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

Guy Aridor¹ Duarte Gonçalves¹ Shan Sikdar²

Abstract

We study a model of user decision-making in the context of recommender systems via numerical simulation. We show that this model provides an explanation for the findings of Nguyen, et. al (2014), where, in environments 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 select 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 (MacKenzie et al., 2013). 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 (Pariser, 2011), and that personalized RS may also lead users to become increasingly homogenized at the same time (Chaney et al., 2018; Hosanagar et al., 2013).

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 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 in order 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, prior to consuming them, users are *uncertain* about how much they value the different items. This is motivated both by the fact that, in many contexts, recommendations regard 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

¹Department of Economics, Columbia University, New York, NY ²Everquote, Cambridge, MA. Correspondence to: Guy Aridor <g.aridor@columbia.edu>.

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 (Schulz et al., 2019).

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 (Nguyen et al., 2014) 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 tradeoff 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.

Related Work. The first set of related works studies the extent and implications of filter bubbles. (Pariser, 2011) 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 (Flaxman et al., 2016; Hosanagar et al., 2013; Möller et al., 2018; Nguyen et al., 2014). The most relevant to our study is (Nguyen et al., 2014) 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. Furthermore, 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. (Celma & Cano, 2008; Treviranus & Hockema, 2009) show that incorporating content popularity into RS can lead to increased user homogenization. (Chaney et al., 2018) shows how user homogenization may arise from training RS on data from users exposed to algorithmic recommendations. (Fleder & Hosanagar, 2009) 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.

Another set of papers studies the impact of human decision-making on the design and evaluation of RS. (Chen et al., 2013) 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 (Bettman et al., 1998; Lichtenstein & Slovic, 2006) whereby users develop preferences over time through the context of a decision process. 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, (Celma & Herrera, 2008; Cremonesi et al., 2013; Pu et al., 2011) 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.

2. Model Illustration

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 beliefs are given as follows where ρ denotes the correlation between the utilities of the items, μ is the mean belief, and Σ is the covariance matrix:¹

$$\mu = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \quad \Sigma = \sigma^2 \begin{pmatrix} \rho^0 & \rho^1 & \rho^2 & \rho^0 \\ \rho^1 & \rho^0 & \rho^1 & \rho^2 \\ \rho^2 & \rho^1 & \rho^0 & \rho^1 \\ \rho^1 & \rho^2 & \rho^1 & \rho^0 \end{pmatrix}$$

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 3 items.² For concreteness, we suppose that $\sigma = 1$ and $\rho = 0.5$, but the intuitions hold for any value of σ and

¹The prior beliefs are therefore given by $\mathcal{N}(\mu, \Sigma)$ and the realized utilities are drawn from this distribution.

²The 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

$\rho > 0$. 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:

$$\bar{\mu} = \begin{pmatrix} \mathbb{E}[x_1] \\ \mathbb{E}[x_2] \\ \mathbb{E}[x_3] \end{pmatrix} = \begin{pmatrix} \rho y \\ \rho^2 y \\ \rho y \end{pmatrix} = \begin{pmatrix} \frac{1}{4} \\ \frac{1}{8} \\ \frac{1}{4} \end{pmatrix}$$

$$\bar{\Sigma} = \begin{pmatrix} \frac{3}{4} & \frac{3}{8} & 0 \\ \frac{3}{8} & \frac{15}{16} & \frac{3}{8} \\ 0 & \frac{3}{8} & \frac{3}{4} \end{pmatrix}$$

Thus, upon learning $x_0 = y$, the user updates beliefs about the remaining items. Note that $\mathbb{E}[x_1] = \mathbb{E}[x_3] > \mathbb{E}[x_2]$ since item 0’s value is more informative about similar items’ values, items 1 and 3, than items further way 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 $\mathbb{E}[x_1] = \mathbb{E}[x_3] = -\frac{1}{4} < -\frac{1}{8} = \mathbb{E}[x_2]$ 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 than item 2. In

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:

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

Thus, 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:

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

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

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.

3. Recommender System Design

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 (Adomavicius & Tuzhilin, 2005).

There has been a recent movement away from such evaluation measures due to the observation that accurate recommendations are not necessarily useful recommendations (McNee et al., 2006). 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 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 the type of data that 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.³ The second type of data that is relevant for RS 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. 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 would the user 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 (Kotkov et al., 2016), calibration (Steck, 2018), coverage (Ge et al., 2010), novelty (Vargas & Castells, 2011), and many others. Our approach most closely follows the set of proposed serendipity measures which are surveyed in (Kotkov et al., 2016). As discussed by (Maksai et al., 2015), 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, (Vargas & Castells, 2011; Kaminskas & Bridge, 2014) propose unexpectedness metrics that look at the dissimilarity of the proposed recommended items compared to what the recommender already

³Additionally, user beliefs contain information that may not have been observed by the recommender that only observes user choices on the platform.

knows the user likes - either via a content-based approach or a collaborative-based approach. 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.

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 recommender to understand how “unexpected” a recommendation would be from the perspective of a user.

References

- Adomavicius, G. and Tuzhilin, A. 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):734–749, 2005.
- Bettman, J. R., Luce, M. F., and Payne, J. W. Constructive consumer choice processes. *Journal of consumer research*, 25(3):187–217, 1998.
- Celma, Ò. and Cano, P. 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*, pp. 5. ACM, 2008.
- Celma, Ò. and Herrera, P. A new approach to evaluating novel recommendations. In *Proceedings of the 2008 ACM conference on Recommender systems*, pp. 179–186, 2008.
- Chaney, A. J., Stewart, B. M., and Engelhardt, B. E. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM Conference on Recommender Systems*, pp. 224–232. ACM, 2018.
- Chen, L., De Gemmis, M., Felfernig, A., Lops, P., Ricci, F., and Semeraro, G. Human decision making and recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiIS)*, 3(3):1–7, 2013.
- Cremonesi, P., Garzotto, F., and Turrin, R. User-centric vs. system-centric evaluation of recommender systems. In *IFIP Conference on Human-Computer Interaction*, pp. 334–351. Springer, 2013.
- Flaxman, S., Goel, S., and Rao, J. M. Filter bubbles, echo chambers, and online news consumption. *Public opinion quarterly*, 80(S1):298–320, 2016.
- Fleder, D. and Hosanagar, K. Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity. *Management science*, 55(5):697–712, 2009.
- Ge, M., Delgado-Battenfeld, C., and Jannach, D. Beyond accuracy: evaluating recommender systems by coverage and serendipity. In *Proceedings of the fourth ACM conference on Recommender systems*, pp. 257–260. ACM, 2010.
- Hosanagar, K., Fleder, D., Lee, D., and Buja, A. Will the global village fracture into tribes? recommender systems and their effects on consumer fragmentation. *Management Science*, 60(4):805–823, 2013.
- Kaminskas, M. and Bridge, D. 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, 2014.
- Kotkov, D., Wang, S., and Veijalainen, J. A survey of serendipity in recommender systems. *Knowledge-Based Systems*, 111:180–192, 2016.
- Lichtenstein, S. and Slovic, P. *The construction of preference*. Cambridge University Press, 2006.
- MacKenzie, I., Meyer, C., and Noble, S. *How retailers can keep up with consumers.*, 2013. URL <https://www.mckinsey.com/industries/retail/our-insights/how-retailers-can-keep-up-with-consumers>.
- Maksai, A., Garcin, F., and Faltings, B. Predicting online performance of news recommender systems through richer evaluation metrics. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pp. 179–186. ACM, 2015.
- McNee, S. M., Riedl, J., and Konstan, J. A. Being accurate is not enough: how accuracy metrics have hurt recommender systems. In *CHI'06 extended abstracts on Human factors in computing systems*, pp. 1097–1101. ACM, 2006.
- Möller, J., Trilling, D., Helberger, N., and van Es, B. 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):959–977, 2018.
- Nguyen, T. T., Hui, P.-M., Harper, F. M., Terveen, L., and Konstan, J. A. Exploring the filter bubble: the effect of using recommender systems on content diversity. In *Proceedings of the 23rd international conference on World wide web*, pp. 677–686. ACM, 2014.

- Pariser, E. *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin, 2011.
- Pu, P., Chen, L., and Hu, R. A user-centric evaluation framework for recommender systems. In *Proceedings of the fifth ACM conference on Recommender systems*, pp. 157–164, 2011.
- Schulz, E., Bhui, R., Love, B. C., Brier, B., Todd, M. T., and Gershman, S. J. Structured, uncertainty-driven exploration in real-world consumer choice. *Proceedings of the National Academy of Sciences*, pp. 201821028, 2019.
- Steck, H. Calibrated recommendations. In *Proceedings of the 12th ACM conference on recommender systems*, pp. 154–162, 2018.
- Treviranus, J. and Hockema, S. 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, sep 2009. doi: 10.1109/tic-sth.2009.5444430.
- Vargas, S. and Castells, P. Rank and relevance in novelty and diversity metrics for recommender systems. In *Proceedings of the fifth ACM conference on Recommender systems*, pp. 109–116. ACM, 2011.