

Filter Bubble or Homogenization? Disentangling the Long-Term Effects of Recommendations on User Consumption Patterns

Md Sanzeed Anwar
University of Michigan
Ann Arbor, MI, USA
sanzeed@umich.edu

Grant Schoenebeck
University of Michigan
Ann Arbor, MI, USA
schoeneb@umich.edu

Paramveer S. Dhillon
University of Michigan
Ann Arbor, MI, USA
dhillonp@umich.edu

ABSTRACT

Recommendation algorithms play a pivotal role in shaping our media choices, which makes it crucial to comprehend their long-term impact on user behavior. These algorithms are often linked to two critical outcomes: homogenization, wherein users consume similar content despite disparate underlying preferences, and the filter bubble effect, wherein individuals with differing preferences only consume content aligned with their preferences (without much overlap with other users). Prior research assumes a trade-off between homogenization and filter bubble effects and then shows that personalized recommendations mitigate filter bubbles by fostering homogenization. However, because of this assumption of a tradeoff between these two effects, prior work cannot develop a more nuanced view of how recommendation systems may independently impact homogenization and filter bubble effects. We develop a more refined definition of homogenization and the filter bubble effect by decomposing them into two key metrics: how different the average consumption is between users (inter-user diversity) and how varied an individual's consumption is (intra-user diversity). We then use a novel agent-based simulation framework that enables a holistic view of the impact of recommendation systems on homogenization and filter bubble effects. Our simulations show that traditional recommendation algorithms (based on past behavior) mainly reduce filter bubbles by affecting inter-user diversity without significantly impacting intra-user diversity. Building on these findings, we introduce two new recommendation algorithms that take a more nuanced approach by accounting for both types of diversity.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Recommender Systems, Diversity, Filter Bubbles, Homogenization

Note: This paper was accepted at the ACM Web Conference 2024 (WWW '24), May 13–17, 2024, Singapore, Singapore.

1 INTRODUCTION

With the advent of the Internet, much of our social interaction and entertainment has moved online, dispersed across various platforms that each curate their own content. Recommendation algorithms help us navigate these content collections, influencing our choices by providing context. However, lingering questions exist about the effects of these algorithms on our media consumption and social behavior. Previous research has examined their role in fostering homophilous communities [10], amplifying a *rich-get-richer* effect in online social ties [19], and potential bias against minority users [12].

This paper aims to deepen our understanding of two key phenomena often linked to recommendation algorithms: *homogenization* and *filter bubbles*. Past studies (e.g., Nguyen et al. [15], Aridor et al. [1]) indicate that personalized recommendations based on past consumption can mitigate filter bubble effects, but they do so at the expense of homogenizing the audience. These findings, however, only look at how homogeneous agents are in terms of the average item consumed by each agent, and do not examine the diversity of consumption of individual users. Thus, the question remains: *do these algorithms diversify or homogenize the set of items any particular individual consumes?* The answer to this question has important implications for recommendation algorithm design.

Given the relative lack of control over confounding factors when using observational data, we explore these questions through a simulation study using agent-based modeling. We start by proposing a novel simulation model consisting of users and items. Each item has a quality and a genre, both represented via real numbers. On the other hand, each user has an underlying preference for what genre of item they like the most, also represented via a real number. Quality indicates how universally desirable the item is, while the genre of an item impacts different users differently as users prefer to consume items nearer their genre preferences. When deciding which item to consume, users estimate and maximize item utility according to a set of available signals, including a recommendation signal provided by the system. We consider seven such recommendation algorithms; four of these act as idealized baselines, while the remaining three are based on past consumption.

Our first contribution is to disentangle the effects of recommendation algorithms on two types of diversity: *inter-user diversity*, which measures how the mean of individual consumption varies across users, and *intra-user diversity*, which measures how diverse an individual's consumption is on average. This insight leads us to operationalize a new definition of the filter bubble effect as a ratio between inter-user and intra-user diversity. The intuition behind our definition is that a weak filter bubble effect exists when all users consume the same blockbuster items (i.e., low inter-user diversity), but also when each individual user consumes items from

a wide range of genres and are not just confined to their own preferences (i.e. high intra-user diversity). Results from our simulations show that the past consumption-based recommendations alleviate the filter bubble effect only by homogenizing the population towards blockbuster items and reducing inter-user diversity, without significantly affecting intra-user diversity.

Next, as our second contribution, we propose two novel recommendation ideas: ***binned consumption-based recommendation*** and ***skewed top pick recommendation***, inspired by the insight that understanding the dynamics between homogenization and filter bubbles requires examining both inter-user and intra-user diversity. Binned consumption-based recommendation recommends the set of most consumed items in each genre, therefore recommending a curated set of items and eliminating bias towards blockbuster items. This recommendation alleviates the filter bubble effect by not only decreasing inter-user diversity, but also significantly increasing intra-user diversity. On the other hand, skewed top pick recommendation prioritizes exposure to more niche items. Rather than alleviating the filter bubble effect, this recommendation focuses on simultaneously increasing inter-user and intra-user diversity.

These novel recommendation algorithms are, on the surface, very similar to prior work which intentionally recommends a diverse slate of items to users [4, 11, 14]. However, there are also important differences: we are studying which items are consumed rather than which items are recommended. How users make use of the system’s recommendations significantly impacts our results. This in turn is subtly affected by the entirety of information the system is providing and the other information available to the agent.

The rest of this paper is organized as follows: Section 2 reviews related work. Section 3 details our simulation framework, and Section 4 introduces our novel measures for homogeneity and the filter bubble effect. Then, Section 5 outlines the empirical setup, with Section 6 presenting the results. Next, Section 7 proposes new recommendation algorithms. Finally, Section 8 discusses our findings and also suggests directions for future research.

2 RELATED WORK

Our work combines multiple streams of research on recommendation algorithms. In particular, our research builds on past literature investigating the role of recommendation algorithms in reducing inter-user diversity through homogenization, causing filter bubbles, and the possible interplay between these two phenomena.

Homogenization and inter-user diversity: The existing literature strongly supports a connection between recommendation algorithms and homogenization, largely attributing this to a popularity bias or feedback loop that continually directs users toward a common set of popular items—often referred to as blockbusters. For instance, Salganik et al. [18] used an experimental method allowing participants to listen to a song and decide to download it based on its popularity. The study revealed a widening disparity in song success, signaling a popularity bias. Fleder and Hosanagar [7] employed a 2D simulation model of consumers and items, showing through the Gini coefficient that sales diversity diminishes with collaborative filtering-based recommendations. Similarly, Chaney et al. [5] used a more complex simulation where consumption choices

are deterministic, based on both recommendation ranking and personal utility signals. Their findings indicate that user consumption overlap, and thus homogenization, increases over time. Mansoury et al. [13] adopted a hybrid simulation using a real movie dataset and leveraged KL divergence to demonstrate convergence in genre distributions among users. Across these studies, it is evident that recommendation algorithms reinforce the popularity of already well-consumed items, pushing the general population toward these choices and perpetuating the cycle.

Filter bubbles: The existence of filter bubbles is far more contentious than that of homogenization. Eli Pariser first introduced the term “Filter Bubble” in 2011 to describe how personalization could limit exposure to content that diverges from user preferences [17]. However, he didn’t provide a definitive framework, resulting in an ongoing debate marked by an absence of a universally accepted, operational definition [3].

One approach to defining filter bubbles is to adopt Pariser’s original concept of literal “bubbles” or “filters” that fully restrict exposure to non-conforming content. Counterarguments suggest that recommendation algorithms actually expand user horizons. Flaxman et al. [6], for instance, found that recommendation algorithms expose consumers to more diverse news than they would find independently. Similarly, Hosanagar et al. [9] discovered that while recommendations reduce the distance between users within the same preference cluster, they also reduce the distance across different clusters. Alternatively, filter bubbles can be conceptualized as focused exposure to content that aligns with user preferences. While algorithms may introduce some diverse content, they predominantly amplify existing preferences. O’Callaghan et al. [16] found that top- K related YouTube channels often mirror the political orientation of the original channel, suggesting concentrated exposure. Bakshy et al. [2] revealed a 15% reduction in exposure to conflicting viewpoints on Facebook due to news feed filtering. Geschke et al. [8] further bolstered this view using agent-based modeling to show that social and technological factors enhance naturally occurring filter bubble effects.

In this paper, we adopt a nuanced perspective that eschews the notion of a literal “bubble.” Instead, we define the filter bubble effect on a continuum, and it intensifies when individuals with different preferences increasingly consume different types of items.

Connections between homogenization and filter bubbles: There is limited prior work that directly examines the interplay between homogenization and filter bubbles. To our knowledge, only two studies—by Nguyen et al. [15] and Aridor et al. [1]—address this trade-off. Nguyen et al. work with the MovieLens dataset, where users get personalized recommendations and rate movies post-viewing. Aridor et al. employ a simulation where items have randomly drawn social and user-specific valuations, with recommendations tailored to the latter. These studies operationalize the filter bubble effect as how concentrated individual user consumption is (at a high level, in our terminology, this is $1/\text{intra-user diversity}$, the inverse of our intra-user diversity measurement). On the other hand, they operationalize homogeneity as increasing overlap between the items consumed by different users (again, at a high level, in our terminology this is $\text{intra-user diversity}/\text{inter-user diversity}$, the inverse of our filter bubble effect measurement). These studies then show that personalized recommendations increase their

measure of homogeneity and decrease their measure of the filter bubble effect, establishing a direct trade-off between the two.

Our research enriches this body of work by revealing that the homogenization and filter bubble effects of recommendations can be disentangled into their impact on both inter-user and intra-user diversity. We argue that the trade-off between homogenization and filter bubbles is not as direct as previously assumed. Specifically, recommendations not only can impact inter-user diversity but also can augment intra-user diversity—a facet unaccounted for in prior studies.

3 RESEARCH DESIGN

3.1 Research question

As previously stated, we are interested in uncovering a more complete picture of the dynamics between the filter bubble and homogenization effects of recommendation algorithms. Hence, we seek to answer the following research question in this paper:

Can we explain the dynamics between homogenization and filter bubble effects of recommendations beyond a simple trade-off by considering both inter-user and intra-user diversity?

3.2 Simulation model

The core building blocks of our simulated world $\mathcal{W}$ are m users and n items. Each user $j = 1, 2, \dots, m$ has an associated genre preference p_j drawn independently from some distribution $\mathcal{P}$, i.e.

$$p_j \sim_R \mathcal{P} \quad \forall j = 1, 2, \dots, m$$

On the other hand, each item i in this world has some inherent quality q_i drawn independently from distribution $\mathcal{Q}$, as well as some genre attribute g_i drawn independently from distribution $\mathcal{G}$.

$$q_i \sim_R \mathcal{Q}, \quad g_i \sim_R \mathcal{G} \quad \forall i = 1, 2, \dots, n$$

We choose to use real values for user preferences and item genres because it allows us to distinguish and observe niche users and items without adding additional complexity to our model. Specifically, users with preferences situated away from the mode(s) of the preference distribution are considered niche. Similarly, items with genres situated away from the mode(s) of the genre distribution are considered niche.

The progression of time t in our simulated world $\mathcal{W}$ is discrete, and continues for T rounds. Initially, the world consists of k_{init} items. At each of the T discrete rounds, k_{new} items are added to the world. Therefore, $n = k_{init} + T \cdot k_{new}$.

User utility in our model consists of two components: a shared quality component corresponding to the quality q_i of item i , and an affinity component corresponding to the loss due to misalignment between the genre g_i of item i and the preference p_j of user j .

Definition. The *utility* received by user j by consuming item i is $U(j, i) = q_i - |p_j - g_i|$.

In each round, each user consumes exactly one item, at which point the said item becomes unavailable to them for future consumption. Following convention, we model users as utility maximizers: in each round, a user j attempts to choose the item i that would yield the maximum utility for them from the set of items they have yet to consume.

However, users do not know an item's true quality or true genre, and therefore cannot directly compute its utility. Instead, in each round t , each user estimates the utility of each previously unconsumed item as a function $\mathcal{F}$ of the following three signals available to them:

- (1) A private signal q_i^j , which is a noisy personal estimate of the quality of item i obtained by adding some noise ξ_i^j drawn from distribution $\mathcal{N}_{qual}$ to the true quality q_i :

$$q_i^j = q_i + \xi_i^j, \quad \xi_i^j \sim_R \mathcal{N}_{qual}$$

- (2) The perceived distance between their preference and the item genre. User j has a noisy personal estimate g_i^j of the true genre of item i , obtained by adding some noise δ_i^j drawn from distribution $\mathcal{N}_{genre}$ to the true genre g_i :

$$g_i^j = g_i + \delta_i^j, \quad \delta_i^j \sim_R \mathcal{N}_{genre}$$

The signal used by user j for estimating utility is $|p_j - g_i^j|$.

- (3) Recommendation $\mathbf{r}_i^j(t)$ provided by the system consisting of one or more pieces of information about the item, such as the number of times the item has been consumed etc. (i.e., $\mathbf{r}_i^j(t)$ is a vector of real numbers)

In other words, in each round t , each user j chooses available item i that maximizes the estimated utility

$$\hat{U}(j, i, t) = \mathcal{F}(q_i^j, |p_j - g_i^j|, \mathbf{r}_i^j(t))$$

The shared quality component means that a user can learn something about each item from all other users. However, the affinity component means that a user is intuitively best informed by other similar users. Therefore, if a user only listens to similar users, they fail to learn as much as possible about the quality component. But if they listen too much to the global consensus, they fail to learn as much as possible about the affinity component.

Overall, the world $\mathcal{W}$ in our simulation framework can be defined via the collection of hyperparameters and distributions used in order to generate the user and item properties:

$$\mathcal{W} = (m, k_{init}, k_{new}, T, \mathcal{P}, \mathcal{Q}, \mathcal{G}, \mathcal{N}_{qual}, \mathcal{N}_{genre})$$

3.2.1 Utility estimation by users. While the exact nature of the estimator $\mathcal{F}$ is unknown, we can use a machine learning model as a suitable replacement for it in our simulation. This model would take private signal q_i^j , perceived distance between preference and genre $|p_j - g_i^j|$ and recommendation $\mathbf{r}_i^j(t)$ as features. In other words, the feature vectors for users j and items i in round t for this ML model are given by

$$\mathbf{x}_{ji}(t) = [q_i^j, |p_j - g_i^j|, \mathbf{r}_i^j(t)]$$

For simplicity, we replace $\mathcal{F}$ with a linear regression model¹: at each step, user j chooses item i to consume based on $\hat{U}(j, i)$ estimated via

$$\hat{U}(j, i, t) = \mathbf{w}_0 + \mathbf{w}^\top \mathbf{x}_{ji}(t) \quad \forall j, i$$

We assume that the coefficients $\mathbf{w}_0, \mathbf{w}$ are not user-specific. The details about how we learn this regression model in order to run our simulation are provided in section 5.1.

¹We experimented with using a more complex multi-layer perceptron neural network model, but the results were qualitatively the same.

4 MEASURING LONG-TERM EFFECTS OF RECOMMENDATIONS

In this section, we define our new measure to quantify the long-term effects of recommendations in terms of *inter-user diversity* and *intra-user diversity*, and then describe the filter bubble effect and homogenization using that measure. However, before we do that, we need to define some preliminary concepts.

We start with the consumption set of a given user, simply the ordered collection of items they consume.

Definition. The *consumption set* of a user j at time t is defined to be an ordered collection $C_j(t)$ of items consumed by the user prior to round t , i.e. $C_j(t) = (c_j^1, c_j^2, \dots, c_j^{t-1})$.

For brevity, we omit t from the notation and assume that C_j represents consumption of user j after the final round, i.e., after round T , unless stated otherwise.

Given the items consumed by a user, we can take the mean of the genres of these items as an indicator of what general genre of items they consume. We define this as the *mean consumed genre*.

Definition. The *mean consumed genre* of a user j , $\mu_j = \frac{1}{T} \sum_{i \in C_j} g_i$, is the mean of the genres of the items consumed by the user C_j .

On the other hand, we can take the variance of the genres of the items consumed by the user to indicate of how broad their consumption is. We define this as the *consumed genre variance*.

Definition. The *consumed genre variance* of a user j , $\sigma_j^2 = \text{Var}[\{g_i | i \in C_j\}]$, is the variance of the genres of the items consumed by the user C_j .

4.1 Inter-user and intra-user diversity

Inter-user diversity measures how diverse average individual consumption is across all users. We previously defined mean consumed genre as a representation of the average consumption of an individual user. Therefore, we can measure inter-user diversity by taking the variance of mean consumed genres across all users. Formally,

Definition. Given a group of users $\mathcal{U}$ with their respective consumption sets, we define their *inter-user diversity* of consumption as the variance of their individual mean consumed genres, mathematically given as $\text{Var}_{j \in \mathcal{U}}[\mu_j]$.

On the other hand, intra-user diversity measures how broad the consumption of a random individual user is, indicating whether individual users are consuming from a very narrow genre range or if they are exposed to many different genres. We previously defined consumed genre variance as a measure of how broad the consumption of an individual is. Therefore, we can measure intra-user diversity by taking the mean of consumed genre variances across all users. Formally,

Definition. Given a group of users $\mathcal{U}$ with their respective consumption sets, we define their *intra-user diversity* of consumption as the mean of their individual consumed genre variances, mathematically given as $\mathbb{E}_{j \sim \mathcal{U}}[\sigma_j^2]$.

4.2 Homogenization and filter bubble effect

At the heart of this paper is the argument that we need to examine both inter-user and intra-user diversity to fully understand the role of recommendations in the dynamics between homogenization and filter bubbles, which necessitates expressing homogeneity and the filter bubble effect in terms these two types of diversity.




Figure 1: Identifying the strength of homogeneity and filter bubble effect for different levels of inter-user and intra-user diversity. The dynamics seen here motivate our novel definitions of homogeneity and the filter bubble effect.

Our high-level interpretation about these two phenomena is as follows: the filter bubble effect is stronger when users with different preferences have less in common in their respective consumption. This happens when individual mean consumptions are more spread out (higher inter-user diversity), or when individuals consume from a narrower genre range (lower intra-user diversity). On the other hand, homogenization is stronger when users with different preferences are consuming more similar items. This happens when individual mean consumptions are very close (lower inter-user diversity), and individuals consume from a narrower genre range (lower intra-user diversity). In particular, the four scenarios arising from low or high inter-user and intra-user diversity, as well how homogenization and the filter bubble effect change, are presented concisely in figure 1.

Motivated by these observations, we propose the following novel definitions of the filter bubble effect and homogeneity:

Definition. The *filter bubble effect* is given by the ratio between inter-user and intra-user diversity, i.e.,

$$\text{Filter bubble effect} = \frac{\text{Inter-user diversity}}{\text{Intra-user diversity}} = \frac{\text{Var}_{j \in \mathcal{U}}[\mu_j]}{\mathbb{E}_{j \sim \mathcal{U}}[\sigma_j^2]}$$

On the other hand, the level of homogenization of content consumption could naturally be measured via

$$\frac{1}{\sqrt{\text{Var}[\bigcup_{j \in \mathcal{U}} \{g_i | i \in C_j\}]}}$$

However, this measure of homogeneity does not explicitly contain inter and intra-user diversity. In order to demonstrate and

explain our primary contribution, and motivated by the dynamics presented in figure 1, we adopt the following operationalization of homogeneity:

Definition.

$$\text{Homogeneity} = \frac{1}{\sqrt{\text{Inter-user diversity}^2 + \text{Intra-user diversity}^2}} = \frac{1}{\sqrt{\text{Var}_{j \in \mathcal{U}}[\mu_j]^2 + \mathbb{E}_{j \sim \mathcal{U}}[\sigma_j^2]^2}}$$

Figure 2 shows our definition of homogeneity against the natural measure. As demonstrated in this figure, the two quantities are highly correlated, with a Pearson correlation coefficient of 0.93. This justifies our definition.




Figure 2: Our definition of homogeneity, formally given by $1/\sqrt{\text{inter-user diversity}^2 + \text{intra-user diversity}^2}$, vs. the inverse of the standard deviation of all consumption for each of the seven recommendation algorithms. The Pearson correlation coefficient between the two values is 0.93, i.e. they are highly correlated.

5 SIMULATION SETUP

In each round t in a world $\mathcal{W}$ in our simulation framework, the system observes past consumption of users, uses some recommendation algorithm to construct a recommendation $\mathbf{r}_i^j(t)$ from this past data about each existing item i for each user j , and sends it to user j in order to guide their consumption choices, e.g., to help user estimation of item utility.

As mentioned in section 3.2.1, we assume that users estimate the utility of each available item using a linear regression model. In order to simulate users, we need to know this regression model—we learn it by using a simulation process with two phases:

Learning users’ utility estimation model: In this first phase, we simply learn $w_0, \mathbf{w}$ from section 3.2.1. The intuition is that users learn how to interpret and combine different signals about an item and estimate its utility from their past consumption experiences. So we simulate user interaction with a similar set of items in order to learn the regression model they come to use to estimate item utility.

Simulating recommendations: In this second phase, we simulate the interactions between users and items to generate simulated data about recommendation algorithms and user consumption for our analysis.

Details of how this two-phase simulation process is implemented are provided in appendix C.

5.1 Recommendation Algorithms

We test seven recommendation algorithms in our simulation: four act as baselines, while the remaining three are past consumption-based. These algorithms are described below. Note that for each item i in world $\mathcal{W}$, we count the number of times it has been consumed at the beginning of round t and denote it via $d_i(t)$. We also define $d_i^{j'}(t) = 1$ if user j' has consumed item i before round t and 0 otherwise.

Baseline recommendations:

- (1) **No recommendation:** Our first baseline, where we can observe user consumption patterns without the effects of any recommendation signals. More precisely, $\mathbf{r}_i^j(t) = 0$ for item i and user j .
- (2) **True genre:** Shows the true genre of an item to users. More precisely, $\mathbf{r}_i^j(t) = (g_i)$ for item i and user j .
- (3) **True quality:** Shows the true qualities of items to users. More precisely, $\mathbf{r}_i^j(t) = (q_i)$ for item i and user j .
- (4) **Perfect recommendation:** Shows both the true qualities and the true genres of items to users. More precisely, $\mathbf{r}_i^j(t) = (q_i, g_i)$ for item i and user j .

Past consumption-based recommendations:

- (1) **Consumption:** Shows the number of times an item has been consumed so far to users. More precisely, in this case, $\mathbf{r}_i^j(t) = (d_i(t))$ for item i and user j .
- (2) **Singular Value Decomposition (SVD):** Shows a weighted version of the consumption number of each item to users. Consumption numbers are weighted by a similarity score between user-user pairs via SVD. More precisely, $\mathbf{r}_i^j(t) = \left(\sum_{j'=1}^m \text{Sim}(j, j') d_i^{j'}(t) \right)$ for item i and user j . Here, $\text{Sim}(j, j')$ is the cosine similarity between the representations of users j and j' obtained from SVD.
- (3) **Hybrid recommendation:** Shows both the consumption signal and SVD signal of items to users. More precisely, $\mathbf{r}_i^j(t) = \left(d_i(t), \sum_{j'=1}^m \text{Sim}(j, j') d_i^{j'}(t) \right)$ for item i and user j .

6 RESULTS

To answer the research question from section 3.1, we examine the simulated data and extract the metrics defined and discussed in section 4. Table 1 describes the specification of the simulation parameters. For each recommendation algorithm discussed here, we run our simulation 15 times with these parameters and report the aggregate results.

Our parameter choices are guided by several assumptions about a realistic user-item interaction. We assume that the number of users is much larger than the number of items (i.e., $m \gg n$) allowing for more information about items and facilitating better learning for algorithms.

Figure 3 shows inter-user diversity vs. intra-user diversity for various recommendation algorithms. As shown in this figure, past consumption-based recommendations (consumption, SVD and hybrid) induce significantly weaker filter bubble effects compared to

Parameter	Value
Q (Quality)	$N(\mu_q, \sigma_q^2)$
$\mathcal{G}$ (Item genre)	$N(0, \sigma_g^2)$
$\mathcal{P}$ (User genre preference)	$N(0, \sigma_u^2)$
N_{qual} (Noise in private quality signal)	$N(0, \sigma_{ps}^2)$
N_{genre} (Noise in private genre signal)	$N(0, \sigma_{gs}^2)$
m (Number of users)	1000
k_{init} (Initial items)	10
k_{new} (New items per round)	5
T (Number of rounds)	100
k_{train} (Training worlds)	10
δ (For skewed top pick recommendation)	1
$k_{top}\%$ (For skewed top pick recommendation)	25%

Table 1: Parameter values used in our model. Note: While item genres and user preferences are drawn independently from normal distributions for the results reported in this paper, our results replicated for bimodal distributions of genres and items. We fixed $\mu_q = 100$, $\sigma_q^2 = \sigma_g^2 = \sigma_u^2 = \sigma_{ps}^2 = \sigma_{gs}^2 = 10$; the results are robust to different values of the mean and standard deviation parameters.

no recommendation. In other words, these recommendations alleviate filter bubbles, consistent with previous results. On the other hand, the baseline algorithms that provide accurate genre information (true genre and perfect recommendations) induce stronger filter bubble effects compared to no recommendation.




Figure 3: Inter-user diversity vs. intra-user diversity for the different recommendation algorithms. As shown, the baseline algorithms induce a direct trade-off between the two types of diversity. Past consumption-based recommendation algorithms deviate from this trade-off line primarily by reducing inter-user diversity—they do not significantly affect intra-user diversity.

In particular, this definition of the filter bubble effect is consistent with our general interpretation of the filter bubble effect (i.e. the filter bubble effect is stronger when users with different underlying preferences consume increasingly more different items). We provide empirical evidence for this in appendix A.




Figure 4: Deviation between mean consumed genre (see section 4) and preference vs. preferences for the recommendation algorithms. Compared to no recommendation, past consumption-based recommendations cause large deviations in mean consumed genre towards 0, by pushing all users towards items with near-mode genres.

6.1 Understanding homogenization-filter bubbles dynamic through effects on diversity

As previously mentioned, our central argument is that to fully understand the role of recommendations in the dynamics between homogenization and the filter bubble effect, we need to examine their impact on both inter-user and intra-user diversity. Therefore, we will now investigate how the algorithms in our simulation affect these two facets of diversity.

To do so, we will rely on three key figures. First, figure 3 shows inter-user diversity (Y axis) against intra-user diversity (X axis) for each of the seven algorithms from section 5.1. Second, figure 4 shows the deviation between preference and mean consumed genre for individual users (Y axis) across varying user preferences (X axis) for all seven recommendation algorithms. Finally, figure 5 shows the consumed genre variance for individual users (Y axis) across varying user preferences (X axis) for all seven recommendation algorithms. For the last two figures, the range of possible user preferences is split into multiple bins, each of size 3. Users from each of the 15 iterations are put into one of these bins. We then report the mean of the relevant statistics for each bin.

We use the no recommendation case as our primary baseline. As shown in figure 3, the true genre recommendation achieves higher inter-user diversity and lower intra-user diversity compared to no recommendation, resulting in a stronger filter bubble effect. Without any recommendations, users rely on their personal knowledge of item qualities and genres. When they have accurate information only about item genres, they prioritize affinity over quality and stick to consuming items closer to their preferences. As a result, they deviate the least from their preferences (figure 4) and consume items from a very narrow genre range (figure 5). We observe similar consumption patterns for perfect recommendation, resulting in a stronger filter bubble effect than no recommendation. When users know both item qualities and genres, they can accurately identify high quality items closer to their preferences and consume those.

On the other hand, true quality recommendation lowers inter-user diversity and increases intra-user diversity compared to no recommendation (figure 3), resulting in a weaker filter bubble effect




Figure 5: Consumed genre variance (see section 4) against user preferences for the recommendation algorithms in section 5.1. Compared to no recommendation, past consumption-based recommendations decrease variance for near-mode users and increase variance for niche users by pushing everyone towards blockbuster items.

and weaker homogeneity. When users have accurate knowledge of item qualities but not of item genres, they are more likely to consume high quality items far away from their preferences. As a result, individual mean consumed genre deviates closer to 0 (figure 4), but users consume items from a wider genre range (figure 5).

Meanwhile, past consumption-based recommendations (consumption, SVD, hybrid) rely on past consumption data to learn item quality and are prone to a feedback loop. Since there are more users with near-mode preferences, items with near-mode genres naturally have higher consumption numbers. As a result, these items are favored by past consumption-based recommendation algorithms, which in turn further increases their consumption numbers and continues the loop. These algorithms shift entire user consumption towards the mode of genre distribution rather than widening the range of consumed genres. Therefore, we see large deviations towards 0 in mean consumed genre for niche users (figure 4), while consumed genre variance for individual users do not change much on average compared to no recommendation (figure 5). Consequently, past consumption-based recommendations largely reduce inter-user diversity compared to no recommendation but do not affect intra-user diversity by much (figure 3), resulting in significantly weaker filter bubble effects, and significantly stronger homogeneity.

Combining our observations so far, we can state the following: **past consumption-based recommendations do indeed alleviate filter bubbles, but they do so by greatly reducing inter-user diversity without much effect on intra-user diversity.** Rather than increasing intra-user diversity and exposing users to items from all genres, these recommendations primarily shift the consumption of any individual user towards the mode of the genre distribution, increasing the similarity in consumption between different users.

7 NOVEL RECOMMENDATION ALGORITHMS

Our simulation results demonstrate the importance of considering effects on both inter-user and intra-user diversity when designing recommendation algorithms. Motivated by this insight, next, we propose two novel recommendation algorithms aimed at affecting both inter-user and intra-user diversity.




Figure 6: Inter-user diversity vs. intra-user diversity for different recommendation algorithms. As shown here, binned consumption-based recommendation greatly increases intra-user diversity and reduces inter-user diversity compared to no recommendation by exposing users towards the popular items from each genre. Skewed top pick recommendation with $\delta = 1$ increases inter-user and intra-user diversity simultaneously compared to no recommendation by exposing near-mode users more towards niche items.




Figure 7: Deviation between mean consumed genre (see section 4) and preference vs. preferences for different recommendations. As shown here, binned consumption-based recommendation causes deviations similar to past consumption-based recommendation. Since every user consumes the popular items from each genre, everyone's mean consumed genre gets close to 0. Skewed top pick recommendation with $\delta = 1$ causes less deviation compared to no recommendation. It keeps niche users close to their preferences, while exposing near-mode users to niche items from both sides of the mode.

7.1 Binned consumption-based recommendation.

Our first proposed algorithm, binned consumption-based recommendation, aims to alleviate the filter bubble effect by not only decreasing inter-user diversity but also increasing intra-user diversity. It offers non-personalized curation by pushing users towards items with high consumption numbers relative to the rest of their genre. Formally, we define it as:

Definition (binned consumption-based recommendation). The binned consumption-based recommendation for item i provided to user j is given by $r_i^j(t) = \left(\frac{d_i(t) - \mu}{\sigma} \right)$. Here, $d_i(t)$ is the number

of consumption, and μ and σ are respectively the mean and the standard deviation of the set $\{d_{i'}(t) | g_{i'} = g_i\}$.

Note that since our model assumes continuous real values for item genre, we discretize the set of possible genres in order to use this recommendation algorithm. The intuition here is to eliminate the implicit bias towards blockbuster items by suppressing the genre information, similar to true quality recommendation.

This algorithm reduces inter-user diversity (figure 6) compared to no recommendation because it nudges all users towards the “popular” items (with high consumption compared to the rest of their genre) and shifts individual mean consumed genre towards 0 (figure 7). However, it increases intra-user diversity compared to no recommendation and past consumption-based recommendations (figure 6) since it helps users consume the popular items and increases the genre range they consume from (figure 8). Finally, with low inter-user diversity and high intra-user diversity, this algorithm significantly weakens the filter bubble effect compared to no recommendation.




Figure 8: Consumed genre variance (see section 4) against user preferences for two novel recommendations, as well as for no recommendation, perfect recommendation and hybrid recommendation. binned consumption-based recommendation significantly increases variance for everyone by pushing everyone towards the popular items from each genre. Skewed top pick recommendation increases variance for near-mode users by pushing them towards more niche items.

7.2 Skewed top pick recommendation.

Our second proposed algorithm, skewed top pick recommendation, focuses on simultaneously increasing inter-user and intra-user diversity rather than alleviating the filter bubble effect. It offers preference-centered exploration to particular groups of users. Formally,

Definition (Skewed top pick recommendation). First, each item i is ranked according to $(q_i^j \cdot |g_i|^\delta)$ in descending order. Then, item i is recommended if it is in the top $k_{top}\%$ in this ranking. More precisely, $r_i^j(t) = 1$ if item i is in the top $k_{top}\%$ of this ranking, and $r_i^j(t) = 0$ otherwise.

Depending on δ , this recommendation is skewed towards either niche items or items with near-mode genres. With $\delta = 1$, this recommendation algorithm nudges near-mode users more towards niche items, but does not significantly affect their mean consumed genre

since they consume niche items from both sides of the mode. Niche users however stick close to their original preferences. As a result, we see small deviations in mean consumed genre from preferences similar to true genre recommendation (figure 7), and a significant increase in the genre range near-mode users consume from (figure 8). Consequently, this recommendation increases both inter-user and intra-user diversity (figure 6) and causes a stronger filter bubble effect and weaker homogeneity compared to no recommendation.

While not reported here, we did simulate this algorithm for different values of δ . Increasing δ means that the algorithm will push users more towards niche items. For sufficiently large δ , users begin to ignore the algorithm. As a result, we observe consumption patterns similar to the no recommendation scenario.

In addition to the effects of these two novel algorithms on user consumption patterns, we also look at their effects on user utility compared to other recommendation algorithms. These comparisons are provided in appendix D.

8 DISCUSSION AND FUTURE DIRECTIONS

We proposed a novel agent-based simulation study to investigate the effects of a select set of recommendation algorithms on user consumption patterns. We developed more refined definitions of the filter bubble and the homogenization effects of recommendations, which decomposed the effect into two components: inter-user diversity and intra-user diversity. Our simulation results show that past consumption-based recommendations only reduce inter-user diversity when alleviating the filter bubble effect—their impact on intra-user diversity is not significant. We then define and examine two novel recommendation algorithms: *binned consumption-based recommendation*, which provides a non-personalized curated set of content and thus significantly increases intra-user diversity while reducing collective diversity, and *skewed top pick recommendation*, which facilitates preference-centered exploration and thus increases inter-user and intra-user diversity simultaneously.

We do not advocate for either minimizing or maximizing the filter bubble effect; we believe such judgments should be made on a case-by-case basis. Instead, our aim is to enable discussions on which strategy is more suitable given a particular context, through our decomposition-based framework. For instance, a news provider may wish to synchronize user perspectives (reduced filter bubble) while also offering them diverse viewpoints (increased intra-user diversity). As demonstrated by our findings, traditional consumption-based recommendations fall short in this regard.

The scope of our current work leaves ample opportunities for future research. Our existing model presupposes that content items are exogenously generated. However, content producers play a crucial role in online ecosystems, influencing the available item pool. Moreover, our model assumes a single, neutral platform, whereas, in practice, multiple platforms, each with distinct objectives, may vie for the attention of users and producers. For example, content creators might migrate to platforms that better serve their genre, leading to genre-specific platforms (e.g., Twitch for live streams, YouTube for long-form videos, TikTok for short clips). Therefore, a logical next step in our research could be to incorporate all three types of agents—users, platforms, and producers—and examine the effects of homogenization and filter bubbles in such a multifaceted ecosystem.

REFERENCES

- [1] Guy Aridor, Duarte Goncalves, and Shan Sikdar. 2020. Deconstructing the Filter Bubble: User Decision-Making and Recommender Systems. In *Proceedings of the 14th ACM Conference on Recommender Systems* (Virtual Event, Brazil) (RecSys '20). Association for Computing Machinery, New York, NY, USA, 82–91. <https://doi.org/10.1145/3383313.3412246>
- [2] Eytan Bakshy, Solomon Messing, and Lada A. Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132. <https://doi.org/10.1126/science.aaa1160> arXiv:<https://www.science.org/doi/pdf/10.1126/science.aaa1160>
- [3] A. Bruns. 2019. Filter bubble. <https://doi.org/10.14763/2019.4.1426>
- [4] Pablo Castells, Neil Hurley, and Saúl Vargas. 2022. *Novelty and Diversity in Recommender Systems*. Springer US, New York, NY, 603–646. https://doi.org/10.1007/978-1-0716-2197-4_16
- [5] Allison J. B. 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* (Vancouver, British Columbia, Canada) (RecSys '18). Association for Computing Machinery, New York, NY, USA, 224–232. <https://doi.org/10.1145/3240323.3240370>
- [6] Seth Flaxman, Sharad Goel, and Justin M. Rao. 2016. Filter Bubbles, Echo Chambers, and Online News Consumption. *Public Opinion Quarterly* 80, S1 (03 2016), 298–320. <https://doi.org/10.1093/poq/nfw006> arXiv:<https://academic.oup.com/poq/article-pdf/80/S1/298/17120810/nfw006.pdf>
- [7] Daniel Fleder and Kartik Hosanagar. 2009. Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity. , 697–712 pages. <https://doi.org/10.1287/mnsc.1080.0974>
- [8] Daniel Geschke, Jan Lorenz, and Peter Holtz. 2019. The triple-filter bubble: Using agent-based modelling to test a meta-theoretical framework for the emergence of filter bubbles and echo chambers. *British Journal of Social Psychology* 58, 1 (2019), 129–149. <https://doi.org/10.1111/bjso.12286> arXiv:<https://bpspsychub.onlinelibrary.wiley.com/doi/pdf/10.1111/bjso.12286>
- [9] 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.
- [10] Jonas Kaiser and Adrian Rauchfleisch. 2020. Birds of a Feather Get Recommended Together: Algorithmic Homophily in YouTube's Channel Recommendations in the United States and Germany. *Social Media + Society* 6, 4 (2020), 2056305120969914. <https://doi.org/10.1177/2056305120969914> arXiv:<https://doi.org/10.1177/2056305120969914>
- [11] Matevž Kunaver and Tomaž Požrl. 2017. Diversity in recommender systems – A survey. *Knowledge-Based Systems* 123 (2017), 154–162. <https://doi.org/10.1016/j.knosys.2017.02.009>
- [12] Anja Lambrecht and Catherine Tucker. 2016. Algorithmic Bias? An Empirical Study into Apparent Gender-Based Discrimination in the Display of STEM Career Ads. *SSRN Electronic Journal* (01 2016). <https://doi.org/10.2139/ssrn.2852260>
- [13] Masoud Mansoury, Himan Abdollahpouri, Mykola Pechenizkiy, Bamshad Mobasher, and Robin Burke. 2020. Feedback Loop and Bias Amplification in Recommender Systems. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management* (Virtual Event, Ireland) (CIKM '20). Association for Computing Machinery, New York, NY, USA, 2145–2148. <https://doi.org/10.1145/3340531.3412152>
- [14] Kari Karpainen Natali Helberger and Lucia D'Acunto. 2018. Exposure diversity as a design principle for recommender systems. *Information, Communication & Society* 21, 2 (2018), 191–207. <https://doi.org/10.1080/1369118X.2016.1271900> arXiv:<https://doi.org/10.1080/1369118X.2016.1271900>
- [15] 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* (Seoul, Korea) (WWW '14). Association for Computing Machinery, New York, NY, USA, 677–686. <https://doi.org/10.1145/2566486.2568012>
- [16] Derek O'Callaghan, Derek Greene, Maura Conway, Joe Carthy, and Pádraig Cunningham. 2015. Down the (White) Rabbit Hole: The Extreme Right and Online Recommender Systems. *Social Science Computer Review* 33, 4 (2015), 459–478. <https://doi.org/10.1177/0894439314555329> arXiv:<https://doi.org/10.1177/0894439314555329>
- [17] Eli Pariser. 2011. *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin.
- [18] Matthew J. Salganik, Peter Sheridan Dodds, and Duncan J. Watts. 2006. Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science* 311, 5762 (2006), 854–856. <https://doi.org/10.1126/science.1121066> arXiv:<https://www.science.org/doi/pdf/10.1126/science.1121066>
- [19] Jessica Su, Aneesh Sharma, and Sharad Goel. 2016. The Effect of Recommendations on Network Structure. In *Proceedings of the 25th International Conference on World Wide Web* (Montréal, Québec, Canada) (WWW '16). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1157–1167. <https://doi.org/10.1145/2872427.2883040>

A COMPARISON WITH GENERAL INTERPRETATION OF FILTER BUBBLES

As discussed previously, the general interpretation of the filter bubble effect is that "the filter bubble effect is stronger when people with different underlying preferences consume increasingly more different items". Within our simulation model, we can represent this interpretation by measuring the total pairwise distance in genre between the items consumed respectively by two users on average. The higher this measure is, the stronger the filter bubble effect is for the corresponding recommendation algorithm.




Figure 9: Total pairwise distance in genre between the respective items consumed by two users, vs. the distance between their respective preferences. The higher the position of a curve, the higher the mean total pairwise distance in consumed genre for the corresponding recommendation, and the stronger the filter bubble effect according to the general interpretation.




Figure 10: Filter bubble effect (see section 4) for each of the seven recommendation algorithms. As shown here, past consumption-based algorithms significantly decrease the filter bubble effect compared to our baseline algorithms, while true genre recommendation significantly increases it.

Figure 9 shows the total pairwise distance in consumed genre (Y axis) between two users as we vary the distance in preference (X axis) between the users. The range of possible distances in user preferences is split into multiple bins, each of size 3. Every pair of users from each of the 15 iterations are placed into one of these

bins based on the distance between them. We then report the mean total pairwise distance in consumed genre for each bin. Each curve in this figure corresponds to a different recommendation algorithm from section 5.1, as identified in the legend.

According to the general interpretation, the higher the position of a curve is in figure 9, the stronger the filter bubble effect is for the corresponding recommendation algorithm. This allows us to rank these algorithms in descending order of the strength of the filter bubble effect. We can also rank these algorithms in descending order of our definition of the filter bubble effect (shown in figure 10). We can then verify that the two rankings are exactly the same. This implies that our definition of the filter bubble effect, constructed using inter-user and intra-user diversity, is consistent with the general interpretation of the filter bubble effect.

B ALTERNATIVE OPERATIONALIZATION OF HOMOGENEITY

In section 4, we provided a definition for homogeneity that involves both inter and intra-user diversity, namely:

$$\frac{1}{\sqrt{\text{inter-user diversity}^2 + \text{intra-user diversity}^2}}$$

As an alternative, we can also use the following definition:

$$\frac{1}{\text{inter-user diversity} + \text{intra-user diversity}}$$

Figure 11 demonstrates this alternative definition against the natural measure of homogeneity mentioned in section 4. The Pearson correlation coefficient for these two measures is 0.99992089, i.e., they are highly correlated, justifying the use of this definition.




Figure 11: $1/(\text{inter-user diversity} + \text{intra-user diversity})$ against the inverse of the standard deviation of all consumption for each of the seven recommendation algorithms. The Pearson correlation coefficient between the two values is 0.99992089, i.e. they are highly correlated.

C SIMULATION IMPLEMENTATION DETAILS

In section 5, we briefly described how we leverage a two-phase simulation process to learn the regression model users use to make their decisions. Detailed description of how each of these two phases function are provided below.

C.1 Learning users' utility estimation model

During the training phase, we initiate k_{train} training worlds $\mathcal{W}^\ell$, $\forall \ell = 1, 2, \dots, k_{train}$. Each training world simulates the interaction between the users from the true world $\mathcal{W}$ and a set of items similar to those from $\mathcal{W}$. This works as a proxy for the past interactions of users with items, and allows us to learn the regression model used by users in decision making. In particular, we have the following:

- Each training world $\mathcal{W}^\ell$ has the same set of users as $\mathcal{W}$, because we want to learn the regression model used by these particular users.
- Each training world has n items: item qualities and genres in each training world $\mathcal{W}^\ell$ are drawn from the same distributions $\mathcal{Q}, \mathcal{G}$ as $\mathcal{W}$. This is because users learn how to interpret various signals about an item (i.e. the regression model) from their past consumption of similar distributions of items.

Therefore, mathematically, each training world can be defined as a collection of m users and necessary distributions:

$$\mathcal{W}^\ell = (m, k_{init}, k_{new}, T, \{p_1, p_2, \dots, p_m\}, \mathcal{Q}, \mathcal{G}, \mathcal{N}_{qual}, \mathcal{N}_{genre})$$

$$\forall \ell = 1, 2, \dots, k_{train}$$

Procedure Learning phase pseudocode

```

1 Initialize  $\mathcal{W}, \mathcal{W}^\ell \forall \ell = 1, 2, \dots, k_{train}$ , each with  $k_{init}$  items;
2 Initialize  $m$  users, shared across all the simulation worlds;
3 begin Training phase
4   for round  $t = 1, 2, \dots, T$  do
5     Construct private signals  $q_i^{j\ell}$  and recommendation signals  $r_i^{j\ell} \forall$  world  $\ell$ , user  $j$ , available item  $i$  and standardize 2;
6     Construct  $X = \{\mathbf{x}_{jit}\}_{j,i,\ell}$  where
        $\mathbf{x}_{jit} = [q_i^{j\ell}, |p_j - g_i^{j\ell}|, r_i^{j\ell}]$ ,  $Y = \{U^\ell(j, i)\}_{j,i,\ell}$ 
        $\forall$  world  $\ell$ , user  $j$ , available item  $i$ ;
7     Learn new  $\mathbf{w}_0^s, \mathbf{w}^s$  from  $(X, Y)$ ;
8     for training world  $\ell = 1, 2, \dots, k_{train}$  do
9       Add  $k_{new}$  items to  $\mathcal{W}^\ell$ ;
10      Construct private signals  $q_i^{j\ell}$  and recommendation signals  $r_i^{j\ell}$ 
         $\forall$  world  $\ell$ , user  $j$ , available item  $i$  and standardize;
11      for user  $j = 1, 2, \dots, m$  do
12        User  $j$  predicts utility for each available item
           $i: \hat{U}^\ell(j, i) = \mathbf{w}_0^s + \mathbf{w}^{s\top} \mathbf{x}_{jit}$  where
           $\mathbf{x}_{jit} = [q_i^{j\ell}, |p_j - g_i^{j\ell}|, r_i^{j\ell}]$ ;
13        User  $j$  chooses unconsumed item
           $i = \arg \max_i \{\hat{U}^\ell(j, i)\}$ ;
14      Update the consumption numbers of each item  $i$ ;

```

Each training world $\mathcal{W}^\ell$ also has T discrete rounds of progression, similar to our true simulated world $\mathcal{W}$. Each round in a training world progresses similarly to the true simulated world $\mathcal{W}$, except for the following additional mechanism:

At the beginning of each round s , consumption data from all k_{train} training worlds is aggregated in order to construct the following training data:

- (1) Set X of feature vectors: one feature vector $\mathbf{x}_{jit}$ for each triplet of user j , available item i and training world $\mathcal{W}^\ell$. The features in each vector are the signals available to user j about item i in training world $\mathcal{W}^\ell$:
 - (a) private quality signal $q_i^{j\ell}$ of user j about the quality of item i in world $\mathcal{W}^\ell$
 - (b) perceived distance between the preference of user j and the genre of item i in training world $\mathcal{W}^\ell$, $|p_j - g_i^{j\ell}|$
 - (c) recommendation $r_i^{j\ell}$ about item i for user j in training world $\mathcal{W}^\ell$ provided by the system
 Formally, $X = \{\mathbf{x}_{jit}\}_{j,i,\ell}$ where $\mathbf{x}_{jit} = [q_i^{j\ell}, |p_j - g_i^{j\ell}|, r_i^{j\ell}]$.
- (2) Set Y of target values: the true utility user j would receive from item i in world $\mathcal{W}^\ell$, namely, $U^\ell(j, i) = q_i^\ell - |p_j - g_i^\ell|$. Formally, $Y = \{U^\ell(j, i)\}_{j,i,\ell}$.

From (X, Y) , we learn a new regression model (i.e., $\mathbf{w}_0^s, \mathbf{w}^s$), which users then use to estimate item utilities and choose their consumption throughout the remainder of round s in each training world. A pseudocode of this phase is provided in Learning phase pseudocode.

C.2 Simulating recommendations

After the final round of the training phase, we end up with estimates $\hat{\mathbf{w}}_0 = \mathbf{w}_0^T, \hat{\mathbf{w}} = \mathbf{w}^T$ of $\mathbf{w}_0, \mathbf{w}$, and are ready to run the simulation in the "true" simulation world $\mathcal{W}$. At each round s , each user j estimates the utility of each available item i , and chooses to consume the item with the maximum estimated utility. A pseudocode of this phase is provided in Simulation phase pseudocode.

Procedure Simulation phase pseudocode

```

1 begin Deployment phase
2   for round  $t = 1, 2, \dots, T$  do
3     Add  $k_{new}$  items to  $\mathcal{W}$ ;
4     Construct private signals  $q_i^j$ 
        $\forall$  user  $j$ , available item  $i$  and standardize;
5     Construct recommendation signals  $r_i^j$ 
        $\forall$  user  $j$ , available item  $i$  and standardize;
6     for user  $j = 1, 2, \dots, m$  do
7       User  $j$  predicts utility for each available item  $i$ :
          $\hat{U}(j, i) = \hat{\mathbf{w}}_0 + \hat{\mathbf{w}}^\top \mathbf{x}_{ji}$  where
          $\mathbf{x}_{ji} = [q_i^j, |p_j - g_i^j|, r_i^j]$ ;
8       User  $j$  chooses unconsumed item
          $i = \arg \max_i \{\hat{U}(j, i)\}$ ;
9     Update the consumption numbers of each item  $i$ ;

```

D UTILITY COMPARISON

In addition to the effects on inter-user and intra-user diversity, we can also observe the effects of these two novel recommendations on user utility. Figure 12 shows mean affinity component, $-|p_j - g_i|$, of individual utility against mean quality component, q_i (along with corresponding standard deviations) for the two novel recommendations, as well as for the representative baseline and past consumption-based recommendations.

As demonstrated in this figure, binned consumption recommendation achieves utility similar to that of past consumption-based recommendations. Since this recommendation pushes users towards the same items, i.e., items with high consumption numbers relative to the rest of their genre, it performs worse compared to no recommendation for affinity, but achieves significantly higher quality. On the other hand, skewed top pick recommendation achieves utility similar to that of no recommendation. This recommendation focuses on driving near-mode users to explore niche items, and does significantly worse in terms of the quality component compared to past consumption-based recommendations and binned consumption recommendation.




Figure 12: Breakdown of mean individual utility into quality and affinity components for two novel recommendations, as well as for no recommendation, perfect recommendation and hybrid recommendation. As shown here, binned consumption recommendation performs similarly to past consumption-based recommendations. On the other hand, since skewed top pick recommendation prioritizes exploration over anything else, its performance is significantly worse, and comparable to no recommendation.