample variance of τ , which is equal to

$$Var(\tau) = \frac{1}{N} \sum \left(\tau_n - \frac{1}{N} \right)^2. \tag{26}$$

Hence

$$H = \frac{1}{N} + N \text{Var}(\tau). \quad (27)$$

Inspection reveals that the lower bound on the variance follows from $\tau_n = \frac{1}{N}$ for all n , that is, $\lambda_n = \lambda_j$ for all n and j . In this case $\text{Var}(\tau) = 0$. However, as our assumptions on λ require that $\lambda_n \neq \lambda_j$ for at least some n and j , $\text{Var}(\tau) > 0$. An upper bound on the variance requires the differences between τ_n and $\frac{1}{N}$ to be maximized, as this increases the value of the summation. Consider the case where $\lambda_j > 0$ for some j and $\lambda_n = 0$ for all $n \neq j$. Then τ is a degenerate distribution; $\tau_j = 1$ and $\tau_n = 0$ for $n \neq j$. In this case $\text{Var}(\tau) = \frac{1}{N}[(1 - \frac{1}{N})^2 + (N-1)(\frac{1}{N})^2] = \frac{N-1}{N^2}$. It can be shown that letting λ be positive for more products only increases $\text{Var}(\tau)$. Given these bounds, $\text{Var}(\tau) \in (0, \frac{N-1}{N^2}]$ and we conclude that $H \in (\frac{1}{N}, 1]$.

We next turn to the decomposition of the Joint Herfindahl,

$$\begin{aligned} JH &= \sum \tau_n^a \tau_n^b \\ &= \sum \left(\frac{1}{N} + (\tau_n^a - \frac{1}{N}) \right) \left(\frac{1}{N} + (\tau_n^b - \frac{1}{N}) \right) \\ &= \sum \frac{1}{N^2} + \sum \frac{1}{N} (\tau_n^a - \frac{1}{N}) + \sum \frac{1}{N} (\tau_n^b - \frac{1}{N}) + \sum (\tau_n^a - \frac{1}{N}) (\tau_n^b - \frac{1}{N}). \end{aligned} \quad (28)$$

Where the first summation adds up to $1/N$, and the second and third up to zero. The fourth summation equals N times the sample covariance, which is given by

$$\text{Cov}(\tau^a, \tau^b) = \frac{1}{N} \sum (\tau_n^a - \frac{1}{N}) (\tau_n^b - \frac{1}{N}). \quad (29)$$

Hence

$$JH = \frac{1}{N} + N \text{Cov}(\tau^a, \tau^b). \quad (30)$$

Inspection reveals that a lower bound on the covariance requires that both τ_n^a and τ_n^b differ from $\frac{1}{N}$ in opposite directions for some n , as this ensures that some terms of the summation are negative. Consider the case where $\lambda_j^a > 0$ for some j and $\lambda_n^a = 0$ elsewhere, and conversely $\lambda_k^b > 0$ for some $k \neq j$ and $\lambda_n^b = 0$ elsewhere. Then $\tau_j^a = 1$, $\tau_{n \neq j}^a = 0$, $\tau_k^b = 1$ and $\tau_{n \neq k}^b = 0$.

Notice that $\tau_n^a > \frac{1}{N} > \tau_n^b$ for $n = j$ and $\tau_n^b > \frac{1}{N} > \tau_n^a$ for $n = k$. In this case $Cov(\tau^a, \tau^b) = \frac{1}{N}(2(1 - \frac{1}{N})(0 - \frac{1}{N}) + (0 - \frac{1}{N})^2(N - 2)) = -\frac{1}{N^2}$. It can be shown that letting λ^a and λ^b differ for more elements only increases the covariance.

An upper bound on the covariance requires that τ_n^a and τ_n^b differ from $\frac{1}{N}$ in the same direction for some n , as this increases the value of some terms in the summation. It then follows that the covariance is maximized when both distributions coincide. In this case, $Cov(\tau^a, \tau^b) = Var(\tau^a) = Var(\tau^b)$. However, as our assumptions on λ require that $\lambda_n \neq \lambda_j$ for at least some n and j , both distributions cannot coincide, $\tau^a \neq \tau^b$, and $Cov(\tau^a, \tau^b) < Var(\tau^a) = Var(\tau^b)$. So an upper bound for the covariance is given by the upper bound on the variance, $\frac{N-1}{N^2}$. Thus $Cov(\tau^a, \tau^b) \in [-\frac{1}{N^2}, \frac{N-1}{N^2}]$ and $Cov(\tau^a, \tau^b) \neq 0$ given that $\lambda_n \neq \lambda_j$ for at least some n and j . We conclude that $JH \in [0, 1)$ and $JH < H$.

7.2 Proofs

Proof of Proposition 3. We start to solve the game by characterizing search with evaluations. Note that our previous result on the utility of an evaluation draw u_e^i (9) continues to hold for all consumers as taste intensity β is common for both types. Let us assume for now that a positive mass of evaluating consumers of both types is present in the market. We will show below that this is always the case in equilibrium. The aggregate distribution of evaluating consumers σ will now be composed by consumers of both types. Denote by $c_e^{t,i}$ the indifferent evaluator of type t . We define s_e^t as the share of consumers of type t among the population of consumers that search with evaluations, where

$$s_e^t = \frac{(c_e^{t,i}/\bar{c})s^t}{(c_e^{t,i}/\bar{c})s^t + (c_e^{-t,i}/\bar{c})s^{-t}}. \quad (31)$$

The aggregate sales distribution generated by evaluating consumers σ will depend on the distribution of matches of both consumer types. The probability that an evaluating consumer matches product n given the share of types is $(1/N)(s_e^a \lambda_n^a + s_e^b \lambda_n^b)$ and the probability of a match over all products is given by β for both types. This implies

$$\sigma_n = \frac{(1/N)(s_e^a \lambda_n^a + s_e^b \lambda_n^b)}{\beta} = s_e^a \tau_n^a + s_e^b \tau_n^b. \quad (32)$$

We now turn to the problem of unmatched consumers in the third stage given a price level p . These consumers may seek recommendations or choose to stay out of the market. Assume that a positive mass of consumers has searched with evaluations in the second stage. The expected probability of a match on each recommendations draw for a consumer of type t , denoted by α^t , is given by

$$\alpha^t = \sum \sigma_n \lambda_n^t = \sum (s_e^t \tau_n^t + s_e^{-t} \tau_n^{-t}) (\beta N \tau_n^t) = \beta N (s_e^t H + s_e^{-t} JH). \quad (33)$$

This expression measures the concentration of the distribution of evaluating consumers with respect to the taste distribution of consumers of type t . The result implies that the match probability when seeking recommendations will differ across types. As $\partial \alpha^t / \partial s_e^t > 0$, the larger the share of a consumer's own type among evaluating consumers, the larger the match probability on a recommendations draw. And conversely, $\partial \alpha^t / \partial s_e^{-t} < 0$. As s_e^t is a function of the indifferent evaluators of both types, this also implies that $\partial \alpha^t / \partial c_e^{t,i} > 0$ and $\partial \alpha^t / \partial c_e^{-t,i} < 0$.

Note that both $\alpha^t > \beta$ or $\alpha^t \leq \beta$ are possible depending on s_e^t , given that $1/N < H \leq 1$ and $0 \leq JH < H$ (these results are derived in subsection 7.1). Thus a recommendation draw may yield a higher or lower probability of a match than an evaluation draw, but will always be more costly due to r . As α^t is increasing in the share of evaluating consumers of type t , we can identify the critical share \bar{s}^t for $\alpha^t = \beta$,

$$\beta N (\bar{s}^t H + (1 - \bar{s}^t) JH) = \beta, \quad (34)$$

which is symmetric for both types, so $\bar{s}^t = \bar{s}$. Decomposing H and JH as a function of Var and Cov (see subsection 7.1) and rearranging,

$$\bar{s} = \frac{Cov(\tau^t, \tau^{-t})}{Cov(\tau^t, \tau^{-t}) - Var(\tau^t)}. \quad (35)$$

Hence $\alpha^t \leq \beta$ if and only if $s_e^t \leq \bar{s}$. Note that $\bar{s} > 0$ only if $Cov(\tau^t, \tau^{-t}) < 0$, as $Var(\tau^t) > Cov(\tau^t, \tau^{-t})$ due to the fact that both distributions differ by a permutation of their elements. So recommendations may not be preferred over evaluations for consumers of type t when both taste distributions are negatively correlated (negative covariance) and the share of evaluating consumers of their own type is low. In this case, most product recommendations originate from consumers of

the other type and are unlikely to yield a match. Furthermore, the upper bound of $Var(\tau^t)$ and the lower bound of $Cov(\tau^t, \tau^{-t})$ imply that the maximum feasible value for \bar{s} is $\frac{1}{N}$. Therefore $\bar{s} < 1/2$ given that $N \geq 3$ and recommendations always yield a higher match probability than evaluations for at least one consumer type in the market.

If $s_e^t \leq \bar{s}$, a consumer of type t would benefit of discarding a recommended product from the set of products to sample. For these consumers, recommendations identify products that are unlikely to yield a match. Hence recommendations are still an informative signal. The informativeness of recommendations in this case, however, quickly decreases with N with respect to the case $s_e^t > \bar{s}$. To see this, define by β_{-j} the taste intensity over the reduced assortment of $N - 1$ products after discarding product j ,

$$\beta_{-j} = \frac{\sum_{-j} \lambda_n}{N - 1}. \quad (36)$$

The probability of a match when sampling is now β_{-j} . The informativeness of the signal can be expressed as the improvement gained in the match probability when sampling with and without the recommendation,

$$\beta_{-j} - \beta = \frac{\sum_{-j} \lambda_n}{N - 1} - \frac{\sum \lambda_n}{N} = \frac{(\sum \lambda_n) - N\lambda_j}{N^2 - N}. \quad (37)$$

When N is large however,

$$\text{Limit}_{N \rightarrow \infty} \beta_{-j} - \beta = 0, \quad (38)$$

and the improvement gained by discarding products converges to zero. On the other hand, when $s^t > \bar{s}$ and recommendations are valuable to identify relevant products,

$$\text{Limit}_{N \rightarrow \infty} \alpha^t - \beta > 0. \quad (39)$$

We conclude that for large assortments and a positive recommendation cost r , the discarding strategy will not pay off.

The expected utility of seeking a new recommendation for consumer i of type t in the third stage of the game is given by

$$u_r^{t,i} = \alpha^t(u - p) - r - c^i. \quad (40)$$

Note that the $u_r^{t,i}$ differs both across types due to α^t and within types depending on c^i . For any given consumer, however, the utility of a successive draw is constant throughout the search. 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 by $c_r^{t,i}$,

$$c_r^{t,i} = \alpha^t(u - p) - r. \quad (41)$$

Consumers of type t with a sampling cost $c^{t,i} \leq c_r^{t,i}$ choose to search in the third stage and consumers 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 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 may differ between them, as in general $\alpha \neq \beta$, consumers will evaluate the utility of the complete search process to locate a match. 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 to locate a match, $u_{t,i}^r = u_{t,i}^e$. The expected number of draws required for a match with evaluations and recommendations are given by β^{-1} and α^{-1} respectively,

$$\begin{aligned} u - p - \frac{r + c_e^{t,i}}{\alpha^t} &= u - p - \frac{c_e^{t,i}}{\beta} \\ c_e^{t,i} &= \frac{\beta r}{\alpha^t - \beta}. \end{aligned} \quad (42)$$

Note that $c_e^{t,i}$ is given by an implicit equation, as α^t is a function of the shares of evaluating consumers, s_e^t and s_e^{-t} , which in turn are a function of $c_e^{t,i}$ and $c_e^{-t,i}$. So the equilibrium is defined by a system of implicit equations, one for each consumer type. We next argue that the solution to this system must satisfy $c_e^{a,i} < c_e^{b,i}$ and $s_e^a > s_e^b$. First, consider the case $c_e^{a,i} = c_e^{b,i}$. This requires that $\alpha^a = \alpha^b$, which implies that $s_e^a = s_e^b$ by 33. But on the other hand, equation (31) implies

that $s_e^a > s_e^b$ if $c_e^{a,i} = c_e^{b,i}$ as there is a larger share of consumers of type A in the population, which is a contradiction. Next, consider the case $c_e^{a,i} > c_e^{b,i}$. This requires that $\alpha^a < \alpha^b$, which implies that $s_e^a < s_e^b$ by 33. But in this case equation (31) implies that $s_e^a > s_e^b$, which again is a contradiction. Hence the only feasible solution must satisfy $c_e^{a,i} < c_e^{b,i}$ and $s_e^a > s_e^b$, which is compatible with equation (31). Although $s_e^a > s_e^b$, the solution implies that mainstream consumers are underrepresented in the population of evaluators, $s_e^a < s^a$ and $s_e^b > s^b$.

Consumers of type t with an evaluation cost $c_e^{t,i} \in [0, c_r^{t,i})$ prefer to search with evaluations in the second stage over seeking recommendations. Hence two cases exist for each consumer type, either $c_e^{t,i} \leq c_r^{t,i}$ or $c_e^{t,i} > c_r^{t,i}$. Which cases hold in equilibrium will depend on the match probability with recommendations α^t and the price level p . Given that in equilibrium $s_e^a > s_e^b$, this implies that $\alpha^a > \alpha^b$ by (33) and $\alpha^a > \beta$ given that $\bar{s} < 1/2$. On the other hand, α^b may be above or below β , depending on s_e^b and \bar{s} . If $s_e^b \leq \bar{s}$ then $\alpha^b \leq \beta$ and consumers of type B will search with evaluations only. In this case, evaluations dominate recommendations due to a higher match probability on each draw and a lower cost.

We next analyze the impact of p on $c_r^{t,i}$. As $c_r^{t,i}$ is a decreasing function of price level p , we can identify the boundary price \bar{p}^t such that $c_e^{t,i} = c_r^{t,i}$ by equating both expressions,

$$\bar{p}^t = u - \frac{r}{\alpha^t - \beta}. \quad (43)$$

If the firm's price is low, $p < \bar{p}^t$, then $c_e^{t,i} < c_r^{t,i}$ and consumers of type t will seek recommendations whenever $s_e^t > \bar{s}$. If $p \geq \bar{p}^t$, then $c_e^{t,i} \geq c_r^{t,i}$ and consumers of type t search only with evaluations. Also note that $\bar{p}^a > \bar{p}^b$, as we have established that in equilibrium $\alpha^a > \alpha^b$.

We can now characterize the search strategy of all consumers as a function of p and s_e^t given \bar{p}^t and \bar{s} . All possible cases are listed in the following table:

	Type A	Type B
$p \in [0, \bar{p}^b)$ and $s_e^b > \bar{s}$	Seek recommendations	Seek recommendations
$p \in [0, \bar{p}^b)$ and $s_e^b \leq \bar{s}$	Seek recommendations	Evaluations only
$p \in [\bar{p}^b, \bar{p}^a)$	Seek recommendations	Evaluations only
$p > \bar{p}^a$	Evaluations only	Evaluations only

If consumers of type t seek recommendations in equilibrium then $c_e^{t,i}$ and $c_r^{t,i}$ are given by (42) and (41) respectively. Consumers of type t with sampling cost $c^i \in [0, c_e^{t,i})$ search with evaluations, consumers such that $c^i \in [c_e^{t,i}, c_r^{t,i})$ seek recommendations and consumers such that $c^i > c_r^{t,i}$ stay out of the market. If types t search only with evaluations then $c_e^{t,i}$ is given by the consumer strictly indifferent between evaluating and not participating as derived for the case of homogeneous taste in 10,

$$c_e^{t,i} = \beta(u - p), \quad (44)$$

and only consumers such that $c^i \in [0, c_e^{t,i})$ participate in the market and search with evaluations.

Next, we derive the sales distribution generated by consumers of type t seeking recommendations, ρ^t . The probability that a consumer matches product n when drawing a recommendation is given by $\sigma_n \lambda_n^t$ and the probability of a match over all products is given by α^t . This implies

$$\rho_n^t = \frac{\sigma_n \lambda_n^t}{\alpha^t} = \frac{\tau_n^t [s_e^t \tau_n^t + s_e^{-t} \tau_n^{-t}]}{s_e^t H + s_e^{-t} JH}. \quad (45)$$

The probability that a consumer of type t matches with product n depends on both τ_n^t and τ_n^{-t} , as recommendations originate from consumers of both types. This interaction between types only impacts the sales distribution generated by recommendations whenever taste similarity is positive, that is, $\tau_n^t \tau_n^{-t} > 0$ for some n given that jointly relevant products exist. Consider the case with positive taste similarity. If consumers have quasiuniform taste, $\tau_n^t \tau_n^{-t} = \tau_n^t \tau_n^t$ for jointly relevant products and $\tau_n^t \tau_n^{-t} = 0$ for all others. The sales distribution generated by recommendations is then skewed towards jointly relevant products.

If consumers have non-quasiuniform taste, then $\tau_n^t \tau_n^{-t} > 0$ for jointly relevant products but may differ from $\tau_n^t \tau_n^t$. For products where $\tau_n^t \tau_n^{-t} \neq \tau_n^t \tau_n^t$, either $\tau_n^t > \tau_n^{-t}$ or $\tau_n^t < \tau_n^{-t}$ and the interaction between types weighs the sales share of jointly preferred products according to the taste and evaluating share s_e^t of both consumer types. In equilibrium we have established that $s_e^a > s_e^b$. Thus the interaction has a larger impact on types B , and consumers of type B have a higher probability of matching with jointly relevant products preferred by type A than conversely. Therefore the sales distribution generated by recommendations is not only skewed towards products with high match probability, as in the basic model with non-quasiuniform taste, but the sales share

of jointly preferred products is influenced by the taste of the other consumer type, with a larger impact for consumers of type B .

We next turn to the first stage of the game and analyze the firm's pricing problem. The demand curve is composed of linear components, it is continuous and (non-strictly) convex. The number of components is either three if consumers of type B seek recommendations at positive prices or two otherwise (a price range for consumers of type A to seek recommendations always exists). If $s_e^b > \bar{s}$, the demand curve has three components and is non-differentiable at both \bar{p}^a and \bar{p}^b . If $s^b \leq \bar{s}$ the demand curve has two components and is non-differentiable at \bar{p}^a .

For the firm choosing a price to maximize profits, each component describes a concave profit curve. Each profit curve lies above the others in its own price range and intersect at the points where the demand curve is non-differentiable.

Consider firm profits in the price range $p \in [0, \bar{p}^b)$ when $s_e^b > \bar{s}$. Consumers of both types seek recommendations, so participation is given by c_r^a and c_r^b ,

$$\pi_{r,r} = \left[\frac{c_r^a}{\bar{c}} s^a + \frac{c_r^b}{\bar{c}} s^b \right] (p - t) = \left[\frac{\alpha^a(u-p) - r}{\bar{c}} s^a + \frac{\alpha^b(u-p) - r}{\bar{c}} s^b \right] (p - t). \quad (46)$$

The profit-maximizing price for this demand curve is

$$p_{r,r} = \frac{u+t}{2} - \frac{r}{2(s^a \alpha^a + s^b \alpha^b)}. \quad (47)$$

Firm profits in the price range $p \in [\bar{p}^b, \bar{p}^a)$ when $s^b > \bar{s}$ and the range $p \in [0, \bar{p}^a)$ when $s^b \leq \bar{s}$ are characterized by consumers of type A seeking recommendations and consumers of type B searching only with evaluations. Participation is then given by c_r^a and c_e^b ,

$$\pi_{r,e} = \left[\frac{c_r^a}{\bar{c}} s^a + \frac{c_e^b}{\bar{c}} s^b \right] (p - t) = \left[\frac{\alpha^a(u-p) - r}{\bar{c}} s^a + \frac{\beta(u-p)}{\bar{c}} s^b \right] (p - t). \quad (48)$$

The profit-maximizing price in this case is

$$p_{r,e} = \frac{u+t}{2} - \frac{s^a r}{2(s^a \alpha^a + s^b \beta)}. \quad (49)$$

If the demand curve has two components, the optimum price is given by $p_{r,e}$. If the demand

curve has three components, the firm compares profits at price levels $p_{r,r}$ and $p_{r,e}$. In this case, boundary conditions must be checked to ensure each maximum is contained within the range in which it is well defined. If $p_{r,r} \leq \bar{p}^b$, the solution must be given by $p_{r,e}$. Conversely, if $p_{r,e} \geq \bar{p}^b$ the firm's solution is $p_{r,r}$.

Our restriction on r ensures the firm's profit-maximizing price falls in either of these demand ranges, which in turn ensures that consumers of type A always seek recommendations in equilibrium. For the profit maximizing price not to fall in the range $p \geq \bar{p}^a$, profits attainable in this range must be lower than those in the range $p < \bar{p}^a$. Denote by $p_{e,e}$ the profit maximizing price in this range, which simplifies to the solution obtained in the homogeneous case with evaluations, $p_{e,e} = p_e$ 13. A sufficient condition given the properties of the profit curves is that $p_{e,e} < \bar{p}^a$, which reduces to $r < \frac{1}{2}(u-t)(\alpha^a - \beta)$. Given that $s_e^a > 1/2$ in equilibrium and substituting α^a as a function of H and JH , $r \leq \frac{1}{4}((H + JH)N - 2)\beta(u-t)$ ensures $p_{e,e} < \bar{p}^a$.

We next show that a positive mass of consumers of both types always search with evaluations in equilibrium. Whenever consumers of type t search with recommendations, $\alpha_r^t > \beta$ must hold and it follows that $c_e^{t,i} = \frac{\beta r}{\alpha_r^t - \beta} > 0$. Whenever types t search only with evaluations, $c_e^{t,i} = \beta(u-p) > 0$ given that $p < u$. Hence in equilibrium word of mouth recommendations always originate from consumers of both types. ■

Proof of Proposition 4. We solve the equilibrium with the recommender system. Denote by σ^t the sales distribution generated by consumers of type t when searching with evaluations. By Proposition 1, $\sigma_n^t = \tau_n^t$. The sales distribution generated by all evaluating consumers is given by σ in 32. Let us assume for now that a positive mass of evaluating consumers of both types is present in the market. We will show below that this is always the case in equilibrium with the recommender. Consider the match probability for consumers when searching with the recommender. Recommendations are drawn from σ^t with probability q , and from σ with probability $1 - q$. The probability of a match on each recommendation draw is then given by

$$\alpha_{rs}^t = \sum[q\sigma_n^t + (1-q)\sigma_n]\lambda_n^t = \beta N[(q + (1-q)s_e^t)H + (1-q)s_e^{-t}JH]. \quad (50)$$

Clearly, the match probability improves with quality as $\partial\alpha_{rs}^t/\partial q > 0$. Also note that $\frac{\partial\alpha_{rs}^t}{\partial q\partial s_t} < 0$

and the improvement gained with the recommender is larger for the type with a lower share in the mass of evaluating consumers. Recall that the match probability with word of mouth is given by α^t in 33. It follows that $\alpha_{rs}^t - \alpha^t > 0$ if $q > 0$, increasing in q and decreasing in s_e^t . When the recommender system is of minimum quality and $q = 0$, this expression is equivalent to that of word of mouth. When the $q = 1$ the recommender ensures all recommendations originate from evaluating consumers of the same type, and α_{rs}^t is equivalent to the match probability derived in the homogeneous consumer case, $\alpha_{rs}^t = \beta NH$. This implies that given a common cost r of obtaining a recommendation from the recommender or from word of mouth, consumers strictly prefer the recommender system when $q > 0$ and are indifferent when $q = 0$.

The critical share \bar{s} for consumers of type t to be indifferent between searching with evaluations and using the recommender system is now given by

$$\bar{s}_{rs} = \frac{Cov(\tau^t, \tau^{-t})}{[Cov(\tau^t, \tau^{-t}) - Var(\tau^t)](1 - q)} - \frac{q}{1 - q}. \quad (51)$$

And $\bar{s}_{rs} < \bar{s}$ if $q > 0$ and $\frac{\partial \bar{s}_{rs}}{\partial q} < 0$. This implies that, with the recommender, consumers require a smaller share of evaluating consumers of their type to benefit from recommendations.

The indifferent recommendation seeker is given by,

$$c_r^{t,i} = \alpha_{rs}^t (u - p) - r. \quad (52)$$

And the participation of consumers of type t in the market, all other factors equal, increases with the quality of the recommender as $\frac{\partial \alpha_{rs}^t}{\partial q} > 0$.

The indifferent evaluator of type t is given by

$$c_e^{t,i} = \frac{\beta r}{\alpha_{rs}^t - \beta}. \quad (53)$$

Note that if $q = 0$ then $\alpha_{rs}^t = \alpha^t$ and the solution for $c_e^{t,i}$ is equivalent to that of word of mouth. If $q \in (0, 1)$ the equilibrium satisfies the same properties as our previous analysis, given that α_{rs}^t is increasing in s_e^t , decreasing in s_e^{-t} and $\alpha_{rs}^a = \alpha_{rs}^b$ if and only if $s_e^a = s_e^b$. Hence the solution to the system of implicit equations is still characterized by $c_e^{a,i} < c_e^{b,i}$ and $s_e^a > s_e^b$. In the equilibrium

with the recommender, as $\alpha_{rs}^t > \alpha^t$, recommendations are more valuable for all consumers and $c_e^{t,i}$ is lower for both types. The solution satisfies that $c_e^{a,i} - c_e^{b,i}$ is decreasing in q , so the recommender reduces the asymmetry between both types. If $q = 1$ then $\alpha_{rs}^a = \alpha_{rs}^b = \beta NH$ and the solution must be symmetric for both types and $c_e^{a,i} = c_e^{b,i}$.

The boundary price for consumers of type t to seek recommendations is now

$$\bar{p}_{rs}^t = u - \frac{r}{\alpha_{rs}^t - \beta}. \quad (54)$$

Note that $\bar{p}_{rs}^t > \bar{p}^t$ whenever $q > 0$ given that $\alpha_{rs}^t > \alpha^t$, and the price ranges where consumers search with recommendations are larger due to the recommender.

We can now characterize the equilibrium with the recommender system.

	Type A	Type B
$p \in [0, \bar{p}_{rs}^b)$ and $s_e^b > \bar{s}_{rs}$	Seek recommendations	Seek recommendations
$p \in [0, \bar{p}_{rs}^b)$ and $s_e^b \leq \bar{s}_{rs}$	Seek recommendations	Evaluations only
$p \in [\bar{p}_{rs}^b, \bar{p}_{rs}^a)$	Seek recommendations	Evaluations only
$p > \bar{p}_{rs}^a$	Evaluations only	Evaluations only

The profit-maximizing price for the firm in the range $p \in [0, \bar{p}_{rs}^b)$ when $s_e^b > \bar{s}_{rs}$ is

$$p_{r,r} = \frac{u+t}{2} - \frac{r}{2(s^a \alpha_{rs}^a + s^b \alpha_{rs}^b)}. \quad (55)$$

The profit-maximizing price in the range $p \in [0, \bar{p}_{rs}^b)$ and $s_e^b \leq \bar{s}_{rs}$ or $p \in [\bar{p}_{rs}^b, \bar{p}_{rs}^a)$ is

$$p_{r,e} = \frac{u+t}{2} - \frac{s^a r}{2(s^a \alpha_{rs}^a + s^b \beta)}. \quad (56)$$

Our restriction on r from section 4 continues to ensure that the firm's optimal price falls in either of these demand ranges (which now holds true for larger values of r). It can be shown that $\frac{\partial c_r^{t,i}}{\partial \alpha_{rs}^t} > 0$ given $p_{r,r}$ or $p_{r,e}$, so whenever consumers of type t search with recommendations their market participation is increasing in q . Therefore, both consumer participation and the firm's optimal price are increasing in q , and as a result firm profits are also increasing with q . Furthermore, $\frac{\partial u_r^{t,i}}{\partial \alpha_{rs}^t} > 0$ given $p_{r,r}$ or $p_{r,e}$, so the reduction in search costs is not offset by the price increase. Hence

the surplus enjoyed by consumers searching with recommendations also increases with q . Consumers searching with evaluations, however, are worse off as prices increase but their search costs to locate a match are unaffected. The aggregate impact on consumer surplus is ambiguous.

We next show that a positive mass of consumers of both types always search with evaluations in equilibrium. Whenever consumers of type t search with recommendations, $\alpha_{rs}^t > \beta$ must hold and it follows that $c_e^{t,i} = \frac{\beta r}{\alpha_{rs}^t - \beta} > 0$, so a positive mass of consumers of type t choose to search with evaluations. Whenever types t search only with evaluations, $c_e^{t,i} = \beta(u - p) > 0$ given that $p_{r,r} < u$ and $p_{r,e} < u$. Hence in equilibrium the recommender always receives input from both consumer types in the market.

Two special cases exist. When $q = 0$, the equilibrium is equivalent to that of word of mouth. When $q = 1$, consumer participation and firm profits are equivalent to the homogeneous case analyzed in section 3 given that taste intensity, taste concentration and search frictions across the consumer population are equivalent in both cases. ■

Proof of Proposition 5. When taste is quasiuniform $\lambda_n^t \in \{0, a\}$ for all n , where $a \in (0, 1]$. Denote the number of relevant products by R , which is common for both types as taste profiles differ by a permutation of elements. It follows that $\tau_n^t = \frac{1}{R}$ if $\lambda_n = a$, and $\tau_n^t = 0$ otherwise.

We start by characterizing the sales distribution generated by consumers of type t with a recommender in the market. When $q = 1$, we have established that both consumer types search with recommendations in addition to evaluations. The sales distribution for any consumer type has two components; the sales distribution generated by evaluating consumers and that generated by consumers searching with recommendations. The sales distribution generated by evaluating consumers of type t is given by $\sigma_n^t = \tau_n^t$. So the mass of evaluating consumers distribute uniformly over products relevant to them. Denote by $\hat{\rho}^t$ the sales distribution generated by consumers of type t searching with recommendations from the recommender. As a perfect recommender ensures that all recommendations originate from consumers of the same type, $\hat{\rho}^t$ is equivalent to the homogeneous consumer case and is given by 21,

$$\hat{\rho}_n^t = \frac{\tau_n^t{}^2}{H} = \tau_n^t. \quad (57)$$

And consumers searching with recommendations follow the same distribution as evaluating con-

sumers. Hence consumers of type t that participate in the market distribute uniformly over products relevant to them.

We next characterize the sales distribution generated by consumers of type t when no recommender is available in the market. As established in Proposition 3, either both consumer types search with recommendations in equilibrium or only mainstream consumers do. The sales distribution generated by evaluating consumers of type t is given by $\sigma_n^t = \tau_n^t$. If consumers of type t search only with evaluations, the sales distribution generated by these consumers is equivalent to the case with the recommender in the market. If consumers of type t also search with recommendations, the sales distribution component generated by recommendations is given by 45 as derived in our analysis of taste heterogeneity,

$$\rho_n^t = \frac{\tau_n^t [s_e^t \tau_n^t + s_e^{-t} \tau_n^{-t}]}{s_e^t H + s_e^{-t} JH}. \quad (58)$$

Two possible cases arise depending on taste similarity. On the one hand, when taste similarity is zero, no jointly relevant products exist and $\tau_n^t \tau_n^{-t} = 0$ for all n . Therefore, $\rho_n^t = \tau_n^t$ and consumers searching with recommendations generate the same sales distribution as evaluating consumers. In this case consumers of type t distribute uniformly over products relevant to them, and the solution is equivalent to the case with the recommender in the market. On the other hand, when taste similarity is positive, jointly relevant products exist and $\tau_n^t \tau_n^{-t} > 0$ for these products. Inspection of 58 reveals that jointly relevant products enjoy larger demand than those relevant only to type t . Thus if consumers of type t search with recommendations in addition to evaluations, the sales distribution they generate is skewed towards jointly relevant products. As at least one consumer type searches with recommendations when no recommender is available, it is always true that the sales distribution generated by consumers differs with the recommender when jointly relevant products exist.

We turn to analyze the impact of the recommender on sales concentration. We consider separately the cases where taste similarity is zero and when it is positive. Let us start by the case in which taste similarity is zero. No jointly relevant products exist. Consider the market share of products without the recommender in the market. Market shares depend both on the share of participating consumers of each type and the sales distribution they generate. We established in

Proposition 3 that a larger share of participating consumers are of type A . Both a larger proportion participate and there is a larger mass of consumers of type A in the population. This implies that the market share of products may take any of three values. Irrelevant products for both consumer types obtain zero market share, as no consumer purchases them. Products relevant only for type B enjoy market share m^b . Products relevant only for type A enjoy market share m^a . Note that $m^a > m^b > 0$ as the share of participating consumers of type A in the market is larger than that of B .

We next argue that introducing the recommender reduces the sales concentration in the market when no jointly relevant products exist. Consider the market shares of products with the recommender. The sales distribution generated by each consumer type in the market is unaffected by the recommender. The mass of participating consumers changes, however, as we established in Proposition 4 that a perfect recommender ensures an equal proportion of consumers of both types participate. There is a *participation effect*. This implies that the share of participating consumers of type B is larger with the recommender. The new market shares of relevant products with the recommender, \hat{s}^a and \hat{s}^b , must satisfy $\hat{s}^a < s^a$ and $\hat{s}^b > s^b$. Note that $\hat{s}^a > \hat{s}^b$ as there is a larger mass of consumers of type A in the population. So the recommender introduces a non rank altering transfer from high to low market share products, reducing the sales concentration in the market.

We next consider the case in which taste similarity is positive. In this case jointly relevant products exist. Consider the market share of products in the market without the recommender. Market shares of products may take any of four values. Irrelevant products obtain zero market share. Products relevant only to type B enjoy market share m^b . Products relevant only to type A enjoy market share m^a . Jointly relevant products obtain share m^{ab} . Note that $m^{ab} > m^a$ and $m^{ab} > m^b$, as jointly relevant products are demanded by consumers of both types and their share in the sales distribution of each type is at least as high as that of products relevant only for that type. Also note that either $m^a > m^b$ or $m^a < m^b$. Although a larger share of participating consumers are of type A , they search with recommendations and the sales distribution they generate is skewed towards jointly relevant products. The sales distribution of type B is not skewed if they search only with evaluations, and in this case products relevant only to type B may enjoy a larger market share than those relevant only to type A .

We next argue that introducing the recommender reduces the sales concentration in the market when jointly relevant products exist. Consider the market shares of products with the recommender. There is a change in the sales distribution generated by consumer types that search with recommendations in absence of the recommender, a *sales distribution effect*. Furthermore, the share of participating consumers of type B is also larger with the recommender. The *participation effect* is also present. We consider both effects in sequence. If both effects reduce the concentration of sales, the net effect of introducing the recommender must also reduce sales concentration in the market.

We first introduce the sales distribution effect. We need only consider consumer types that search with recommendations in absence of the recommender, as we have shown that only the sales distribution generated by these types is affected. Without the recommender, their sales distribution is skewed towards jointly relevant products. With the recommender, the sales distribution is uniform over relevant products. This implies that the change in the sales distribution with the recommender shifts weight from jointly relevant products to relevant products (for this type), which must result in a product market share shift in the same direction. Denote the market shares with the recommender when accounting only for the sales distribution effect by \tilde{s}^{ab} , \tilde{s}^a and \tilde{s}^b . These market shares must satisfy $\tilde{s}^{ab} < s^{ab}$ and $\tilde{s}^t > s^t$ for the types affected. Note that $\tilde{s}^{ab} > \tilde{s}^a > \tilde{s}^b$ must hold, as jointly relevant products benefit from the demand of both types and there is a larger share of participating consumers of type A . The sales distribution effect implies a non rank altering transfer from high to low market share products, reducing the sales concentration in the market.

We next introduce the participation effect. This effect ensures that the share of participating consumers of type B is larger with the recommender. Denote the market shares with the recommender when accounting for both the sales distribution and the participation effects by \hat{s}^{ab} , \hat{s}^a and \hat{s}^b . These market shares must satisfy $\hat{s}^{ab} = \tilde{s}^{ab}$, $\hat{s}^a < \tilde{s}^a$ and $\hat{s}^b > \tilde{s}^b$ given the increased share of participating consumers of type B . The market shares of jointly relevant products are unaffected as their share in the sales distribution of both types is common (given that we have already accounted for the sales distribution effect). Note that $\hat{s}^a > \hat{s}^b$ as consumers now participate in equal proportion but there is a larger mass of consumers of type A in the population. So the participation effect implies a non rank altering transfer from high to low market share products, reducing the sales

concentration. Hence the recommender reduces the concentration of sales in the market whenever consumers exhibit quasiuniform taste. ■

References

- [1] Adler, Moshe (1985) ‘Stardom and Talent’, *American Economic Review*, Vol. 75, No. 1, pp. 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*, Vol. 17, No. 6, pp. 734-749.
- [3] Anderson, Chris (2004) ‘The Long Tail’, *Wired*.
- [4] Anderson, Chris (2006) ‘The Long Tail’, Hyperion press, New York, NY.
- [5] Avery, Chris, Paul Resnick and Richard Zeckhauser (1999) ‘The Market for Evaluations’, *American Economic Review*, Vol. 89, No. 3, pp. 564-584.
- [6] Bakos, Yannis J. (1997) ‘Reducing Buyer Search Costs: Implications for Electronic Marketplaces’, *Management Science*, Vol. 43, No. 12, pp. 1676-1692.
- [7] 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, MIT Sloan.
- [8] Brynjolfsson, Erik, Yu Jeffrey Hu and Michael D. Smith (2006) ‘From Niches to Riches: Anatomy of the Long Tail’, *MIT Sloan Management Review*, Vol. 47, No. 4, pp. 67-71.
- [9] Chevalier, Judith A. and Dina Mayzlin (2006) ‘The Effect of Word of Mouth on Sales: Online Book Reviews’, *Journal of Marketing Research*, Vol. 43, pp. 345-354.
- [10] Dellarocas, Chrysanthos (2006) ‘Strategic Manipulation of Internet Opinion Forums: Implications for Consumers and Firms’, *Management Science*, Vol. 52, No. 10, pp. 1577-1593.

- [11] Ehrenberg, Andrew S. C., Gerald J. Goodhardt and T. Patrick Barwise (1990) ‘Double Jeopardy Revisited’, *Journal of Marketing*, Vol. 54, No. 3, pp. 82-91.
- [12] Elberse, Anita and Felix Oberholzer-Gee (2006) ‘Superstars and Underdogs: An examination of the Long Tail Phenomenon in Video Sales’, Working Paper, Harvard Business School.
- [13] Fleder, Daniel and Kartik Hosanagar (2007) ‘Blockbuster Culture’s Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity’, Working Paper, The Wharton School.
- [14] MacDonald, Glenn M. (1988) ‘The Economics of Rising Stars’, *American Economic Review*, Vol. 78, No. 1, pp. 155-166.
- [15] Resnick, Paul and Hal R. Varian (1997) ‘Recommender Systems’, *Communications of the ACM*, Vol. 40, No. 3, pp. 56-58.
- [16] Rosen, Sherwin (1981) ‘The Economics of Superstars’, *American Economic Review*, Vol. 71, No. 5, pp. 845-858.
- [17] Salganik, Matthew J., Peter Sheridan Dodds and Duncan J. Watts (2006) ‘Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market’, *Science*, Vol. 311, pp. 854-856.
- [18] Tucker, Katherine and Juanjuan Zhang (2007), ‘Long Tail or Steep Tail? A Field Investigation into How Online Popularity Information Affects the Distribution of Customer Choices’, Working Paper, MIT Sloan.