Deconstructing the Filter Bubble:
User Decision-Making and Recommender Systems

GUY ARIDOR, Columbia University

DUARTE GONÇALVES, Columbia University

SHAN SIKDAR, Everquote

We study a model of user decision-making in the context of recommender systems via numerical simulation. Our model provides an explanation for the findings of Nguyen, et. al (2014), where, in environments where recommender systems are typically deployed, users consume increasingly similar items over time even without recommendation. We find that recommendation alleviates these natural filter-bubble effects, but that it also leads to an increase in homogeneity across users, resulting in a trade-off between homogenizing across-user consumption and diversifying within-user consumption. Finally, we discuss how our model highlights the importance of collecting data on user beliefs and their evolution over time both to design better recommendations and to further understand their impact.

1 INTRODUCTION

Recommender systems (RS) have become critical for assisting users in navigating the large choice sets that they face on many online platforms. For instance, users have to choose from thousands of movies on Netflix, millions of products on Amazon, and billions of videos on YouTube. Users in many cases are not aware of most items, let alone have full information about their preferences over them. To make matters worse, the items in these contexts are usually experience goods whose true value for users can only be learned after consumption.

RS have driven a significant fraction of consumer choice on these platforms with 75% of movies watched on Netflix and 35% of page-views on Amazon coming from recommendations.[1] While there are many positive effects from these systems, there is an increasing worry about their unintended side-effects. There have been claims that personalized RS lead users into *filter bubbles* where they effectively get isolated from a diversity of viewpoints or content [25], and that personalized RS may also lead users to become increasingly homogenized at the same time [6, 15].

Understanding how RS influence user behavior is important not only for characterizing the broader consequences of such systems but also for guiding their design. In this paper, we develop a theoretical model of user decision-making in contexts where RS are traditionally deployed. We utilize previous empirical studies that characterize how RS influence user choice as a benchmark and our theoretical model provides an intuitive mechanism that can explain these empirical results. The key insight of our model is that user *beliefs* drive the consumption choices of users and that recommendations

---

[1]MacKenzie et al. (2013, Oct.), How retailers can keep up with consumers. https://www.mckinsey.com/industries/retail/our-insights/how-retailers-can-keep-up-with-consumers. Retrieved on October 3, 2019.

provide them with information that leads them to update their beliefs and alter their choices. A crucial component of our model is that users' beliefs about items are driven not only by recommendations, but also from their previous experiences with similar items. We use these insights to provide guidance for RS design, highlighting that understanding users' beliefs about the quality of the available items is essential to design recommendations and evaluate their impact.

Our Model. We analyze a model of user choice with four central components.

The first component of our model is that users sequentially consume items and face large choice sets. In our setting of interest, users are long-lived, but they only consume a small fraction of this choice set over their lifetime. This is traditionally the case on online platforms that have thousands or millions of options for users.

The second component is that, prior to consuming them, users are *uncertain* about how much they value the different items. This is motivated both by the fact that recommender systems are traditionally deployed in contexts with experience goods, whose true value can only be learned after consumption, and the fact that such uncertainty is why RS exist in the first place. Thus, users face a sequential decision-making problem under uncertainty.

The third, and most crucial, element is that consumption of an item reveals information that changes user beliefs about their valuation of similar items. Unlike in standard sequential decision-making problems, once an item is consumed all uncertainty about its valuation is resolved and provides information that enables users to update their beliefs about similar items. This exploits the fact that the valuations of similar items are correlated which assists users in navigating the vast product space. The idea that users make similarity-based assessments to guide their choice has grounding in empirical evidence on how users navigate large choice sets [27].

Finally, in our model recommendation provides users with information about the true valuations. We model the realized valuations as being a weighted sum of a common-value and an idiosyncratic component. This formulation gives a stylized notion of predictability of user preferences where the idiosyncratic component is inherently unpredictable given other users' preferences and the common-value component is what the recommender can learn from previous users' data. We suppose that the recommender knows the common-value component for each item and combines it with users' beliefs over the product space when designing personalized recommendation.

Our Contributions. We provide a clear mechanism that explains the empirical results in [24] who show that, in the context of movie consumption, user behavior is consistent with filter-bubble effects even without recommendation and that recommendation leads to users being less likely to fall into such filter bubbles. In this context, filter-bubble effects are defined as users consuming items in an increasingly narrow portion of the product space over time. The simple and intuitive driving force of this is that preferences for similar items are correlated, which implies that when an item is consumed and the user learns its value, it provides information about similar items. Crucially, this not only impacts the underlying belief about the expected value of similar items, but also how uncertain the user is about their valuation of them. Consequently, this learning spillover induces users to consume items similar to those they consumed before that had high realized value, leading to an increasing narrowing of consumption towards these regions of the product space. This effect is further amplified when users are *risk-averse*, a concept from decision theory where all else being equal, users have a preference for items with lower uncertainty to those with higher uncertainty. However, by providing information to users, recommendation leads users to be more likely to explore other portions of the product space, limiting the filter bubble effect.

We find that, while recommendation leads a single user to be more likely to explore diverse portions of the product space, it also coordinates consumption choices across users. This leads to an increase in homogeneity across users, resulting in a trade-off between homogenizing across-user consumption and diversifying within-user consumption. We

explore the relationship between the overall diversity of consumed items and user welfare and find that more diverse sets of consumed items do not always correspond to higher user welfare.

Lastly, we discuss how our model and findings can be used to inform the design and evaluation of RS as well as the data that is traditionally collected for them. This highlights the importance of user beliefs in determining user consumption choices and how both recommendation and informational spillovers determine how these beliefs change over time. By collecting information on user beliefs, RS designers can understand what items a user would consume *without* recommendation and then predict how providing information to the user would change her beliefs and resulting consumption decisions. Thus, our evaluation measure determines the value of a recommendation based on the marginal welfare gain associated with providing a user with a recommendation over what the user would do without it. We discuss how this provides an additional rationale as to why "accurate" recommendations are not always good recommendations.

2 RELATED WORK.

The first set of related works studies the extent and implications of filter bubbles. [25] first informally described the idea of the "filter bubble" which is that online personalization services would lead users down paths of increasingly narrower content so that they would effectively be isolated from a diversity of viewpoints or content. Following this, a number of empirical studies in various disciplines, have since studied the extent to which this phenomenon exists in a wide range of contexts [9, 15, 23, 24]. The most relevant to our study is [24] who study whether this effect exists in the context of movie consumption. They find that even users whose consumption choices are not guided by recommendations exhibit behavior consistent with "filter bubbles" and that RS can actually increase the diversity of the content that users consume. To our knowledge there are no theoretical models that rationalize these empirical findings and we provide a theoretical framework through which to view this problem. Moreover, we provide a clear mechanism that drives such effects and how recommendation interacts with them.

Another set of papers has examined whether RS can lead users to become increasingly homogenized. [4, 29] show that incorporating content popularity into RS can lead to increased user homogenization. [6] shows how user homogenization may arise from training RS on data from users exposed to algorithmic recommendations. [10] show that homogenization can increase due to a popularity recommendation bias that arises from lack of information about items with limited consumption histories. We show similar results as previous work where RS lead to increased user homogenization. However, the mechanisms behind this differ from existing work as homogenization arises due to the fact that recommendation leads users to coordinate their consumption decisions in certain portions of the product space.

A third strand in the literature studies the impact of human decision-making on the design and evaluation of RS. [7] surveys the literature on the relationship between human decision making and RS. The closest set of papers pointed out in this survey are those related to preference construction [2, 19] whereby users develop preferences over time through the context of a decision process. We point out that the true underlying preferences of users may be stable over time, but, due to the nature of items in contexts where RS are deployed, they have incomplete information of their valuations and both consumption and recommendation provide them with information to reduce their uncertainty. Thus, the primary insight of our paper is that user beliefs and how users update their beliefs about similar items after consumption are important and previously unconsidered elements of human decision making that are critical for understanding the design and consequences of RS. Within this literature, [5, 8, 26] focus on "user-centric" approaches to recommendation whereby user evaluation of the usefulness of recommendation is a key evaluation measure. Our evaluation measure is similar, but, unlike previous approaches, emphasizes the importance of user beliefs. Finally,

[14] considers a similar model as ours where users engage in "spatial learning" and exploit the correlation of their preferences in the environment, but consider it in the context of search for a single item.

3 OUR MODEL AND PRELIMINARIES

3.1 Preliminaries on Expected Utility Theory

For every item n in the product space $\mathcal{J}$, we assume that each user i assigns a monetary equivalent $x_{i,n} \in \mathbb{R}$ to the experience of consuming it. Each user can value the same item differently. However, we assume that users have the same utility over money, given by a utility function $u : \mathbb{R} \to \mathbb{R}$, strictly increasing and continuous. So, ex-post, the value of item n for user i is given by $u(x_{i,n})$. Before consuming the item, the user does not know exactly how she will value it. In particular, even users that will end up having the same ex-post valuation of item n may differ in their ex-ante valuation because they hold different beliefs about it. We denote by p_i the beliefs user i has about how she will value each of the items in the product space. Note that this implies that consuming item n is the same as taking a gamble. Each user evaluates the item according to the expected utility associated with the item, i.e. $U_i(n) = \mathbb{E}_{p_i}[u(x_n)]$.

Risk aversion captures how different users react to the risk associated to a particular consumption opportunity. It is formalized as follows: a given gamble x takes real values and follows distribution p. Then, for every gamble x, there is a certain amount of money that makes the user indifferent between taking the gamble or taking the sure amount of money. This sure amount of money is called the *certainty equivalent* of gamble x and is denoted as $\delta(x)$. A user i is more risk-averse than another user j if whenever user j prefers a sure thing to the gamble, then user i does too. Therefore, a more risk-averse user is more willing to avoid the risk of taking the gamble. We assume that the utility function takes a flexible functional form $u(x) = 1 - \exp(-\gamma x)$ for $\gamma \neq 0$ and $u(x) = x$ for $\gamma \to 0$ – known as constant absolute risk-aversion preferences (from hereon CARA). Higher γ implies higher risk-aversion, with $\gamma \to 0$ corresponding to the risk-neutral case and $\gamma > 0$ to the risk-averse one. Our formulations here follow standard economic consumer theory (see [21] for a textbook treatment of these topics).

3.2 Model

Users. We consider a set of users I where each user $i \in I$ faces the same finite set of N items $\mathcal{J} = \{0, 1, ..., N - 1\}$. For simplicity, we assume that users only derive pleasure from item $n \in \mathcal{J}$ the first time they consume it.

We denote by $x_{i,n}$ user i's realized value from consuming item n. In particular, we consider that the realized value derived from a given item can be decomposed in the following manner: $x_{i,n} = v_{i,n} + \beta v_n$, where $v_{i,n}$ denotes an idiosyncratic component – i.e. user i's idiosyncratic taste for item n – and v_n, a common-value component. One can interpret v_n as a measure of how much item n is valued in society in general and, in a sense, $v_{i,n}$ denotes how i diverges from this overall ranking. The scalar $\beta \in \mathbb{R}_+$ denotes the degree to which valuations are idiosyncratic to each user or common across users. If $\beta = 0$, it is impossible to generate meaningful predictions of any one's individual preferences based on others, while if β is large, every individual has similar preferences.

Stacking values in vector-form, we get the vector of values associated with each item

$$\left(x_{i,n}\right)_{n \in \mathcal{J}} =: X_i = V_i + \beta V,$$

where $V_i = \left(v_{i,n}\right)_{n \in \mathcal{J}}$ and $V = (v_n)_{n \in \mathcal{J}}$.

User Decision-Making. We assume the user makes T choices and therefore can only consume up to T items, where T is a small fraction of N. This captures the idea that users are faced with an immense choice set, but that ultimately they

end up experiencing (and learning) about just a small fraction of it. For tractability, we impose that users are myopic and every period consume the item that they have not yet tried (n_i^t) that gives them the highest expected utility given the information from past consumption ($C_i^{t-1} = (n_i^1, ..., n_i^{t-1})$) and their initial beliefs.

User Beliefs. We assume that all the true realized values are drawn at $t = 0$. However, users do not know the realized values before consuming an item, but rather have beliefs over them. Formally, user i starts with some beliefs about X_i, namely that the idiosyncratic and common-value parts of the valuations are independent – $V_i \perp\!\!\!\perp V$ – and that each is multivariate normal:

(1) $V_i \sim \mathcal{N}(\overline{V}_i, \Sigma_i)$; and

(2) $V \sim \mathcal{N}(\overline{V}, \Sigma)$ with $\overline{V} = 0$.

We impose the normality assumption for two reasons. The first is that this allows for simple and tractable belief updating. The second is that it allows us to incorporate an easily interpretable correlation structure between the items. The precise formulation of Σ and Σ_i that we consider is defined below when we discuss user learning.

Recalling that V_i represents idiosyncratic deviations from V, we assume that, on the population level, prior beliefs $\overline{V}_i = (\overline{v}_{i,n})_{n \in \mathcal{J}}$ are drawn independently from a jointly normal distribution, where $\overline{v}_{i,n} \sim \mathcal{N}(0, \overline{\sigma}^2)$ are independent and identically distributed. These $\overline{v}_{i,n}$ denote the prior belief that individual i holds about her valuation over item n. As people are exposed to different backgrounds, their beliefs about how much they value a given item also varies and $\overline{v}_{i,n}$ denotes this idiosyncrasy at the level of prior beliefs.

We assume users are expected utility maximizers. User i's certainty equivalent for item n, the sure value that makes user i indifferent between it and consuming the item n, conditional on the consumption history, is given by $\delta_i(n) \mid C_i^{t-1} = \mu_n - \frac{1}{2}\gamma\Sigma_{nn}$, where μ_n and Σ_{nn} are the expected value and variance for item n that the user has given their initial beliefs and consumption history up until time t. Note that this expression is known to be the certainty equivalent for CARA preferences in a Gaussian environment [21]. As it is immediate from this expression, the user assigns greater value to items for which the expected monetary equivalent, μ_n, is higher, but penalizes those about which there is greater uncertainty Σ_{nn}, the more so the greater the user's degree of risk aversion γ.

User Learning. When a user consumes an item n she learns the realized value for that item. We consider the case where learning about the value of item n reveals more about the value associated to items that are closer to it, which captures the idea that trying an item provides more information about similar items than about dissimilar ones.[2] In order to have a well-defined notion of similarity we need to define a distance function between items, which we define as $d(n, m) := \min\{|m - n|, N - |m - n|\}$ where m and n are indices of items in $\mathcal{J}$. This distance function is not intended to model the intricacies of a realistic product space, but instead to provide a stylized product space to help us understand the effects of informational spillovers on user behavior. The basic intuition, in the context of movie consumption, is that a user's valuation of *John Wick* (item n) provides more information about how much she will like *John Wick: Chapter Two* (item m) than *Titanic* (item q) since $d(n, m) < d(n, q)$.

We consider that the entry of n-th row and the (m)-th column of Σ_i is given by $\sigma_i^2 \rho^{d(n,m)}$, and that of Σ is given by $\sigma^2 \rho^{d(n,m)}$. The scalar $\rho \in [0, 1]$ therefore impacts the covariance structure: a higher ρ implies that learning the utility of n is more informative about items nearby and, for $\rho \in (0, 1)$, this effect is decreasing in distance. The particular

---

[2][27] empirically studies how individuals solve sequential decision-making problems under uncertainty in large choice sets in the context of mobile food delivery orders. They find that individuals engage in similarity-based generalizations where learning about the realized value of a particular item provides them with information about similar items. We incorporate this finding into our model in a stylized manner.

distance function that we rely on leads to a simple covariance structure, where the $(n, n + 1)$-th entry in the covariance matrix is $\sigma^2 \rho$, the $(n, n + 2)$-th entry is $\sigma^2 \rho^2$, etc.[3]

The precise updating rule is as follows. Recall that at time t the user's consumption history is given by C_i^t and we denote the utility realizations of these items as c_t. We denote μ_t as the initial mean beliefs the user has over the items in C_i^t and μ_{N-t} as the initial mean beliefs the user has over the remaining $N - t$ items, $\mathcal{J} \setminus C_i^t$. We partition the covariance matrix as follows:

$$\Sigma = \begin{pmatrix} \Sigma_{(N-t,N-t)} & \Sigma_{(N-t,t)} \\ \Sigma_{(t,N-t)} & \Sigma_{(t,t)} \end{pmatrix}.$$

After consuming the items in C_i^t, the resulting beliefs over the remaining items are given by $\mathcal{N}(\bar{\mu}, \bar{\Sigma})$ where $\bar{\mu}$ and $\bar{\Sigma}$ are as follows:

$$\bar{\mu} \mid c_t = \mu_{N-t} + \Sigma_{(N-t,t)} \Sigma_{(t,t)}^{-1} (c_t - \mu_t);$$

$$\bar{\Sigma} \mid c_t = \Sigma_{(N-t,N-t)} - \Sigma_{(N-t,t)} \Sigma_{(t,t)}^{-1} \Sigma_{(t,N-t)}.$$

An Illustrative Example. We illustrate the main intuitions of our model with a simple example. Suppose that there are four items: 0, 1, 2, 3. The items are in different places of the product space, where 0 is close to 1 and 3 but more distant from 2. For the sake of expositional clarity, suppose that the initial mean beliefs are given by $\mu = (\mathbb{E}[x_n])_{n=0}^{3} = (0)_{n=0}^{3}$.

In period 1, every item is ex-ante identical since they have the same mean and variance and so suppose that the user breaks the tie arbitrarily and consumes item 0. The underlying correlation structure implies that upon observing that $x_0 = y$ the user will update beliefs about the remaining three items according to the previously specified updating rule. For concreteness, we suppose that $\sigma = 1$ and $\rho = 0.5$, but the intuitions hold for any value of σ and $\rho > 0$. First, we consider the case when the realization of $y > 0$ and, specifically, $y = 0.5$ – though the general intuitions hold for any $y > 0$. The resulting beliefs after observing y are then as follows:

$$\bar{\mu} = (\mu \mid x_0 = y) = \begin{pmatrix} \mathbb{E}[x_1 \mid x_0 = y] \\ \mathbb{E}[x_2 \mid x_0 = y] \\ \mathbb{E}[x_3 \mid x_0 = y] \end{pmatrix} = \begin{pmatrix} \rho y \\ \rho^2 y \\ \rho y \end{pmatrix} = \begin{pmatrix} \frac{1}{4} \\ \frac{1}{8} \\ \frac{1}{4} \end{pmatrix}, \qquad \bar{\Sigma} = (\Sigma \mid x_0 = y) = \begin{pmatrix} \frac{3}{4} & \frac{3}{8} & 0 \\ \frac{3}{8} & \frac{15}{16} & \frac{3}{8} \\ 0 & \frac{3}{8} & \frac{3}{4} \end{pmatrix}.$$

Thus, upon learning $x_0 = y$, the user updates beliefs about the remaining items. Note that $\mathbb{E}[x_1 \mid x_0 = y] = \mathbb{E}[x_3 \mid x_0 = y] > \mathbb{E}[x_2 \mid x_0 = y]$ since item 0's value is more informative about similar items' values, items 1 and 3, than items further away in the product space such as item 2. Moreover, $\bar{\Sigma}_{11} = \bar{\Sigma}_{33} < \bar{\Sigma}_{22}$ as the uncertainty about items 1 and 3 is further reduced compared to item 2. Thus, since $y > 0$, the user in the next period will consume items nearby to item 0 since, even though initially she believed that all items had the same mean, the spillover from consuming item 0 leads her to believe that items 1 and 3 have higher expected valuations. Since both the mean is higher for these items and the variance is lower, the user will consume items 1 and 3 regardless of her risk aversion level.

Now we consider the case when item 0 ends up having a negative valuation so that $y = -0.5 < 0$. This results in $\mathbb{E}[x_1 \mid x_0 = y] = \mathbb{E}[x_3 \mid x_0 = y] = -\frac{1}{4} < -\frac{1}{8} = \mathbb{E}[x_2 \mid x_0 = y]$ with $\bar{\Sigma}$ remaining the same as when $y = 0.5$. In this case the risk-aversion levels of the user determine the choice in the next period. If the user is risk-neutral ($\gamma = 0$), then she will go across the product space to consume item 2 in the next period since it has a higher expected value. However, if she is sufficiently risk-averse then she may still consume item 1 or 3 since her uncertainty about these items is lower

---

[3]This exponential decay correlation structure can be related to the tenet of case-based similarity of [12] – see [3] for an axiomatization of exponential similarity.

than item 2. In particular, this will happen when

$$\delta(3) = \delta(1) = \rho y - \frac{1}{2}\gamma\bar{\Sigma}_{11} > \rho^2 y - \frac{1}{2}\gamma\bar{\Sigma}_{22} = \delta(2)$$

Given the aforementioned parametrization and $y = -0.5$, the user will consume item 1 or 3 when $\gamma > \frac{4}{3}$ and will consume item 2 when $\gamma < \frac{4}{3}$. Thus if the user is risk averse enough, then she might be willing to trade-off ex-ante lower expected values for lower risk and stick to consuming nearby items just because these items have lower uncertainty.

This example illustrates the main mechanisms that can lead to excessive consumption of similar items. Once the user finds items in the product space with high valuations she will update her beliefs positively about items in this portion of the product space and continue consuming these items regardless of her level of risk aversion. However, this same updating leads to a reduction in uncertainty of these items and so, if she is sufficiently risk-averse, she still may continue consuming items in this portion of the product space, even if she has bad experiences with them, since they are perceived to be less risky.

Recommendation. Our model of recommendation is stylized in order to provide qualitative insights into how recommendation shapes behavior, instead of focusing on the details of how RS are implemented in practice. We model recommendation as giving users information about the valuation of the items.

We will consider three cases. The case of primary interest is *recommendation* where the recommender observes values accrued and knows V but does not know V_i.[4] However, the recommender does know the users' beliefs $\bar{V}_i$. Thus, at any given period, the recommender provides a personalized recommendation that combines the knowledge of the common value component V with the user beliefs $\bar{V}_i$. Knowing the user's beliefs about her valuation of each item become crucial in this case: just providing the user with information about V may change the user's original ranking, but, without considering the user's beliefs, she will not necessarily follow the recommendation.[5]

We further consider two cases that serve mainly as benchmarks. The first is *no recommendation*, where users get no additional information and make consumption choices based on their beliefs and consumption history. This gives us a benchmark as to how users would behave *without* recommendation so that we can analyze what changes with the introduction of recommendation. The second is the *oracle recommendation* where the recommender knows the true realized utility of each item for each user and can therefore recommend the best remaining item in every period. This gives us a full information benchmark, which is the optimal consumption path for a user if all uncertainty about their preferences was resolved. Comparison to the oracle regime benchmark provides an analog to the standard regret measures utilized in the multi-armed bandit literature, which look at the difference in the expected value of the ex-post optimal action and the expected value of actions that were taken.

Simulation Details. We analyze our model using numerical simulation since the sequential decision-making component paired with the rich covariance structure between the items makes it hard to characterize optimal behavior analytically.[6]

We explore how consumption patterns differ as we consider different recommendation regimes and report representative results from our simulations. We run this simulation over 100 populations of users with 100 users per population for each combination of parameters. A given set of parameters and a user are a single data point in our dataset.

---

[4]We do not consider the acquisition of information for the recommender to know V and suppose that she has sufficient data to learn V with arbitrary precision at $t = 0$.
[5]The notion of recommendation that we consider is idealized where the recommendation does the Bayesian updating for users, but the results are equivalent to if the users did the updating themselves.
[6]The Gaussian assumption allows for closed form belief updating which enables us to simulate our model, but does not provide much help in analytical characterizations.

We simulate over risk-aversion parameters $\gamma \in \{0, 0.3, 0.6, 1, 5\}$, standard-deviation $\sigma \in \{0.25, 0.5, 1.0, 2.0, 4.0\}$, correlation parameters $\rho \in \{0, 0.1, 0.3, 0.5, 0.7, 0.9\}$ and degree of idiosyncrasy of individual valuations $\beta \in \{0, 0.4, 0.8, 1, 2, 5\}$. The range of considered parameter values cover the relevant portions of the parameter space in order to provide full insight into the behavior of the model. When we consider results varying a single parameter, we group the results over the other parameters and provide reports varying only the parameter of interest. We report results for a relatively small consumption history $T = 20$ with a product space size $N = 200$.[7]

4 RESULTS

4.1 Local Consumption and Filter Bubbles

We characterize "filter bubble" effects as the degree to which users engage in *local consumption*. We define local consumption in terms of the average consumption distance between the items consumed by the users at time $t - 1$ and t. Thus, in the context of our model, filter-bubble effects arise when the average consumption distance decreases over time and, across regimes, when the levels are lower for a given recommendation regime compared to another.

Our first finding can be summarized as follows:

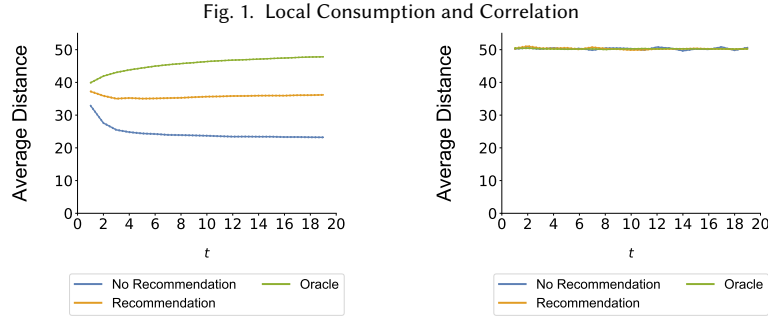FINDING 1. *The impact of recommendation on local consumption:*

(1) *When $\rho = 0$, there is no difference in consumption distance between the three recommendation regimes.*
(2) *When $\rho > 0$, no recommendation induces more local consumption than both recommendation and oracle regimes. This effect is amplified as ρ increases as well as when users are more risk averse (γ increases).*

First, the right panel of Figure 1 shows that, when $\rho = 0$, there is no difference in consumption distance between the three regimes. This is due to the fact that when $\rho = 0$, there is no reason that items that are close in the product space should have similar values and so the optimal consumption path does not depend on the similarity of the items. However this also means that users do not learn anything about neighboring items and so there is limited path-dependence in consumption. Not only is there no difference in the levels between the three regimes, but they all have the same, flat, average consumption distance path. This underscores the fact that if there were no correlation between the realized utilities then there would be no reason for users to consume similar items and thus no narrowing effect, regardless of the information on the true utility of the items that users had.

The left panel of Figure 1 shows that, when $\rho \in (0, 1)$, both recommendation and no recommendation lead to increasingly local consumption compared to the oracle benchmark case. Moreover, the average consumption path between periods is *decreasing* for the no recommendation case whereas it is *increasing* for the oracle case. The recommendation regime decreases the degree of local consumption, but not as much as the oracle benchmark. Due to the correlation of values, the oracle consumption path exploits this and leads to the consumption of more similar items than in the case when $\rho = 0$. However, since these spillovers also impact user learning in the no recommendation case, users *over-exploit* these and increasingly consume items similar to high value items that they have consumed before. This is also illustrated in the top row of Figure 2, which shows how the consumption paths in the oracle and no-recommendation regimes vary as ρ increases and is in line with this intuition.

As shown in the bottom row of Figure 2, this effect is further amplified as the level of risk aversion increases: the degree of local consumption drastically increases as γ increases. This is due to the fact that the spillovers not only impact the mean expected belief about quality but also the degree of uncertainty. Local consumption therefore leads to

---

[7]Due to space constraints, we only report the values for $N = 200$. However, we further conducted the same exercises for $N = 100$ and $N = 500$ and the results are qualitatively similar to those reported here.

Fig. 1. Local Consumption and Correlation

Notes: The figure shows the consecutive consumption path difference between the no recommendation, recommendation, and oracle regime. The figure on the left displays the average consecutive consumption distance aggregating over simulations where $\rho \in (0, 1)$ and the figure on the right displays the average consecutive consumption distance aggregating over simulations where $\rho = 0$. The shaded area represents the 95% confidence interval.

users having less uncertainty about certain areas of the product space and risk aversion may lead them to increasingly consume nearby items.

In sum, filter-bubble effects only arise when there is an inherent correlation between the realized utilities of the items in the product space. When there is a correlation between the realized utilities then filter bubbles can naturally arise due to the nature of how individuals acquire additional information about the remaining items. We have shown that, unless users are provided with additional information to guide their consumption choices, then these information spillovers and user risk-aversion can lead users into filter bubbles. When users consume high valued items, they exploit the underlying correlation across different items' values, stronger for similar items, which leads them to increasingly consume items in narrower and narrower portions of the product space. Risk aversion may lead users into performing local consumption even when they have a low valuation of nearby items just because they know what to expect from the item. Recommendation leads to these effects being mitigated by providing users with additional information on items outside the already explored portions of the product space. Logically, if all uncertainty were resolved as in the oracle regime, then such behavior is not present.
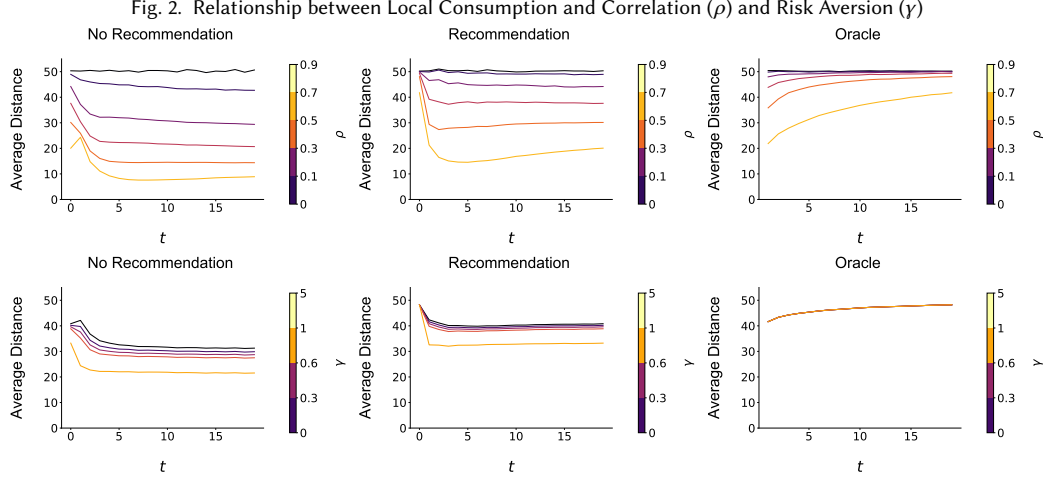
4.2 User Welfare and Item Diversity

In this section we primarily focus on the impact of recommendation on user welfare and the overall diversity of the items that they consume. While in the previous section we looked at the distance between consecutive items, in this section we focus on a diversity measure that considers the entire consumed set of items. The diversity measure we utilize is common in RS literature (e.g. [33]) – the average normalized pairwise distance between the consumed items:

$$D_i := \frac{1}{N} \frac{1}{T(T-1)} \sum_{n,m \in C_i^T : n \neq m} d(n,m)$$

Finding 2 summarizes the main results on item diversity:

FINDING 2. *The impact of recommendation on item diversity:*

(1) When $\rho = 0$, item diversity is the same across all three recommendation regimes;

Fig. 2. Relationship between Local Consumption and Correlation (ρ) and Risk Aversion (γ)

Notes: Each figure plots the average consecutive consumption distance across time as the inherent correlation between the valuation of the items, ρ, varies (top row) and the level of risk-aversion, γ, varies (bottom row). In both rows, the left displays the no-recommendation regime, the center displays the recommendation regime, and the right displays the oracle regime.

(2) When $\rho \in (0, 1)$, item diversity decreases across all recommendation regimes but decreases the most in the no-recommendation regime. This effect is amplified as ρ increases as well as when users become more risk-averse.

As before, when there is no correlation between valuations, item diversity is the same across different recommendation regimes. The over-exploitation of information spillovers when $\rho \in (0, 1)$ leads to item diversity being lowest in the no-recommendation regime. As a result, this effect gets amplified as ρ increases, which leads to the gap in diversity between the regimes to increase as ρ increases. The top row of Figure 2 shows how diversity varies as ρ increases across the three regimes that we consider. There is a similar increasing diversity gap as γ, or the level of risk-aversion, increases as can be seen in the bottom row of Figure 2. The mechanisms behind these effects directly parallels those discussed in the previous section since low average sequential consumption distance directly leads to low diversity.
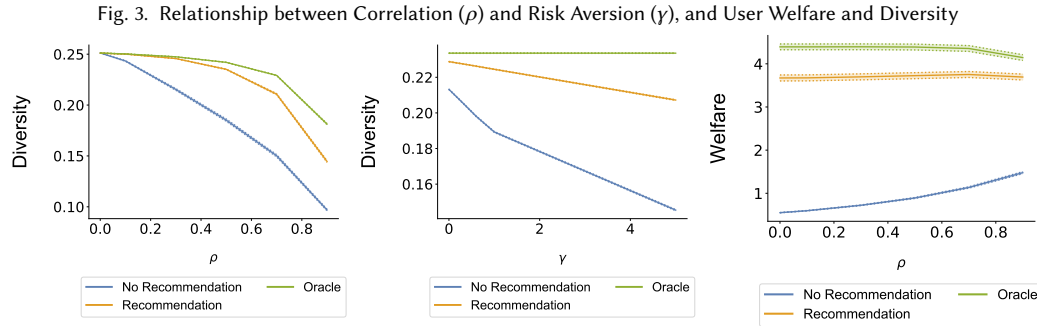
We now study how recommendation impacts user welfare. In our model users make consumption decisions that maximize their current period ex-ante utility that depends on their beliefs in that period. Thus, from an ex-ante perspective, they make optimal decisions, but our primary interest is in understanding how the ex-post, or realized, utility varies across regimes and parameter values. We define user's *ex-post* welfare as the average of the realized values, controlling for the effect of T:

$$W_i := \frac{1}{T} \sum_{n \in C_i^T} x_{i,n}$$

Finding 3 states our main findings of the impact of recommendation on ex-post welfare, which can be seen in the rightmost plot of Figure 3:

FINDING 3. *The impact of recommendation on consumer welfare is as follows:*

(1) *Under oracle recommendation, welfare is invariant with respect to ρ.*

(2) *Under no recommendation, welfare is increasing in ρ.*

10

Fig. 3. Relationship between Correlation (ρ) and Risk Aversion (γ), and User Welfare and Diversity

Notes: The figures on the left and center display the relationship between ρ and overall consumed item diversity (left) as well as γ and overall consumed item diversity (center). The figure on the right displays the relationship between ρ and overall welfare. The shaded area represents the 95% confidence interval.

(3) *Recommendation introduces welfare gains relative to no recommendation, but these gains are decreasing as ρ increases.*

The most interesting observation is that the value of recommendation decreases as ρ decreases. While welfare in the recommendation regime is flat as we increase ρ, it is increasing in the no-recommendation regime and thus the welfare gap between the two shrinks as ρ increases. The intuition is clear as recommendation provides users with information that allows them to better guide their decisions and increase welfare. However, as ρ increases users get increasingly more information from consuming items since the realized utility is now more informative about the utility of nearby items. Thus, since consumption decisions themselves yield valuable information, the information provided by recommendation is less valuable to the user.

One striking observation is that the decrease in diversity does not appear to be associated with a decline in welfare. Indeed, it appears that the opposite is the case - that low diversity is associated with higher welfare and vice versa. We next explore the relationship between welfare and diversity.

Finding 4 summarizes our findings on the relationship between diversity and welfare:

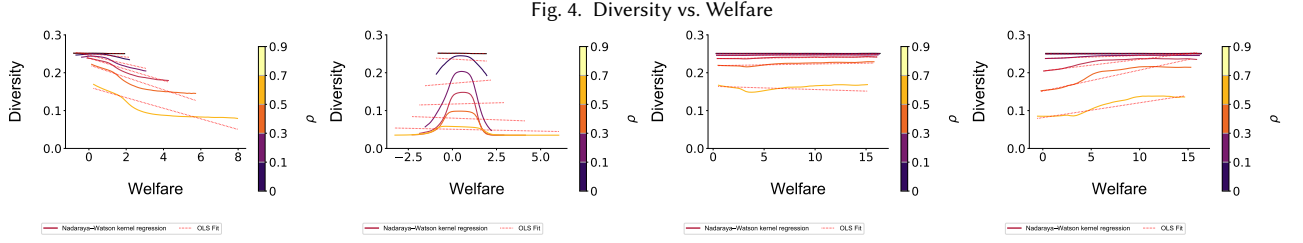FINDING 4. *In the no-recommendation regime, diversity and welfare are:*

(1) *Negatively correlated when users have no risk-aversion;*
(2) *Uncorrelated when users have high levels of risk-aversion.*

In the recommendation regime, diversity and welfare are:

(1) *Uncorrelated when users have no risk-aversion;*
(2) *Positively correlated when users have high levels of risk-aversion.*

In the oracle regime, diversity and welfare are always uncorrelated.

Figure 4 shows how diversity and welfare correlate for the no recommendation case as we vary the degree of risk aversion. When there is no risk-aversion there is a negative correlation between welfare and diversity. This is since, in this case, a user will select the item that she currently believes has the highest expected value. High item diversity in this case can arise from a user who consumes an item she disliked and updates her beliefs about nearby items negatively. As a result, in the following period she will pick an item far away in the product space from the item that was previously

Fig. 4. Diversity vs. Welfare

Notes: The figure plots the relationship between diversity and welfare under the recommendation (top row) and no recommendation (bottom row) regime. In both rows, the figure on the left displays this relationship for when $\gamma = 0$ and the figure on the right for when $\gamma = 5$.

consumed. If instead the user valued highly the item that she had consumed, then she is more likely to pick a nearby item. The information spillovers therefore lead to high item diversity being negatively correlated with welfare.

This only happens since $\gamma = 0$ leads to users only caring about the expected value of the item. However, as we saw in Findings 1 and 2, increasing γ can lead to lower diversity and increasingly local consumption due to the fact that the degree of uncertainty now impacts users' choices. This weakens the negative relationship between diversity and welfare since both negative and positive experiences with an item reduce uncertainty about surrounding items. This leads to the inverted-U shape found in Figure 4 when γ is relatively large (e.g. $\gamma = 5$) though diversity and welfare are virtually uncorrelated in the data. In the recommendation and oracle regimes, under risk neutrality ($\gamma = 0$), welfare and diversity are uncorrelated, while under risk aversion ($\gamma = 5$), it is possible to observe an actual positive relation between diversity and welfare as recommendations are able to reduce uncertainty and facilitate exploration of the product space.

4.3 User Homogenization

In this section, we focus on comparisons across users and investigate how the consumed set of items across users varies across different recommendation regimes and parameter values. In particular we look at the degree of *homogenization* between users. Similar to other papers that study the degree of homogenization in RS (e.g. [6]) we measure homogeneity via the Jaccard index between the consumption sets of users:

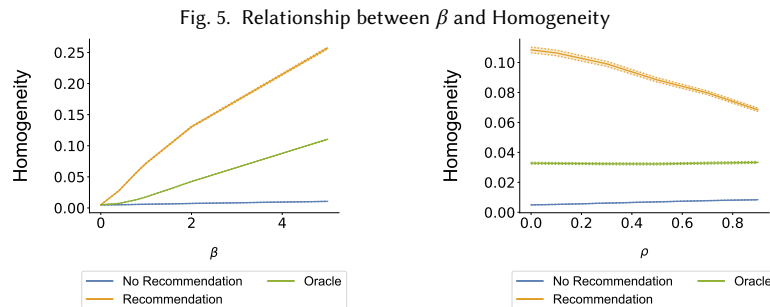$$H := \frac{1}{|I|(|I| - 1)} \sum_{i,j \in I : i \neq j} d_J(C_i^T, C_j^T)$$

where d_J denotes the Jaccard index and $H \in [0, 1]$.

Finding 5 summarizes our findings on the impact of recommendation on user homogeneity:

FINDING 5. *The impact of recommendation on homogeneity is as follows:*

(1) *Highest under recommendation and lowest under no recommendation;*

(2) *Increasing in β, or the weight of the common-value component;*

(3) *Decreasing in ρ for partial recommendation, but weakly increasing in ρ for no recommendation.*

First, we study how the degree of homogenization varies as we increase β, the weight of the common value component. As β increases we expect that users should become increasingly homogeneous as the realized utilities of the items are now increasingly similar. Figure 5 confirms that as the weight of the common-value component β increases, users consume increasingly similar sets of items. The homogenization effect is strongest under the recommendation regime since the revelation of the common-value component induces users to consume items in similar areas of the product

Fig. 5. Relationship between β and Homogeneity

Notes: This figure displays the value of the homogeneity measure as we vary the weight of the common value component, β (left) and correlation between valuations, ρ (right). Each line represents this plot for a single recommendation regime. The shaded area represents the 95% confidence interval.

space. As β increases, some amount of homogenization is optimal as can be seen from the oracle case. However, since users in the no-recommendation regime do not know the common-value component they engage in local consumption in different areas of the product space which leads to less than optimal homogeneity.

We next study how the degree of homogeneity varies with ρ. Figure 5 shows how homogeneity decreases as ρ increases in the recommendation regime. As was highlighted in Findings 1 and 2, the degree of local consumption increases with ρ. Even though the revelation of the common-value component induces them to search in similar parts of the product space, their idiosyncratic components induce them to consume items in a more localized area of the product space as ρ increases which leads to a decline in homogeneity.

5 RECOMMENDER SYSTEM EVALUATION

In this section we discuss how the insights from our model of user decision-making can inform the evaluation and design of recommender systems. The classic approach to evaluation is to predict user ratings for items and to compare how accurate this prediction is to recorded ratings data, either explicitly given by users or inferred from behavioral data. The RS should then recommend the items with the highest predicted ratings [1].

There has been a recent movement away from such evaluation measures due to the observation that accurate recommendations are not necessarily useful recommendations [22]. Our model illustrates a mechanism behind this observation. Consider the domain of movie recommendation and suppose a user has just watched the movie *John Wick* and rated it highly. A RS attempting to predict user ratings may then predict that this user is very likely to enjoy the sequel, *John Wick: Chapter Two*, as well. However, the user herself may also have made this inference since the two movies are very similar to each other. Thus, recommending this movie would not be not useful since the recommendation gives the user little information that she did not already know. The key insight is it is not useful since *it ignores the inference the user themselves made and their updated beliefs*. The user may watch *John Wick: Chapter Two*, then, even without recommendation, and the value of the recommendation was small.

This intuition implies that RS should collect additional data beyond that which is traditionally recorded. The first and most crucial type of data to collect is individual user *beliefs* about items that they have not yet consumed. As illustrated by our model, these beliefs are what drive the future consumption decisions of users and understanding these beliefs is crucial for determining the value of recommending certain items.[8] The second type of data that is relevant for RS

---

[8]Additionally, user beliefs contain information that may not have been observed by the recommender that only observes user choices on the platform.

designers to collect is how user beliefs change over time and, in particular, not just how individuals value the item they just consumed, but also how it impacts their beliefs about the quality of similar items.[9] The third type of data is the risk-aversion levels of users as our model illustrates that the risk preferences of users are important for understanding what information RS can provide that materially leads users to alter their consumption patterns.

A natural follow-up question is how this additional data should be utilized in the design of good recommendations. Our model posits that recommendation provides value to users by providing them with information about the true valuation of an item if they were to consume it. Thus, the prediction problem for the recommender becomes predicting what item the user would choose with no recommendation and, correspondingly, what would be the most useful information to provide to the user that would lead her to consume a *better* item than she would without recommendation. This links back to the intuition our model provided for the *John Wick* example whereby collecting user beliefs and measuring how the user updated beliefs about similar items would lead the recommender to understand that the user would consume *John Wick: Chapter Two*. Our approach would therefore imply that, with this as a starting point, the recommender's problem would be to predict what is the most useful information to give the user leading them to change the item that they eventually consume.

There have been a number of alternative recommendation evaluation metrics proposed in the literature with the aim of providing more useful recommendations than those provided by accuracy metrics, such as serendipity [18], calibration [28], coverage [11], novelty [30], and many others. Our approach most closely follows the set of proposed serendipity measures which are surveyed in [18]. As discussed by [20], serendipitous recommendations are said to "have the quality of being both unexpected and useful" which is in line with the primary intuition behind our approach. The primary difference between our proposed approach and those existing in the literature is that ours crucially hinges on understanding user beliefs and the risk-preferences of users. For instance, [16, 30] propose unexpectedness metrics that look at the dissimilarity of the proposed recommended items compared to what the recommender already knows the user likes. This metric depends only on the proposed item-set and not necessarily on the user's beliefs or how such a recommendation will change the item that the user consumes. [17] provide a comprehensive overview of possible definitions of serendipity and ours is closest to their "motivational novelty" definition, which is that the user was persuaded to consume an item as a result of recommendation.

Indeed, our approach allows us to give precise definitions for what it means for a recommendation to be *unexpected* and *useful* in the spirit of serendipitous recommendations. Our evaluation measure leads to useful recommendations since it leads users towards better items than they would consume without recommendation. It further results in "unexpected" recommendations since it explicitly incorporates user beliefs and thus allows the RS to understand how "unexpected" a recommendation would be from the perspective of a user. Finally, such a measure may lead to a perceived broadening of user preferences as has been discussed in [13, 32]. However, under our interpretation, it may be that their underlying preferences are unchanged and, instead, that recommendation and consumption themselves provides information that encourages users to explore different portions of the product space.

---

[9]Characterizing the similarity between items has been an important goal of designing content-based recommendations, though as noted by [31], how users perceive similarity between items is not always in line with how similarity is computed in content-based RS. Understanding how this impacts which items users update their beliefs about is an important direction for future work.

REFERENCES

[1] Gediminas Adomavicius and Alexander Tuzhilin. 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge & Data Engineering* 6 (2005), 734–749.

[2] James R Bettman, Mary Frances Luce, and John W Payne. 1998. Constructive consumer choice processes. *Journal of consumer research* 25, 3 (1998), 187–217.

[3] Antoine Billot, Itzhak Gilboa, and David Schmeidler. 2008. Axiomatization of an exponential similarity function. *Mathematical Social Sciences* 55, 2 (2008), 107–115.

[4] Òscar Celma and Pedro Cano. 2008. From hits to niches?: or how popular artists can bias music recommendation and discovery. In *Proceedings of the 2nd KDD Workshop on Large-Scale Recommender Systems and the Netflix Prize Competition*. ACM, 5.

[5] Òscar Celma and Perfecto Herrera. 2008. A new approach to evaluating novel recommendations. In *Proceedings of the 2008 ACM conference on Recommender systems*. 179–186.

[6] Allison JB Chaney, Brandon M Stewart, and Barbara 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.

[7] Li Chen, Marco De Gemmis, Alexander Felfernig, Pasquale Lops, Francesco Ricci, and Giovanni Semeraro. 2013. Human decision making and recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 3, 3 (2013), 1–7.

[8] Paolo Cremonesi, Franca Garzotto, and Roberto Turrin. 2013. User-centric vs. system-centric evaluation of recommender systems. In *IFIP Conference on Human-Computer Interaction*. Springer, 334–351.

[9] Seth Flaxman, Sharad Goel, and Justin M Rao. 2016. Filter bubbles, echo chambers, and online news consumption. *Public opinion quarterly* 80, S1 (2016), 298–320.

[10] Daniel Fleder and Kartik Hosanagar. 2009. Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity. *Management science* 55, 5 (2009), 697–712.

[11] Mouzhi Ge, Carla Delgado-Battenfeld, and Dietmar Jannach. 2010. Beyond accuracy: evaluating recommender systems by coverage and serendipity. In *Proceedings of the fourth ACM conference on Recommender systems*. ACM, 257–260.

[12] Itzhak Gilboa and David Schmeidler. 1995. Case-based decision theory. *The quarterly Journal of economics* 110, 3 (1995), 605–639.

[13] Jonathan L Herlocker, Joseph A Konstan, Loren G Terveen, and John T Riedl. 2004. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems (TOIS)* 22, 1 (2004), 5–53.

[14] Charles Hodgson and Gregory Lewis. 2019. You Can Lead a Horse to Water: Spatial Learning and Path Dependence in Consumer Search. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC '19)*. Association for Computing Machinery, New York, NY, USA, 215. https://doi.org/10.1145/3328526.3329595

[15] Kartik Hosanagar, Daniel Fleder, Dokyun Lee, and Andreas Buja. 2013. Will the global village fracture into tribes? Recommender systems and their effects on consumer fragmentation. *Management Science* 60, 4 (2013), 805–823.

[16] Marius Kaminskas and Derek Bridge. 2014. Measuring surprise in recommender systems. In *Proceedings of the workshop on recommender systems evaluation: dimensions and design (Workshop programme of the 8th ACM conference on recommender systems)*. Citeseer.

[17] Denis Kotkov, Joseph A Konstan, Qian Zhao, and Jari Veijalainen. 2018. Investigating serendipity in recommender systems based on real user feedback. In *Proceedings of the 33rd Annual ACM Symposium on Applied Computing*. 1341–1350.

[18] Denis Kotkov, Shuaiqiang Wang, and Jari Veijalainen. 2016. A survey of serendipity in recommender systems. *Knowledge-Based Systems* 111 (2016), 180–192.

[19] Sarah Lichtenstein and Paul Slovic. 2006. *The construction of preference*. Cambridge University Press.

[20] Andrii Maksai, Florent Garcin, and Boi Faltings. 2015. Predicting online performance of news recommender systems through richer evaluation metrics. In *Proceedings of the 9th ACM Conference on Recommender Systems*. ACM, 179–186.

[21] Andreu Mas-Colell, Michael Dennis Whinston, Jerry R Green, et al. 1995. *Microeconomic theory*. Vol. 1. Oxford university press New York.

[22] Sean M McNee, John Riedl, and Joseph A Konstan. 2006. Being accurate is not enough: how accuracy metrics have hurt recommender systems. In *CHI'06 extended abstracts on Human factors in computing systems*. ACM, 1097–1101.

[23] Judith Möller, Damian Trilling, Natali Helberger, and Bram van Es. 2018. Do not blame it on the algorithm: an empirical assessment of multiple recommender systems and their impact on content diversity. *Information, Communication & Society* 21, 7 (2018), 959–977.

[24] Tien T Nguyen, Pik-Mai Hui, F Maxwell Harper, Loren Terveen, and Joseph A Konstan. 2014. Exploring the filter bubble: the effect of using recommender systems on content diversity. In *Proceedings of the 23rd international conference on World wide web*. ACM, 677–686.

[25] Eli Pariser. 2011. *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin.

[26] Pearl Pu, Li Chen, and Rong Hu. 2011. A user-centric evaluation framework for recommender systems. In *Proceedings of the fifth ACM conference on Recommender systems*. 157–164.

[27] Eric Schulz, Rahul Bhui, Bradley C Love, Bastien Brier, Michael T Todd, and Samuel J Gershman. 2019. Structured, uncertainty-driven exploration in real-world consumer choice. *Proceedings of the National Academy of Sciences* (2019), 201821028.

[28] Harald Steck. 2018. Calibrated recommendations. In *Proceedings of the 12th ACM conference on recommender systems*. 154–162.

[29] Jutta Treviranus and Stephen Hockema. 2009. The value of the unpopular: Counteracting the popularity echo-chamber on the Web. In *2009 IEEE Toronto International Conference Science and Technology for Humanity (TIC-STH)*. IEEE. https://doi.org/10.1109/tic-sth.2009.5444430

[30] Saúl Vargas and Pablo Castells. 2011. Rank and relevance in novelty and diversity metrics for recommender systems. In *Proceedings of the fifth ACM conference on Recommender systems*. ACM, 109–116.

[31] Amy A Winecoff, Florin Brasoveanu, Bryce Casavant, Pearce Washabaugh, and Matthew Graham. 2019. Users in the loop: a psychologically-informed approach to similar item retrieval. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 52–59.

[32] Yuan Cao Zhang, Diarmuid Ó Séaghdha, Daniele Quercia, and Tamas Jambor. 2012. Auralist: introducing serendipity into music recommendation. In *Proceedings of the fifth ACM international conference on Web search and data mining*. 13–22.

[33] Cai-Nicolas Ziegler, Sean M McNee, Joseph A Konstan, and Georg Lausen. 2005. Improving recommendation lists through topic diversification. In *Proceedings of the 14th international conference on World Wide Web*. ACM, 22–32.