PRODUCT RECOMMENDATIONS AND MARKET CONCENTRATION[†]

EMILIO CALVANO, GIACOMO CALZOLARI, AND VINCENZO DENICOLÒ AND SERGIO PASTORELLO[*]

SEPTEMBER 2021

PRELIMINARY AND INCOMPLETE
PLEASE DO NOT CIRCULATE WITHOUT AUTHORS' PERMISSION.

We explore the impact of algorithms that use personal data to predict tastes and make recommendations, on consumption choices and in particular on the level of market concentration. We model consumer preferences in a flexible way that allows for varying degrees of vertical and horizontal product differentiation. The analysis confirms that these algorithms produce a significant increase in market concentration even with well-behaved models of demand and against well defined economic benchmarks. The paper goes on tackling the causes and consequences of such increase. It shows that the increase is not due to the "feedback loop" created by the endogeneity of the data. Moreover, we show that adopting these algorithms leads to lower equilibrium prices and higher consumer surplus despite the increased level of concentration. We discuss the implications of these findings for competition policy.

Keywords: Artificial Intelligence, Recommender Systems, Market Concentration, Product Differentiation.

J.E.L. codes: L41, L13, D43, D83.

1. INTRODUCTION

To navigate through the sea of products, consumers today use more sophisticated tools than they did in the past. Just as sailors have replaced the compass and the sextant by the GPS, consumers are replacing individual search and the word of mouth by algorithms, called *recommender systems* (hereafter RSs), which predict individual tastes and filter content. Because these algorithms can help consumers discover new products and find better matches to their preferences, they are believed to boost consumer engagement and sales. That is the reason why RSs are used extensively by platforms such as Netflix, Spotify, Amazon, Facebook and many others. Upon logging on these platforms, consumers are greeted with a tailored selection of items from the catalogue that they might enjoy or need.[1]

Evidently, RSs affect consumption choices and thus have the potential to change the structure of markets and the competition among the firms. From an industrial organization perspective, these developments therefore raise important and yet unresolved questions. Will markets become more or less competitive as consumers increasingly resort to RSs? Will incumbents become more powerful, or will it be easier for new entrants to challenge their dominance?

Little is known on these issues, and the existing views are often contradictory. The optimistic view regards RSs as means for reducing search costs and hence suggests that RSs may have pro-competitive effects. The pessimistic view, on the other hand, argues that RSs may generate a "rich get richer" dynamics whereby a handful of items are increasingly favored over time. This view has been substantiated by the computer science literature, which has largely documented RSs' tendency to increase concentration.[2] This may allegedly result in political polarization in social media, loss of diversity in culture, and so on. In economics, the risk is that higher market concentration may result in more market power and higher prices.

To shed light on this debate, in this paper we draw on both the economic and computer science literatures. The computer science literature has exposed the fallacy of conceptualizing RSs as a reduction in the cost of individual search. In fact, RSs change the process of information acquisition qualitatively. That is, they collect information from a large number of consumers and combine this information by exploiting the estimated similarities

[1] RSs are applied in many different environments, such as recommending news to users in social networks, financial products to investors (Robo-advising), job offers to job seekers, and possible referees to editors of academic journals and new articles to read to scholars. For an assessment of the business value of RSs see Jannach and Jugovac (2019).

[2] See, among many others, Fleder and Hosanagar (2009) and, more recently, Ciampaglia et al. (2018) and Goldfarb and Tucker (2019).

and dissimilarities among consumers. Trying to predict the economic consequences of RSs by mechanically applying models of individual search misses the 'collaborative' nature of these tools and may be misleading. In particular, if consumers searched more extensively but still individually, their choices would presumably better reflect the idiosyncrasies of their personal tastes. We should therefore observe more diversity, not less.

From the economic literature, we import a consistent model of consumer preferences and product characteristics that can generate various combinations of vertical and horizontal product differentiation. We also import meaningful benchmarks and criteria to assess the economic consequences of RSs, which are often lacking in the computer science literature. In particular, we contrast the outcomes produced by RSs with those arising when consumers can only rely on individual search.

Our research strategy is to construct simple but realistic RSs and analyze by means of numerical simulations their effects in a controlled economic environment. The analysis confirms that RSs produce an increase in market concentration even with well-behaved models of demand and against well defined economic benchmarks. Taking product characteristics and prices as given, that is to say, the market is more highly concentrated when consumers follow the recommendations of a RS than if they rely on individual search. This is true when products are vertically differentiated and also when they are purely horizontally differentiated – a case in which a decrease in the cost of individual search should have no effect on concentration. Our findings therefore confirm that it is misleading to regard RSs merely as devices to reduce individual search costs.

The paper then analyzes the causes and consequences of the increase in market concentration. As for the causes, in the computer science literature the standard explanation for the increase in concentration is the endogeneity of the data used to feed the RSs – the so called "feedback loop" effect. RSs, so the story goes, initially identify a few "good" products and recommend them to many consumers because they overestimate the similarity of their tastes. As time passes, RSs keep recommending the products they have recommended in the past because they acquire better information precisely on them. This reinforcement process amplifies small initial differences in products' perceived quality and induces RSs to create their own "champions" almost from scratch.

In contrast to this explanation, we show that market concentration arises already when the data used to feed the RSs are generated exogenously, in a purely random manner, rather than being endogenously produced by the RS itself. In some cases, the feedback-loop effect produces a further increase in concentration, but its additional contribution is relatively small. The increases in concentration must therefore be an effect of the inherent structure of the RSs, as discussed in greater detail below.

As for the consequences, it is well known that high levels of market concentration may have a dual meaning. On the one hand, they may be a signal of the weakness of competition in markets where few large firms possess significant market power. On the other hand, they may be a signal of vigorous competition among heterogeneous firms. The high market shares enjoyed by relatively few firms, that is to say, may reflect the process of output reallocation among producers that occurs when healthy competition selects the more efficient firms and weeds out the less efficient.

The difference between these two interpretations is that higher levels of market concentration tend to translate into higher prices in the former case, and into lower prices in the latter. To clarify the economic implications, we therefore look at the impact of RSs on equilibrium prices. To this end, we derive the system of demand functions implied by the process of individual search on the one hand and by the recommendations of our RSs on the other, and we then calculate the corresponding equilibrium prices. Proceeding in this way, we show that RSs generally lead to lower prices and higher consumer surplus. The effect is quite robust and is obtained for different combinations of vertical and horizontal differentiation. In particular, when vertical differentiation prevails, the adoption of RSs increases the market shares of the dominant products but at the same time reduces their prices.

After establishing these results, we inquire into the underlying mechanism. In particular, our analysis shows that the feedback loop is not the main cause of the increase in market concentration, but this raises the problem of what are, then, the true causes. We rule out by design two possible candidates. The first one is that platforms may deliberately manipulate recommendations so as to favor certain products (e.g., those with the highest margins, or their own products) and damage others. In our analysis, we use "neutral" RSs that are not instructed to support or penalize specific products. The second possible candidate is that RSs basically estimate a misspecificed model of consumer preferences. But the result holds even if the RSs estimate a correctly specified model of consumer preferences.

Still, RSs make mistakes in estimating consumers' true preferences, for various reasons. First and foremost, RSs can observe only a subset of all possible consumer/product pairs, and typically a very small one. Second, even for those pairs, RSs may observe only a noisy signal of the true utility that the consumer obtains by consuming that product. The reason why these mistakes systematically produce levels of market concentration higher than in the individual search model is briefly discussed the last part of the paper.

2. LITERATURE REVIEW

The economic literature on RSs is limited. As noted, in Economics RSs are often naturally conceptualized as tools that reduce search costs (Goldfarb and Tucker, 2019). Were this interpretation appropriate, RSs may allow consumers to discover rare or niche products, leading to a long-tail effect (Anderson 2008, Brynjolfsson et al. 2011). At the same time, low search cost can lead to a superstar effect (Rosen 1981), i.e., the discovery of superior products that attract more consumers. Bar-Isaac, Caruana, and Cunat (2012) show that the long-tail effect and the superstar effect may coexist when products are both vertically and horizontally differentiated, as they are in our model. However, RSs do not simply reduce search costs; as discussed in the introduction, they change the process of information acquisition qualitatively.

Another strand of the economics literature studies platforms' incentives to manipulate recommendations and steer consumer decisions toward certain products. In this literature, which includes Hagiu and Jullien (2011), Calvano and Jullien (2018) and Teh and Wright (2020), among others, RSs are modeled in a highly stylized way. Our analysis focuses on more sophisticated RSs but abstracts from the possibility that recommendations may be manipulated strategically.

A more indirectly related literature studies the effect of improved information about the distribution of product-match utility as perceived by consumers and the firms' price reaction (see e.g. Zhou, 2020 and Zhong, 2018). The possible connection with RSs comes from the fact that the informational improvement could be interpreted as the consequence of the adoption of RSs. However, the connection is loose because in those papers the change in the structure information is taken as exogenous and does not reflect the specificities of RSs.

The relevant computer science literature is much larger: see, for instance, Abdollahpouri et al. (2021) for a recent survey. This literature has identified a "popularity bias" of common RSs, particularly the ones based on collaborative filtering (as those considered in this paper). This popularity bias is the tendency of algorithms to disproportionately recommend popular items and disfavor others in spite of the fact that a significant number of users would actually prefer the latter as they better fit their idiosyncratic preferences.[3]

The computer science literature has categorized the possible sources of the popularity bias as "bias in the data" and "bias in the algorithms" (Abdollahpouri and Mansoury, 2020).

[3]Most analyses identify this bias by means of numerical simulations, as we do. However, Ciampaglia et al. (2018) demonstrates this bias theoretically in a stylized model where, depending on the users' cost of exploration, algorithms that recommend popular items may hinder the discovery of high-quality items.

The bias in the data is the fact that many users may intrinsically prefer the same items. As economists, we would not regard this as a bias but rather as a feature of consumers' true preferences, possibly reflecting vertical product differentiation.

The bias in the algorithms is more interesting and is generally attributed to the interaction between algorithms and data. It has been noted that this interaction may produce a feedback loop, as the algorithms are periodically retrained on the data that they have contributed to generate through their past recommendations (see e.g. Felder and Hosanagar, 2009; Chaney et al., 2018; Lee and Hosanagar, 2019; and Abdollahpouri and Mansoury, 2020).[4] The conventional wisdom is that this feedback loop induces homogenization of consumer decisions and increases market concentration (Zhang et al. 2020).

3. MODEL-BASED RECOMMENDER SYSTEMS

Ideally, one would like to study the RSs actually used by prominent firms, but the design of these RSs is kept confidential. As an alternative, here we focus on algorithms that figure prominently in the computer science literature, in the belief that they share the same underlying structure of those used in practice.

In particular, we concentrate on latent-factor, collaborative-filtering methods, also known as model-based systems.[5] Model-based systems are, indeed, popular in computer science and are "considered to be state-of-the art in recommender systems" (Aggarwal 2016, p. 91). Moreover, model-based collaborative-filtering systems are relatively simple and transparent, which facilitates the analysis of the ultimate drivers of observed outcomes. An additional advantage of these tools is that they are consistent with standard structural econometric modeling of demand, as discussed in greater detail below.

3.1. *The basic problem*

Consider a finite set of users $i = 1, 2, ..., I$ who have preferences over a finite set of items $j = 1, 2, ...J$. The preferences are represented by real numbers $\{r_{ij}\}$, one for each user-item pair. User i prefers item j over item k if $r_{ij} \geq r_{ik}$. In the jargon of computer science, r_{ij} is i's "rating" of item j. In economics, it is natural to interpret r_{ij} as the utility that consumer i obtains from product j.

[4]To identify the bias in the algorithms, the computer science literature typically compares two cases, a benchmark case where the RS is trained only once, and the case where the RS is periodically retrained over new, endogenously generated data. As described in greater detail below, our procedure is different: the algorithm is retrained in each period in both cases, but in the exogenous data case the new data is generated randomly, whereas in the endogenous data case it is created by the RS's past recommendations.

[5]For a thorough treatment of RSs in computer science, see Aggarwal (2016).

In its simplest formulation, a RS's problem is to predict user ratings $\{r_{ij}\}$ starting from a subset of imperfectly observed ratings. The problem is typically formulated as a "matrix completion" problem. Let $\tilde{\mathbf{R}}$ denote a sparse $I \times J$ matrix whose entries are the observed ratings $\tilde{r}_{ij}$. The goal of the RS is to estimate the missing entries, completing the matrix $\mathbf{R}$.

3.2. *Latent factors*

Collaborative-filtering RSs fill the rating matrix $\mathbf{R}$ by exploiting row and column correlations of the sparse observed matrix $\tilde{\mathbf{R}}$. The implicit assumption is that users whose observed ratings are correlated are likely to also rate the unobserved items in a similar way. Likewise, items whose ratings are correlated across observed users are likely to be rated similarly by the unobserved ones.

This simple idea may be articulated and implemented in different ways, giving rise to different types of collaborative-filtering systems. In this paper, we focus on the so-called latent-factor RSs. These models posit a specific structure of the ratings, i.e., a relatively small set of latent variables, whose goal is to parsimoniously explain the observed ratings and predict the unobserved ones.

Specifically, the rating of item j by user i is viewed as the inner product of a vector of user-specific factors $\mathbf{t}_i$ and a vector of item-specific factors $\mathbf{v}_j$:

$$r_{ij} = \sum_{h=1}^{K} t_{ih} v_{jh}, \tag{1}$$

where K is the number of latent factors. In matrix notation:

$$\mathbf{R} = \mathbf{T}\mathbf{V}', \tag{2}$$

where $\mathbf{T}$ and $\mathbf{V}$ are the $(I \times K)$ and $(J \times K)$ matrices formed by the I vectors $\mathbf{t}_i$ and the J vectors $\mathbf{v}_j$, respectively.

Generally speaking, the variable t_{ih} may be thought of as user i's proclivity for factor h, and the variable v_{jh} as product j's affinity to factor h. In applications, the K factors typically have a semantic interpretation. For the algorithms, however, they need not have any specific meaning.

The algorithm chooses the number of factors $\hat{K}$ and the parameters $\hat{\mathbf{T}}$ and $\hat{\mathbf{V}}$ so as to

minimize some measure of the distance between the predicted ratings,

$$\hat{r}_{ij} := \sum_{h=1}^{\hat{K}} \hat{t}_{ih}\hat{v}_{jh} \tag{3}$$

and the observed ones, $\tilde{r}_{ij}$.

It therefore appears that the latent factor model that is the basis of our RSs is consistent with the Random Utility Model that underpins much of the structural empirical work on demand estimation. The main difference is that in econometric estimations, it is typically assumed that the choice attributes (i.e., the v_{jh}s) are observable while preferences (i.e., the t_{jh}s) are not. RSs, in contrast, assume that neither product attributes nor consumers tastes are observable. Thus, RSs *jointly* estimate both attributes and tastes.[6]

3.3. *Estimation procedure*

Generally speaking, the problem of the RS is to minimize a measure of prediction error on observables. Using the Euclidean distance, the estimated parameters $\hat{\mathbf{T}}$ and $\hat{\mathbf{V}}$ solve:

$$\min_{\hat{\mathbf{T}},\hat{\mathbf{V}}} \sum_{ij\in S} \left(\tilde{r}_{i,j} - \sum_{h=1}^{K} \hat{t}_{ih}\hat{v}_{jh} \right)^2, \tag{4}$$

where S denotes the set of all pairs (i, j) for which r_{ij} is (at least imperfectly) observed.[7] Once the latent factors are estimated, one can reconstruct the matrix of estimated ratings, $\hat{\mathbf{R}}$, by simply using the estimator $\hat{\mathbf{R}} = \hat{\mathbf{T}}\hat{\mathbf{V}}'$. With this matrix at hand, the algorithm is then able to rank all items for all users, and hence to make recommendations.

In real-life applications, the set S of observed ratings tends to be small; that is, the matrix $\tilde{\mathbf{R}}$ of observed ratings is very sparse. In this case, the estimation procedure (4) may cause overfitting. A common approach for addressing this problem is to use regularization techniques.[8] The idea is to penalize large values of the coefficients in the matrices $\hat{\mathbf{T}}$ and

[6]The case where some product attributes are observed by the RS can be handled by using hybrid RSs (also known as ensemble RSs). By symmetry, these algorithms can also deal with the case where some consumer tatses are observed.

[7]In most of the analysis, we assume that the RS known the true number of factors K. The case where K is also estimated by the RS will be taken up among the extensions.

[8]This approach is used not only for RSs but also in many other machine learning algorithms.

$\hat{\mathbf{V}}$ so as to decrease the variance of $\hat{\mathbf{R}}$. The actual minimization problem then becomes

$$\min_{\hat{\mathbf{T}},\hat{\mathbf{V}}} \sum_{ij\in S} \left(\tilde{r}_{i,j} - \sum_{h=1}^{K} \hat{t}_{ih}\hat{v}_{jh} \right)^2 + \lambda_t \sum_{i=1}^{I}\sum_{h=1}^{K} \hat{t}_{ih}^2 + \lambda_v \sum_{j=1}^{J}\sum_{h=1}^{K} \hat{v}_{ih}^2, \tag{5}$$

where the non-negative parameters λ_t and λ_v are the regularization weights (one of the hyper-parameters of the algorithm).[9] The downside of this regularization procedure is that it may introduce some bias in the estimates.

From a computational viewpoint, the minimization of the loss function (5) proceeds by calculating the gradient vector of partial derivatives w.r.t. the $J \times K + I \times K$ decision variables and using numerical methods to find values of $\hat{\mathbf{T}}, \hat{\mathbf{V}}$ such that the gradient vanishes. In vector notation, the gradient is

$$\mathbf{G} = \begin{pmatrix} \mathbf{E}\hat{\mathbf{V}} - \lambda_t \hat{\mathbf{T}} \\ \mathbf{E}'\hat{\mathbf{T}} - \lambda_v \hat{\mathbf{V}} \end{pmatrix} \tag{6}$$

where $\mathbf{E}$ is the $I \times J$ matrix of residuals $e_{ij} = \tilde{r}_{i,j} - \sum_{h=1}^{K} \hat{t}_{ih}\hat{v}_{jh}$. The numerical method used to find $\hat{\mathbf{T}}, \hat{\mathbf{V}}$ such that $\mathbf{G} = 0$ is, typically, some version of the Gradient Descent algorithm. The idea is to start with some arbitrary initial values for $\hat{\mathbf{T}}_0, \hat{\mathbf{V}}_0$ and the associated $\mathbf{G}_0$ and iterate according to the following recursive equation

$$\begin{pmatrix} \mathbf{T}_{t+1} \\ \mathbf{V}_{t+1} \end{pmatrix} = \begin{pmatrix} \mathbf{T}_t \\ \mathbf{V}_t \end{pmatrix} - \eta \mathbf{G}_t \tag{7}$$

where $\eta > 0$ is the step size and $\mathbf{T}_{t+1}, \mathbf{V}_{t+1}$ are the new values of the variables of interest at the end of the iteration. The iterations are repeated to convergence.[10]

4. ECONOMIC ENVIRONMENT

This paper analyzes by means of numerical simulations the effects of deploying the collaborative filtering RSs discussed in the previous section in a controlled economic environment. We now provide a description of this artificial environment.

Before going into the details, it may be worth noting that we focus on RSs that are not manipulated strategically and whose estimates are not systematically biased. We believe that this is a natural starting point for the economic analysis of RSs. While the strategic

[9] In our calculations, we set $\lambda_t =$ and $\lambda_v =$.

[10] This computation procedure, which is generally referred to as matrix factorization, comes in many different variants. For our simulations, we use a publicly available, state-of-the-art library (Chin et al. 2016) designed for parallelized computation.

use of RSs, or the possibility that their estimates may be systematically biased, raise important economic issues, these issues can be properly addressed only after the case without strategic or statistical biases is understood.

4.1. *Products and preferences*

In keeping with the approach just outlined, we model consumer preferences and product characteristics in a way that is consistent with the latent factor model used by the RS. Thus, we assume that consumer i's expected utility from consuming product j in period τ is

$$u_{ij} = \sum_{h=1}^{K} t_{ih} v_{jh}. \tag{8}$$

From an industrial organization perspective, it is natural to interpret (8) as a model of vertical product differentiation with multi-dimensional product quality. The coefficient $v_{jh} > 0$ is the h-th dimension of product j's quality, and the coefficient $t_{ih} > 0$ is the value that consumer i attaches to that dimension of quality.

The model is, indeed, one of pure vertical differentiation if $K = 1$. In this case, all consumers rank all products in the same way, albeit with varying intensity of preferences, and this makes the RS's task quite easy. As soon as $K \geq 2$, however, the model may exhibit some degree of horizontal product differentiation, and the RS's problem may become more difficult.

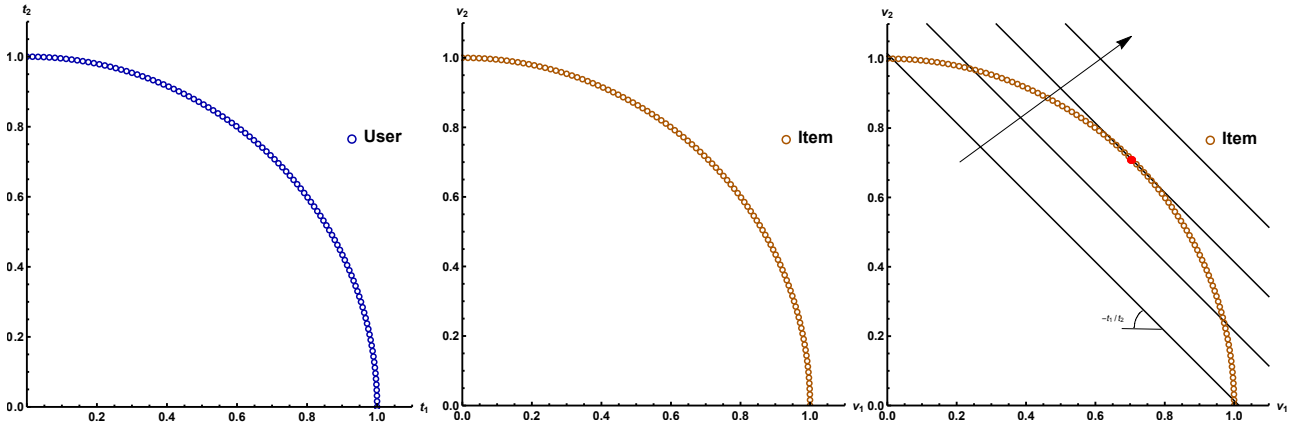
To illustrate, assume that $K = 2$, that the number of products equals the number of consumers (i.e., $I = J$), and that consumers $(t_{i1}, t_{i2}) \in \mathbb{R}^2_+$ are equally spaced on the portion of the unit circle lying in the positive orthant (Figure 1, center).[11,12] These assumptions are part of our baseline model of the product market, which is then closed by making different assumptions about the location of the products. These assumptions determine the extent to which products are horizontally or vertically differentiated.

A model of pure horizontal differentiation is obtained when products are placed on the same locus as consumers,[13] being equally spaced as well (Figure 1, left). Product differentiation is purely horizontal because each consumer prefers a distinct product. In particular, a consumer with coordinates (t_{i1}, t_{i2}) strictly prefers the product with the same coordi-

[11] The implicit normalization $t_{i1}^2 + t_{i2}^2 = 1$ is with no loss of generality.

[12] To be precise, in our analysis we assume that one consumer is located exactly on the 45 degree line. A half of the remaining consumers are equally spaced on the left arc, and a half on the right arc.

[13] Again, the implicit normalization $v_{j1}^2 + v_{j2}^2 = 1$ is with no loss of generality.



(a) User location (b) Item location (c) User $t_1 = t_2$ indifference map and optimal item

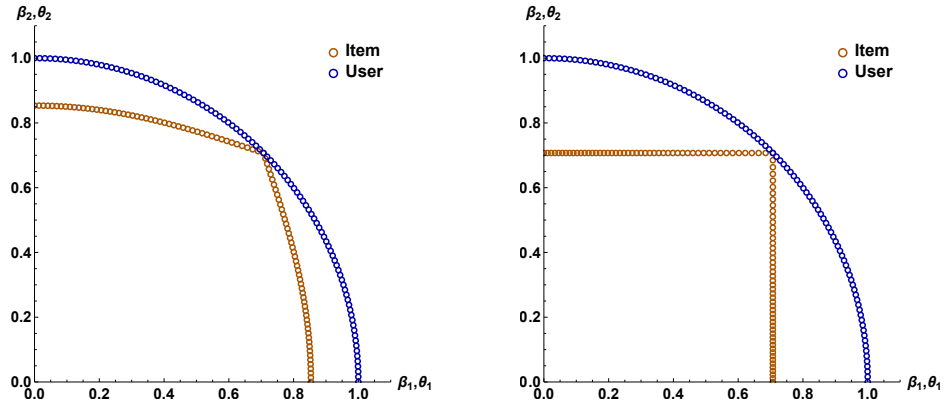
Figure 1: Distribution of user preferences

nates (i.e., $v_{i1} = t_{i1}$ and $v_{i2} = t_{i2}$) to any other product. (This follows from the fact that indifference curves are straight lines with slope $-\frac{t_{i1}}{t_{i2}}$.) The consumer's utility from more distant products is a decreasing function of the distance between the consumer's location and that of the product. For example, Figure 1 (right) depicts the map of indifference curves of a consumer with $t_1 = t_2$ and shows that in this case the optimal product is the one with $v_1 = v_2$. In this model of pure horizontal differentiation, in the surplus maximizing allocation each product is chosen by exactly one consumer. As a result, all products should have the same market share, equal to $1/J$.

On the other hand, one may derive a model of pure vertical differentiation by assuming that products are located on the square $\left[0, \frac{1}{\sqrt{2}}\right]^2$ (Figure 2, right panel). In this case, all consumers agree on their favorite product, which is the one that lies on the 45° line. This product will be referred to as the central, or dominant, product.

Intermediate cases between pure horizontal and pure vertical product differentiation may be obtained by varying the shape of the product locus smoothly, from a circle to a square. This can be done by means of a parameter, α, that ranges from $\alpha = 0$ (circle, horizontal differentiation) to $\alpha = 1$ (square, product differentiation) (Figure 2 left). To be precise, for each product (v_{j1}, v_{j2}) on the circle, define its radial projection on the square as $v'_j = \frac{1}{\sqrt{2}} \left(\frac{v_{j1}}{\max[v_{j1}, v_{j2}]}, \frac{v_{j2}}{\max[v_{j1}, v_{j2}]} \right)$. By taking the convex combination $\alpha v_j + (1 - \alpha) v'_j$ of the circle and the square, one can then generate all intermediate combinations of vertical and horizontal differentiation.

Note that when $\alpha > 0$, a bunch of consumers (i.e., those close to the 45° line) will prefer



(a) Mixed differentiation (b) Pure vertical differentiation

Figure 2: Item differentiation

the same product, i.e. the central one, whereas consumers closer to the axes will prefer more peripheral products. Products close to, but different from the central one would not be chosen by any consumer. As α increases, the fraction of consumers who prefer the central product gets larger, and so does the fraction of nearby products that are not chosen by any consumer.

In this setting, the RS's problem may be far from trivial, especially if the number of consumers and products is large and α is small. In our baseline setting, we consider $I = 99$ consumers, $J = 99$ products, and in future work we plan to consider the case of more consumers.[14]

When $K = 3$, the consumer locus is a sphere, and when $K > 3$ an hypersphere (or, to be precise, the intersection of the hypersphere with the positive orthant). The product locus may vary from an hypersphere to an hypercube, depending on parameter α defined as above. While our baseline case has $K = 2$, we have also run simulations with $K = 5$ and $K = 10$, as reported below.

[14]The fact that the number of consumers equals the number of products may indeed sound implausible, as there typically are more consumers than products. Keep in mind, however, that we assume that all consumers report their ratings in all periods. In practice, only a fraction of them does. Thus, the model is equivalent to one where there are I consumer types, and there are s consumers of each type, but each consumer has only a probability $\frac{1}{s}$ of reporting his or her rating to the RS in each period.

4.2. *Idiosyncratic taste shocks*

To further complicate the RS's problem, and make it more realistic, we allow for the possibility that individual preferences may be subject to idiosyncratic, time-varying shocks:

$$\tilde{u}_{ij\tau} = \sum_{h=1}^{K} t_{ih} v_{jh} + \varepsilon_{ij\tau}. \tag{9}$$

The shocks $\varepsilon_{ij\tau}$ are i.i.d. and follow a Gumbel distribution[15] centered at 0, so they have a zero mean. We assume that consumers experience, and report to the RS, the perturbed utility $\tilde{u}_{ij\tau}$. The higher the variance of the shocks, the more difficult the RS's estimation problem. In the baseline case, the standard deviation of the shocks is set equal to one tenth of the standard deviation of each user's expected utilities across the 99 products.[16]

4.3. *Prices*

Marginal production costs are normalized to zero. Initially, we assume that products are priced at cost. In this case, all prices vanish, and recommendations and consumption choices are driven entirely by the estimated utilities $\{\hat{r}_{ij}\}$.

When we allow prices to vary, recommendations and consumption choices will depend on the net estimated utilities

$$\hat{r}_{ij} - mp_j. \tag{10}$$

For simplicity, the marginal utility of income m is taken to be constant across consumers and over time.[17]

4.4. *Recommendations and consumption choices*

As noted, we abstract from the possibility that recommendations may be manipulated strategically. The reader may think of the RS as run by a platform whose only objective is to maximize the number of subscribers. If the probability of subscription depends on consumer expected utility, the goal of the RS must simply be to predict consumer ratings as best as it can.

[15]This assumption serves to generate a Logit demand system when prices are endogenized.

[16]In general, the variance of the shocks is a parameter that we control and can be varied.

[17]The case when m_i varies across consumers could be dealt with by using the hybrid RSs mentioned in footnote 6 above. (Note that the price can be regarded as an additional factor that is observed by the RS.)

One consequence of the assumption that RSs are unbiased is that users will trust the algorithms and follow their recommendations faithfully. This simplifies the analysis considerably. In the presence of biases of some sorts, users should be wary of RSs' recommendations and filter them using their own private information.

4.5. *Data generation process*

Time is discrete and periods are indexed by $\tau = 0, 1, 2, ...$ The matrix of observed ratings evolves over time as follows. At $\tau = 0$, $\tilde{\mathbf{R}}_0$ is the empty matrix (that is, a matrix with all entries missing). In each subsequent period, each consumer experiences one product and reports the rating $\tilde{r}_{ij}$ to the RSs. Thus, $\tilde{\mathbf{R}}_{\tau+1}$ is obtained by adding to $\tilde{\mathbf{R}}_\tau$ one and only one entry $\tilde{r}_{ij}$ for each of the I rows (i.e. for each consumer). If $\tilde{r}_{ij}$ is not missing in $\tilde{\mathbf{R}}_\tau$, its value is overwritten. All other elements of each row of $\tilde{\mathbf{R}}_\tau$ are unchanged. Thus, $\tilde{\mathbf{R}}_{\tau+1}$ is basically equal to $\tilde{\mathbf{R}}_\tau$ plus one new entry per consumer, which could be either a previous value overwritten or a brand new one. As time passes, the sparsity of $\tilde{\mathbf{R}}_\tau$ therefore decreases.

In our baseline case, the reported rating $\tilde{r}_{ij}$ is taken to be equal to the utility that i obtains from item j at period τ inclusive of the shock, i.e. $\tilde{u}_{ij\tau}$. As noted, the variance of the shock is a parameter that we can control in a robustness analysis.[18]

In each period τ, the RS uses $\tilde{\mathbf{R}}_\tau$ to produce an estimate $\hat{\mathbf{R}}_\tau$ according to the estimation procedure described in the previous section. Using the estimated matrix $\hat{\mathbf{R}}_\tau$, the RS can then identify, for each consumer i, the product with the highest rating:

$$j^\star(i,\tau) = \max_j \{\hat{r}_{ij\tau}\}. \tag{11}$$

At this point, we distinguish between two possible modes of data generation. In the exploitation model, the RS recommends to each consumer i what it regards as the best product for him or her, i.e. $j^\star(i,\tau)$, and consumers try the recommended product.[19] Thus, the new entries in the matrix $\tilde{\mathbf{R}}_{\tau+1}$ are the ratings $\tilde{r}_{ij^\star} = u_{ij^\star\tau}$. In this case, the data used to feed the RS are generated by the RS itself, allowing for the feedback-loop effect discussed in the introduction.

[18] We have considered the case with no shocks (that is, where users report the unperturbed utility $\sum_{h=1}^{K} t_{ih} v_{jh}$), and the results do not differ significantly from the case in which user report their perturbed utility. In future work, we plan to study the case in which the information used by the RS is coarser. For example, users might report their ratings on a discrete scale. In the limit, all the RS might observe is whether the product has been consumed or not (a binary variable).

[19] In the first period, when the RS has no information at all on user ratings, recommendations are made in a purely randon fashion.

In the exploration model, on the other hand, in each period each consumer randomly draws a product j with equal probability. The consumer tries that product and reports his or her rating $\tilde{r}_{ij} = \tilde{u}_{ij\tau}$. As in the exploitation model, this new rating is added, or overwritten, to the matrix of observed rating $\tilde{\mathbf{R}}_\tau$ to produce $\tilde{\mathbf{R}}_{\tau+1}$. Note that even if in the exploration model new products are tried randomly, we can still identify the product that the RS would recommend were it in exploitation mode. Therefore, we can analyze the choices that would be implied by the knowledge that the RS has acquired so far.

These two models can be somehow combined. For example, one could assume that the RS be in exploitation mode with some probability, and in exploration mode with the complementary probability. In particular, it seem natural to assume that initially the RS explores more frequently and then gradually reverts to exploiting the information acquired in the past. Here, however, we focus on the two polar cases so as to better assess whether and to what extent results are driven by the feedback-loop effect.

4.6. *Benchmarks*

We compare the choices made with the assistance of the RS with two benchmarks. The first one is the individual search benchmark, where each consumer can rely only on his or her own search and past experience. To be precise, this benchmark, which is meant to be a stylized representation of a world without RSs, is obtained as follows. At the beginning of each period τ, each consumer i randomly draws an item j and observes $\tilde{u}_{ij\tau}$. The consumer recollects all past draws and chooses the product with the highest expected utility.[20]

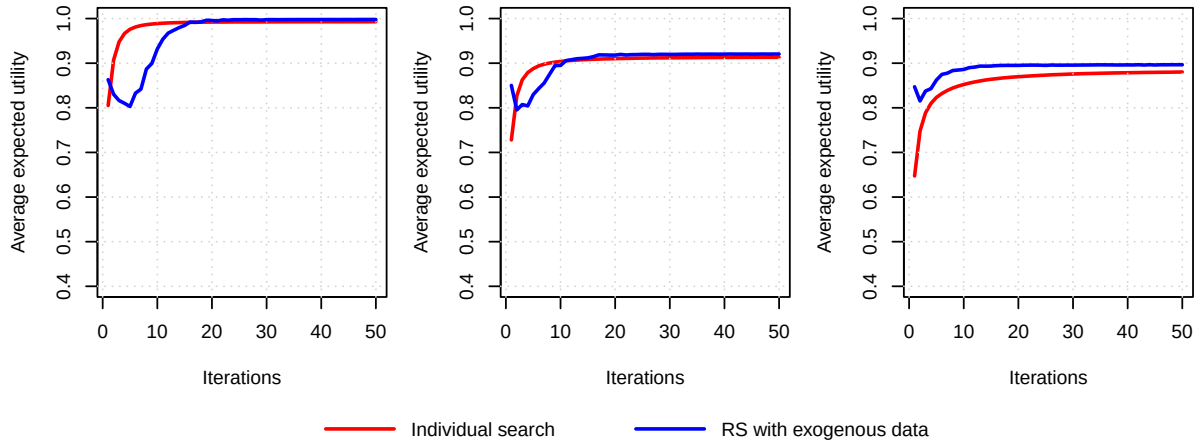
Another benchmark that we shall consider at times is the case of complete information. In this case, consumers are aware of all products, know their expected utilities $\sum_{h=1}^{K} t_{ih} v_{jh}$, and make their choices accordingly.

4.7. *Numerical simulations*

We refer to a combination of preferences, product characteristics, hyper-parameters of the algorithm, and mode of generating the data that are fed into the algorithm as an "experiment." Since the data generation process is random, for each experiment we run $1,000$ independent sessions. Each session consists of $T = 50$ periods indexed by τ. We then average the results across sessions so as to eliminate the noise.

[20] As for RSs, we assume sampling with replacement. When the same item is tried more than once, only the last value of $\tilde{u}_{ij\tau}$ is recalled.

Figure 3: The expected utility generated by the Rs's recommendations (average across 99 consumers and 1,000 sessions).



Horizontal (left), Vertical (right) and Mixed (center) differentiation

While we consider a relatively long time horizon, we are especially interested in outcomes associated with sparse matrices $\tilde{\mathbf{R}}_\tau$.

5. GOODNESS OF FIT

In this section, we analyze the goodness of fit of the RS's estimates, by comparing its recommendations with the choices based on an individual search process. We focus on the exploration model, where the data generation process is exactly the same as in the individual search benchmark and thus the comparison is more direct.

To assess the quality of choices, we use the consumer expected utility at each period τ, averaged across users and across the 1,000 sessions. Note that the maximum utility, which is obtained when a consumer obtains his or her preferred product, is 1 with $\alpha = 0$ but is lower than 1, on average, when $\alpha > 0$.

Figure 3 shows the performance of our RSs relative to the benchmark, for three different values of parameter α: $\alpha = 0$ (pure horizontal differentiation), $\alpha = \frac{1}{2}$ and $\alpha = 1$ (pure vertical differentiation). When the rating matrix is sparse ($\tau \leq 10$), the individual search benchmark does better but is eventually overcome by the RS. As τ increases, the RS approaches the full information benchmark.

Even if RSs eventually perform well, their "cold start" is noteworthy. It may be due to the mechanism of regularization described above, but other explanations are also possible. In any case, after 20 periods or more RSs improve the quality of choice for all values of α: the more so, the higher α.

Relative to the individual search benchmark, RSs perform better, the higher α (that is, the more prevalent is vertical product differentiation). This is intuitive as RSs exploit the similarities across consumers, which are strong when $\alpha = 1$ but do not exist when $\alpha = 0$.

6. MARKET CONCENTRATION

Turning to the more substantive issues addressed in the paper, in this section we analyze the level of market concentration produced by RSs' recommendations. Starting from the exploitation model, we document a significant increase in market concentration relative to the individual search benchmark. This confirms the finding of the computer science literature reviewed above, in an environment that is consistent with standard economic theory, and against a sensible economic benchmark.

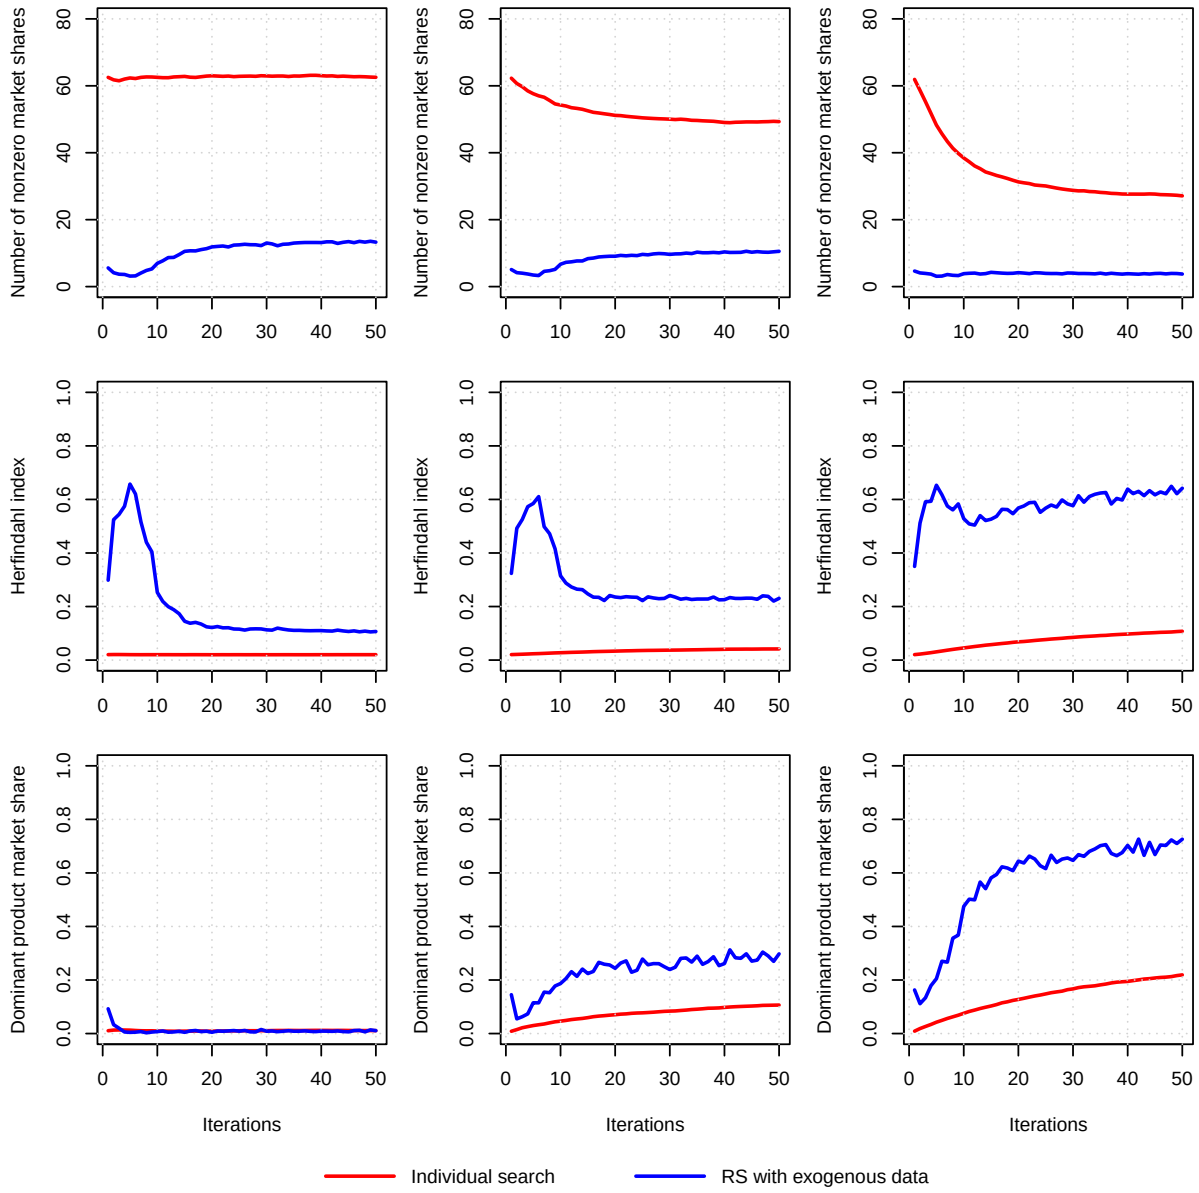
Figure **??** shows, for different values of α, three different measures of market concentration: the number of products consumed by at least one consumer (i.e., the products with a positive market share), the Hirsh-Herfindahl Index (HHI) (a standard measure of market concentration in industrial organization),[21] and the market share of the central, or dominant product. The first variable is a measure of diversity, and hence an inverse measure of concentration, the last two are direct measures. By any measure, the increase in concentration is apparent.

In particular, our first index of concentration captures the diversity of consumer choice. For perfectly informed consumers, the number of products with positive market share would be 99 for $\alpha = 0$ and 1 for $\alpha = 1$. Naturally, the RS lies below the theoretical benchmark for $\alpha = 0$ and above it for $\alpha = 1$, but in any case it lies significantly below the individual search benchmark. In other words, RSs concentrate their recommendations on a much smaller set of products than are consumed in the individual search benchmark. With pure horizontal differentiation, for instance, with RSs consumers end up consuming about 10 products, whereas under individual search they would consume six times as many.[22]

On the other hand, consider the market share of the central product. This product plays no special role when $\alpha = 0$ but is the one that would be chosen, under complete information, by the largest number of consumers when $\alpha > 0$, and by all consumers when $\alpha = 1$. When $\alpha = 0$, neither RSs nor the individual search process attach any prominence to the central product. In other words, the increase in concentration caused by RS does not benefit the central product, but products chosen randomly. With pure vertical differentiation,

[21] A similar pattern is exhibited by the Gini index, another standard measure of concentration.

[22] The number of products consumed in the individual search benchmark does not tend to increase over time because of the idiosyncratic shocks that confound consumers' perception of utility.



Horizontal (left), Mixed (center) and Vertical (right) differentiation

Figure 4: Measures of market concentration (average across 1,000 sessions) for the baseline model. The red curves represent the individual search benchmark, the green curves the case with RSs in the exploitation model (endogenous data).

however, RSs eventually give the dominant product a market share of about 80%, whereas under individual search it would be slightly above 20%. RSs, that is to say, allow consumers to identify with much greater precision a superior product when it exists.

Finally, the HHI is much larger with the RS than in the individual search benchmark, for all values of α and both in the short and the long run. With pure horizontal differentiation, for instance, the HHI is around 2,000 even after 50 iterations, when the matrix of observed rating is filled for more than 30%. This value contrasts with a complete information benchmark of 99, and an individual search benchmark of less than 200.

Summarizing, RSs raise concentration significantly compared to the individual search benchmark. The same pattern is observed with a greater number of factors (i.e., $K = 5$ or $K = 10$).

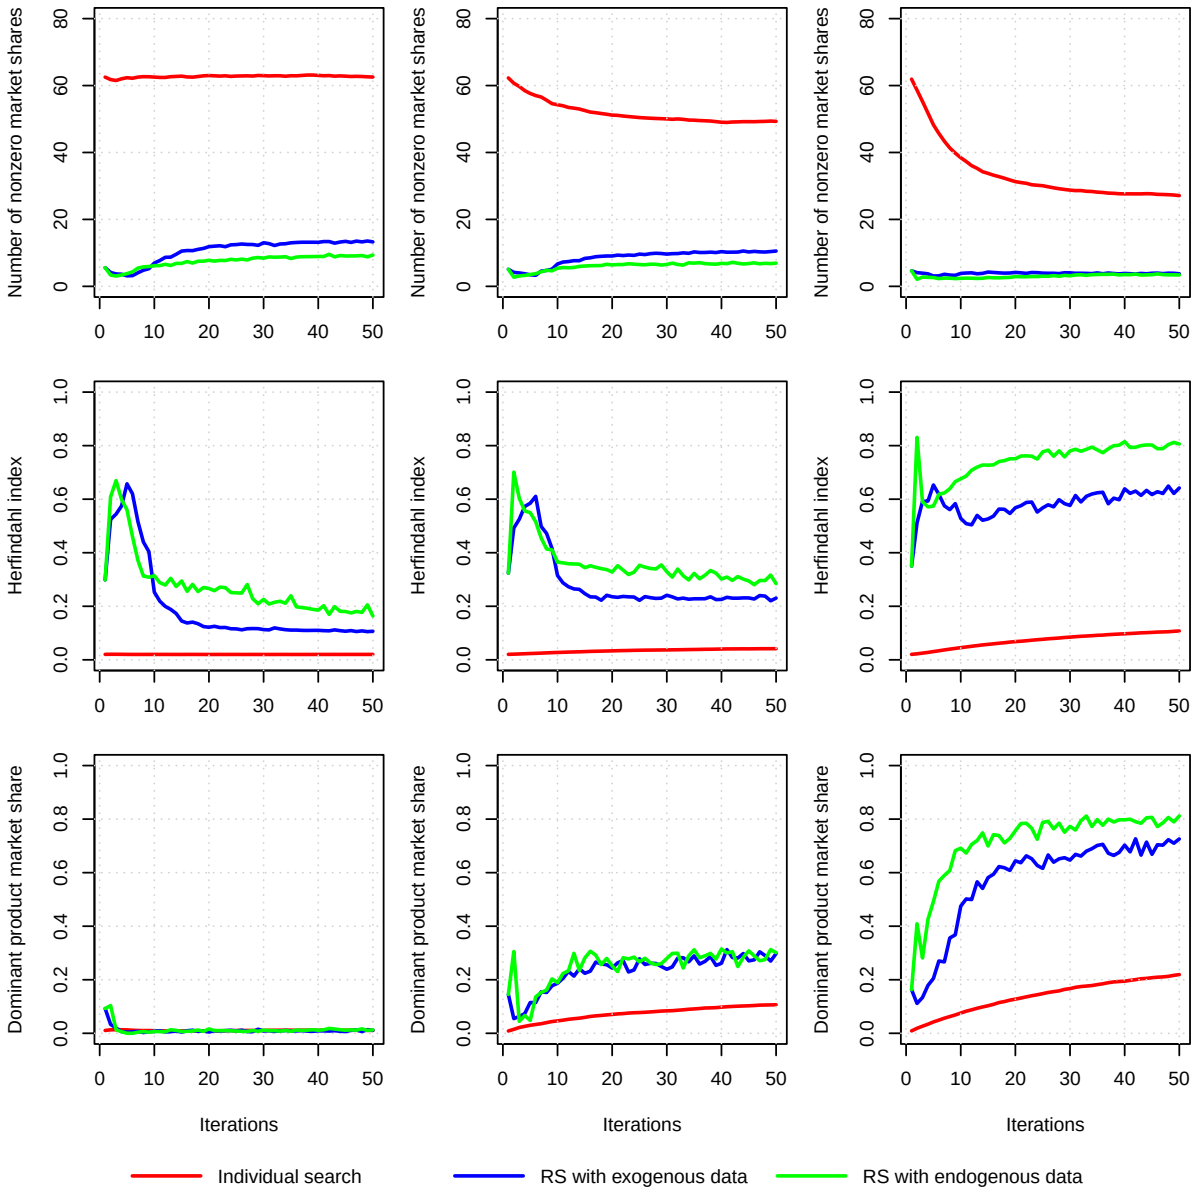
7. CAUSES OF HIGHER CONCENTRATION

The computer science literature has argued that the source of the increase in market concentration is the feedback-loop channel, i.e., the fact that the algorithms are fed with the ratings of the products that they themselves have recommended in the past. If this explanation were correct, the increase in concentration should disappear when the data generation process is purely random, as in our exploration model, because the feedback-loop effect would then be muted.

Figure 5 is similar to Figure 4 but it plots our measures of market concentration also for the case of exogenous data (the exploration model). It appears that market concentration increases, with respect to the individual search benchmark, even if the data used by the RS is not generated by the RS itself.

The endogeneity of data does contribute to the increase in concentration, but the effect is relatively small. For example, with pure horizontal differentiation, the RS recommends eventually around 15 different products in the exploration model, agains the 10 products recommended in the exploitation model. In the individual search benchmark, however, more than 60 products would be consumed. With pure vertical differentiation, on the other hand, the market share of the dominant product increases from about 75% with exogenous data to just above 80% with endogenous data. By comparison, the share would be around 20% under individual search.

In sum, the endogeneity of the data used to feed the RSs is not the only reason why RSs lead to more market concentration, and not even the most important one. This conclusion is important because the feedback-loop narrative suggests that RSs may create market power where there would be none, or they may artificially amplify small initial competitive



Horizontal (left), Mixed (center) and Vertical (right) differentiation

Figure 5: Measures of market concentration (average across 1,000 sessions) for the baseline model. The red curves represent the individual search benchmark, the green curves the case with RSs in the exploitation model (endogenous data), and the blue curves the case with RSs in the exploration model (exogenous data).

advantages, possibly leading to an entrenchment of monopoly. The fact that the feedback-loop effect in itself has only a modest impact on market concentration suggests a more benign interpretation of the competitive effects of RSs. We now turn to a more in depth analysis of such effects.

8. CONSEQUENCES OF HIGHER CONCENTRATION

As noted in the introduction, an increase in the level of market concentration may be a symptom of either stronger or weaker competition. In other words, more concentration may be a signal that fewer firms possess more market power, or that the process of market selection, whereby the more efficient firms get larger and the less efficient ones get smaller, has become more effective. Depending on which explanation applies, the welfare and policy implications can be quite opposite.

To shed light on this important issue, we now analyze how RSs affect producers' incentives to better serve their customers. We focus on firms' pricing decisions, but the same insights apply also to quality choice.[23]

Assume that each product is supplied by a separate, profit-maximizing firm. For simplicity, assume that all consumers have the same marginal utility of income m (normalized to one), so consumer i' net utility from consuming product j is

$$u_{ij\tau} - p_{j\tau}.$$

The RS observes prices and estimates consumer utility from the different products as in the previous sections. Now, however, the recommended product is

$$j^\star(i,\tau) = \max_j \{\hat{r}_{ij\tau} - p_{j\tau}\}. \tag{12}$$

The same changes apply to the individual search benchmark.

When setting prices, firms anticipate the demand implied by the RS's recommendation, or by the individual search process. In the latter case, the system of demand functions is obtained as follows. In period τ, each consumer has a consideration set $C_{i\tau}$ that comprises all products that he or she tried in the past. We assume that consumers observe the prices of all products in their consideration sets and assess each product according to their rating $\hat{r}_{ij\tau}$ (which coincides with the latest realized utility $\tilde{u}_{ij\tau}$ for that product). Thus, the net utility expected by consumer i for product j is $\hat{r}_{ij\tau} - p_{j\tau}$.

[23]This is especially relevant when final consumers are not charged separately for each product they consume.

The demand functions implied by these net expected utilities would be stepwise decreasing. We smooth these demand functions out by assuming that firms anticipate the presence of idiosyncratic shocks as in (9). As a result, firms anticipate that the probability that consumer i will purchase product $j \in C_{i\tau}$ in period τ is

$$q_{ijt} = \frac{e^{\left(\frac{\hat{r}_{ij\tau}-p_{j\tau}}{\sigma}\right)}}{\sum_{j\in C_{i\tau}} e^{\left(\frac{\hat{r}_{ij\tau}-p_{j\tau}}{\sigma}\right)} + 1}. \tag{13}$$

For products $j \notin C_{i\tau}$, the probability is zero. This formulation allows for the possibility that the consumer may choose not to buy – an outside option with net utility of zero.

In the presence of RSs, the system of demand function is obtained in a similar way, with two differences. First, the considerations sets coincide with the set of all products. Second, the estimated utilities $\hat{r}_{ij\tau}$ are those produced by the RS. Strictly speaking, if firms believe that consumers will follow the RS's recommendations mechanically, demand should not be subject to idiosyncratic shocks. However, we sill use the Logit formulation (13) to obtain smooth demand functions.[24]

In all cases, the aggregate demand function of product j is obtained by summing the probabilities of selling to all I consumers. Given the system of demand functions, we finally numerically calculate the static Nash equilibrium prices $\mathbf{p}^{\star}_{\tau}$.[25,26]

Clearly, these prices are an inverse index of the intensity of competition. Therefore, by looking at the impact of RSs on prices one can tell whether the source of the increase in market concentration is stronger or weaker competition.

Figure 6 represents the average prices obtained in the presence of RSs, both with exogenous and endogenous data, and compares it to those arising in the individual search benchmark.[27] Prices are generally very low, reflecting the fact that the economic environment that we are studying is quite competitive, especially when α is low (that is, products are mainly horizontally differentiated).

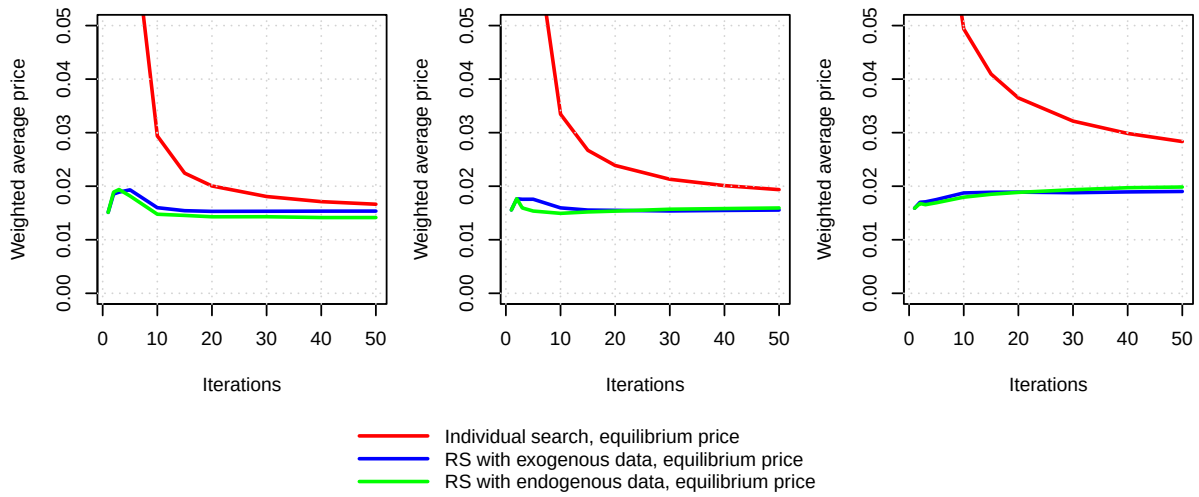
[24]The interpretation is that firms know $\hat{\mathbf{R}}_{\tau}$ and believe that there are many incarnations of user i, each one associated with different realizations of an idiosyncratic shock $\tilde{\varepsilon}_{ij\tau}$, which follows a Gumbel distribution.

[25]We assume that in each period τ, firms maximize their profits statically. Thus, we do not consider the possibility that firms may try to manipulate future recommendations, for example by inflating consumption with low current prices.

[26]To this end, we use the Iterated-Best-Response algorithm (Gauss-Seidel): see Davis and Garces (2009) pp: 417-18 for details. Ifn the few cases where the procedure converges to limit cycles (which could be interpreted as Edgeworth cycles) we take the mean of the firm's prices in the cycle. Cycles however are rare and only occur in the initial periods, i.e., with very sparse rating matrices.

[27]We use a weighted mean, with weights equal to the equilibrium market share of each product.

Figure 6: Average weighted price in the presence of RSs, both with exogenous data (blue curves) and endogenous data (green curves). The individual search benchmark is depicted in red.



Horizontal (left), Vertical (right) and Mixed (center) differentiation

That said, it appears that with RSs, prices are lower than in the individual search benchmark. This shows that RSs actually intensify competition.

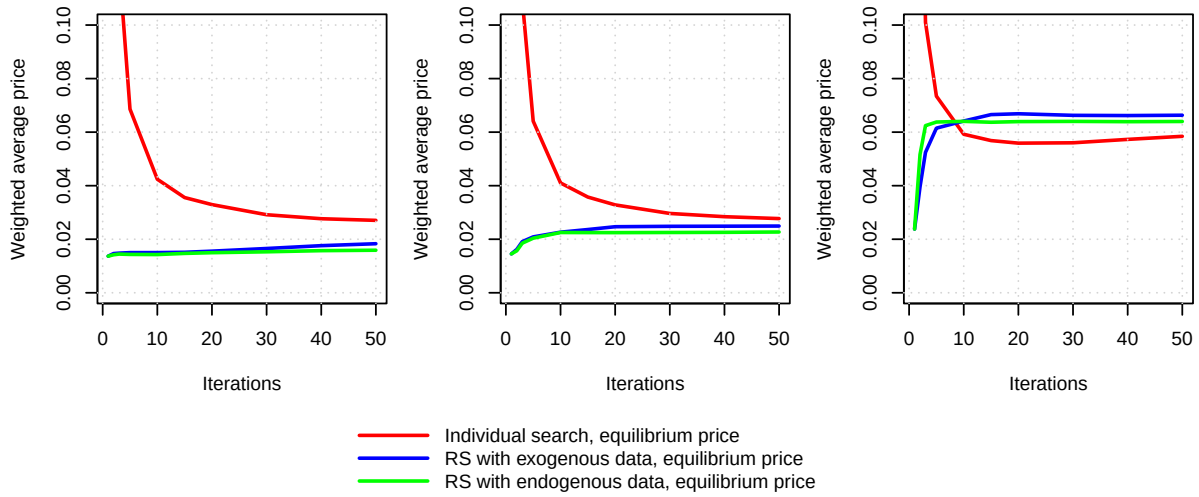
The effect of RSs on prices is stronger, the sparser the rating matrices. This is due to the fact that in the individual search benchmark prices tend to decrease over time, whereas with RSs they are much more stable.

The effect of RSs on prices is stronger, the higher α. Even with pure vertical differentiation ($\alpha = 1$), with RSs the dominant product is disciplined by nearby products, which are potentially pretty good substitutes. In the individual search benchmark, in contrast, consumers may not be aware of these close substitutes and this gives the dominant product more market power.

In terms of prices, the differences between the exploitation and exploration models seem very limited. For the baseline case, the exploitation model leads to lower prices when product differentiation is mainly horizontal, and the matrix of observed rating is rather sparse. When $K = 5$, however, these patterns seem reversed (see Figure 7). In any case, it seems clear that the impact of RSs on the intensity of competition is not due to the feedback-loop effect.

Figure 7 shows that the impact of RSs on the intensity of competition when $K = 5$ is similar to the baseline case. One notable difference is that when $\alpha = 1$ prices can become now higher with RSs than in the individual search benchmark (when sparsity reduces). However, this is entirely due to a composition effect. Prices are in fact higher

Figure 7: Average weighted price ($K = 5$)



Horizontal (left), Vertical (right) and Mixed (center) differentiation

in the individual search benchmark than with RSs, both for the dominant product and for the dominated products. However, the prices of dominated products are much lower than that of the dominant product, and in the individual search benchmark dominated products command a higher market share than with RSs. Another difference is that the price-reducing effect of RSs is more persistent than in the case $K = 2$.

Turning to the case $K = 2$, it may be interesting to ask what would be the level of market concentration accounting for the endogenous prices $\mathbf{p}_\tau^\star$. To answer this question, we have calculated the market shares entailed by the endogenous prices. Figure 8 represent the corresponding HHI index. Relative to the case of zero prices, market concentration decreases. This is intuitive, as products that are favored by the RS will carry higher prices and this tends to reduce their market share.

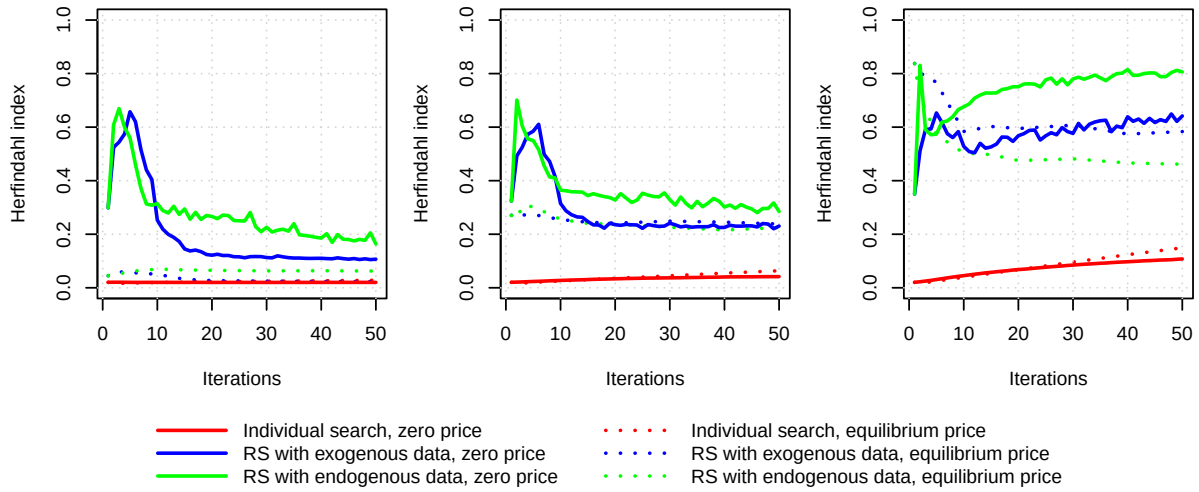
Summarizing, this section has shown that RSs have a pro-competitive, price-reducing effect. The increase in market concentration generated by RSs is therefore a signal of more intense competition.

9. MECHANISM

In this section, we speculate on potential drivers of recommendation bias by looking at how the coefficients' estimates $\hat{\mathbf{T}}$ and $\hat{\mathbf{V}}$ compare with the data generating process. Concentration obviously reflects a systematic failure to recover the true parameters due to small and/or non random sampling.

To provide geometric intuition, figure 9 offers a handpicked example depicting one set of

Figure 8: The HHI with the market shares entailed by endogenous prices.



Horizontal (left), Vertical (right) and Mixed (center) differentiation

estimates for both consumers' and products' latent vectors $\{(\hat{t}_{i1}, \hat{t}_{i2})\}_{i=1}^{I}$ and $\{(\hat{v}_{i1}, \hat{v}_{i2})\}_{j=1}^{J}$ drawn for one particular session in period $\tau = 15$ under pure horizontal differentiation (and nil prices). $\hat{t}_i$ are basically 'virtual' consumers that are directed towards the 'virtual' products $\hat{v}_j$ that maximize their utility. The 99 green dots laying on the circle with unit radius represent the true location of consumers and products. The blue empty circles represent the 99 $\hat{t}_i$ estimates obtained by using the estimation procedure described in section 3.3 with input $\mathbf{R}_{\tau=14}$. The solid black and solid red disks depict the subset of the J estimates (one for each product) that are not dominated (a product is dominated if it exists another product that has higher estimated coordinates).[28] Red disks represent products which are *not* recommended to any consumer despite not being dominated. The larger orange disks around the solid black products have a radius proportional to the (virtual) products' market shares. The dashed lines stemming from the origin represent indifference loci. Consumers on these loci (if any) would be indifferent between different virtual products. Consumers lying between two loci, strictly prefer one particular product over the others. For example, the three virtual consumers in the bottom right 'slice', strictly prefer the product represented by the black dot at the lowest bottom right position (the market share is small and not depicted).

We distinguish at least three economically meaningful sources of failure. First, the algorithms may fail to identify consumers' tastes over factors by placing them 'too close' to a particular product. Geometrically this means that the estimated coefficients in the (t_1, t_2) space are bunched together. In the figure the product with the largest market share is

[28]Dominated products are never recommended as they cannot possibly rank first.

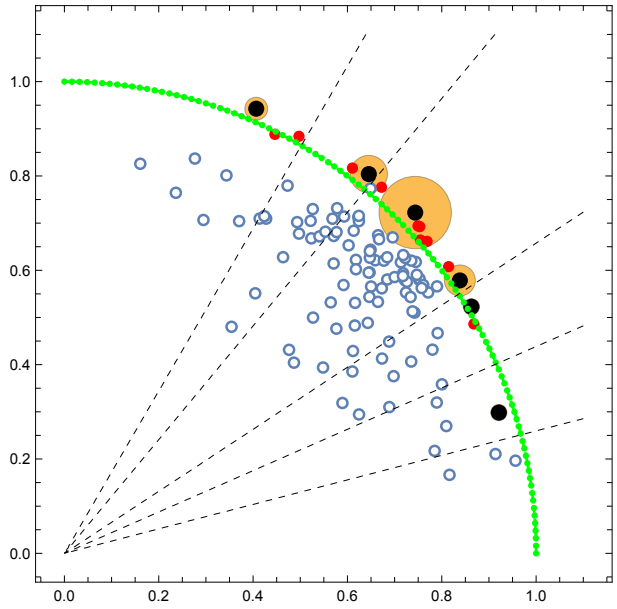


Figure 9: $\hat{t}$, $\hat{v}$ and market share of non-dominated products at $\tau = 15$

strictly preferred by the large cloud of virtual consumers in the middle slice.

Even if the algorithm were to correctly estimate $\hat{\mathbf{T}} = \mathbf{T}$ thus placing virtual consumers on the unit circle, concentration would still occur if it were to systematically fail to estimate product characteristics $\hat{\mathbf{V}}$. This could occur in one of two ways: first, product placed outside of the circle (in the north east direction) would automatically attract more consumers (geometrically they would command larger 'slices'). These are products that would be incorrectly assigned a 'quality' advantage over rival products. Finally, even if all products were to lay on the circle so that no product would command a larger 'quality', if some products were more 'isolated' then others, this would lead to a larger market share.

10. CONCLUSIONS

In future work, we plan to conduct and extensive robustness analysis. Among the extensions that we plan to consider are:

- the case where the number of consumers and products are different;
- the case of misspecified algorithms (for example, the algorithms may use an exogenously given number of factors that is different from the true one, or the number of relevant factors could be estimated jointly with the $\mathbf{T}$ and $\mathbf{V}$ matrices, with the possibility of estimation errors);
- different models of vertical/horizontal product differentiation;
- different amounts of ratings for different users, which could allow to study the im-

plication of users' data-portability.

REFERENCES

Abdollahpouri H. and M. Mansoury, (2020), Multi-sided Exposure Bias in Recommendation, International Workshop on Industrial Recommendation Systems with ACM KDD 2020.

Abdollahpouri H., M. Mansoury, R. Burke, B. Mobasher, E. C. Malthouse, (2021) User-centered Evaluation of Popularity Bias in Recommender Systems, arXiv:2103.06364.

Anderson C. (2008) The Long Tail: Why The Future Of Business Is Selling Less Of More (Hyperion Books, New York).

Bar-Isaac, H., G. Caruana, and V. Cunat (2012) Search, design, and market structure, American Economic Review 102 (2), 1140-60.

Brynjolfsson E, Hu YJ, Simester D (2011) Goodbye Pareto principle, hello long tail: The effect of search costs on the concentration of product sales. Management Sci. 57(8):1373-1386.

Calvano E. and B. Jullien (2018) Can we trust the algorithms that recommend products online? A theory of biased advice, wp.

Chaney A. JB. B.M. Stewart, and B.E. Engelhardt (2018) How algorithmic confounding in recommendation systems increases homogeneity and decreases utility, in Proceedings of the 12th ACM Conference on Recommender Systems. ACM, 224-232.

Chin W.S., B.W. Yuan, M.Y Yang, Y. Zhuang, Y.C. Juan and C.J. Lin (2016) LIBMF: A Library for Parallel Matrix Factorization in Shared-memory Systems, Journal of Machine Learning Research 17: 1-5

Ciampaglia G.L. and A. Nematzadeh, F. Menczer, and A. Flammini, (2018) How algorithmic popularity bias hinders or promotes quality. Nature, Scientific reports 8, 1:17.

Davis, P., and E. Garces, (2009) Quantitative Techniques for Competition and Antitrust Analysis. Princeton: Princeton University.

Fleder D. and K. Hosanagar, (2009) Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity, Management Science, vol. 55, no. 5, pp. 697-712.

Goldfarb A. and C. Tucker (2019) Digital Economics, Journal of Economic Literature, 57(1), 3-43.

Hagiu, A., and B. Jullien (2011) Why do intermediaries divert search?, The RAND Journal of Economics, 42(2), 337-362.

Jannach D. and M. Jugovac (2019) Measuring the business value of recommender systems, ACM Transactions on Management Information Systems (TMIS), 10(4):1-23.

Lee D. and K. Hosanagar, (2019) How do recommender systems affect sales diversity? a cross-category investigation via randomized field experiment, Information Systems Research, vol. 30, no. 1, pp. 239-259.

Rosen, S. (1981). The Economics of Superstars, American Economic Review. 71: 845-858.

Teh, T.-H., and J. Wright (2020) Intermediation and steering: Competition in prices and commissions, American Economic Journal: Microeconomics.

Zhang, J., Adomavicius, G., Gupta, A. and Ketter, W. (2020). Consumption and Performance: Understanding Longitudinal Dynamics of Recommender Systems via an Agent-Based Simulation Framework. Information Systems Research, 31(1), pp.76-101.