Recommender systems and the value of user data*

Gunhaeng Lee† Julian Wright‡

September 18, 2021

Abstract

Many online platforms and other technology-enabled firms including the likes of Amazon, Goodreads, LinkedIn, Netflix, StitchFix, Tinder and YouTube make use of some type of recommender system to help their users discover relevant products, people, or content, thereby increasing the value of the services they offer. In this paper we model how a recommender system works, and use the model to investigate, both theoretically and empirically, what is the value provided by user data. Value is created in our setting by customization, selection and screening. We determine how the value of data varies with the number of target items available for the user, the number of users that data is gathered from, the number of features that are taken into account in the customized prediction, and the degree of misalignment between the platform's and user's interests.

Keywords: collaborative filtering, machine learning, big data.

1 Introduction

Recommender systems which provide users with customized recommendations on items to try have become widely used by tech firms, both big and small, and represent one of the most important applications of big data and machine learning in the economy. Despite their widespread adoption, there is surprisingly little economic analysis of how they work and what value they create. In this paper, we build a model of a recommender system that makes customized recommendations of items to users based on their and other users' past feedback (i.e. it uses collaborative filtering). Our objective is to try to quantify the value created by such recommender systems, and dig into the sources of the value creation and the shape of the resulting learning curve. We do this first using theory and then based on data.

*We thank Ismael Moreno Martinez, Maximilian Schäfter, Juwon Seo and Tat-How Teh, as well as seminar and conference participants at NUS, SAET2021 and APIOC2021. All errors are ours.

†Department of Economics, National University of Singapore ecsgl@nus.edu.sg

‡Department of Economics, National University of Singapore jwright@nus.edu.sg

In the model there is a set of observable features which a platform has user feedback on. We will focus on the case that the user feedback takes the form of ratings on items, but the setting can accommodate other features such as user profiles or users' browsing history as long as there are a finite number of possibilities. The platform makes a recommendation to a new user about which item to try (possibly none) from a set of items that the user hasn't tried before. We call these items target items. To do so, the system learns, using the ratings from existing users, how the items, including ratings on the target items, are correlated with each other, and subsequently estimates the probability that the new user will like each target item based on the user's ratings on other items. Based on the learning, the system selects an item with the highest estimated probability and recommends it to the user if the probability is above a certain threshold. Once the item is recommended to the user, the user decides whether to try it or not. Users get a certain positive payoff if they try and like the item, a certain negative payoff if they try and dislike the item, and a zero payoff if they don't try the item. We focus on a class of recommender systems that use a statistically consistent estimator for the true probability the user likes the target item.

Our paper provides three main sets of results. The first of these is to quantify the value that our class of recommender systems creates and explore how this varies with the number of target items available for the user, the number of users that data is gathered from, the number of items that are taken into account in the prediction, and the threshold adopted for making recommendations. We do this by decomposing the value created by the recommender system into three components according to the value creation process: *customization*, *selection* and *screening*. The system selects and screens items for users, and more detailed customization helps by improving the precision offered by these two functions. Intuitively, data from a larger user base enhances learning about the correlation structure among the items. Based on the learning about the correlation between items, customizing recommendations contributes to the value created for users by segmenting users into groups of similar historical feedback, and then by making group-specific recommendations. This allows the platform to select the best item for each group and, at the same time, to screen out items deemed not suitable for the group. After presenting the value and its decomposition, we show that when the threshold level is optimally chosen for users, additional customization, i.e., finer segmentation of users, always (weakly) increases the value to users, and we provide a necessary and sufficient condition under which the marginal value of customization is strictly positive. However, when the threshold is not aligned with users' interest, we show how there is scope for customization (or more customization) to strictly hurt users. Finally, we show that with one additional mild assumption on the class of recommender systems considered and under our benchmark payoff specification, there is a positive but diminishing data network effect. Specifically, the marginal value to a new user of having more other users (to train the system on) is always positive, but diminishing.

Our second main contribution is to propose a specific Bayesian recommender system that imposes few constraints on the learning process, yet is amenable to theoretical and empirical analysis.

This involves Bayesian learning from data about the unconstrained correlation structure of multiple items and Bayesian estimation of the conditional probability that a user will like each of the target items based on the user's ratings on other items. Our probabilistic approach has several advantages over traditional methods such as matrix factorization or the nearest neighbor approach. Notably, we are able to present a closed-form representation of the estimated probability instead of an algorithmic representation. This not only allows our estimation to be more easily interpreted but also enables us to take our theory to the data and consider a variety of counterfactual experiments.

Our third main contribution is to empirically quantify the level and statistical significance of the theoretical results. To do so, we use a publicly available dataset containing over four million anonymous joke ratings from 73,421 users, and consider several counterfactual experiments. First, for each target item, we evaluate and compare the value created for users under three different scenarios: users trying an item without receiving any information from a recommender system, deciding whether to try the item after reviewing the average rating of the item from past users, and following the customized recommendations of the system we characterize. We find significant improvements in user value when switching from the first scenario to the second (30.4%), and also from the second scenario to the third (33.7%). The value addition is even more significant when the recommender systems can choose from multiple target items.[1] We observe an increase of 132.4% in the average utility when switching from the first scenario to the second and an increase of 7.8% when switching from the second scenario to the third. Second, after first showing that the marginal value of additional users and the marginal value of customization are both positive but diminishing in our data, we assess the extent of complementarity or substitutability between these two dimensions of learning. We find the number of previous users exhibits strong complementarity with the degree of customization when both the number of users and the degree of customization is low, but this complementarity disappears quickly as the system accumulates data from more users. Finally, we estimate the users' utility loss due to the possible misalignment of interests between the platform that collects user data and its users, and show how this misalignment of interests can reduce the value created from customization.

The rest of the paper proceeds as follows. In Section 2 we survey the related literature. In Section 3 we define a recommender system and develop a theory of value created by this recommender system. In Section 4 we introduce the specific Bayesian learning and prediction model, which we apply to data in Section 5 to obtain our empirical findings. Section 6 explores a few possible extensions of our framework, while Section 7 briefly concludes.

2 Related Literature

There is by now an extensive literature in computer science that proposes and compares different recommender systems. Lu et al. (2015) provides a survey. Rather than finding the recommender

[1] We consider the situation that the systems choose one item from three target items.

system that works best for a particular application, we are primarily interested in understanding the implications of a fairly generic recommender system for the value of data (including evaluating the value of customization in recommendations and the misalignment of interests between platforms and users). Thus, we focus on a large class of recommender systems based on collaborative filtering that lends itself to theoretical analysis. We also develop a Bayesian recommender system that belongs to this class to empirically explore new questions that the computer science literature has not addressed.

Our paper differs from the substantial economics literature on learning from past users, most of which either assume users' values over the items are independent (e.g. Gittins (1979), Weitzman (1979) Olszewski and Weber (2015), Papanastasiu et al. (2018) and Wolitzky (2018)) or assumes users have perfectly correlated values (e.g. the literature on social learning and herding, such as the work of Banerjee (1992), Smith and Sørensen (2000) and Acemoglu et al. (2011)). Of particular note is Kremer et al. (2014) and Che and Hörner (2018) who study the recommendations of a designer that maximizes social welfare. They take a mechanism design approach in which the designer incentivises users to try a product by manipulating their belief in order to promote exploration of the product (or multiple uncorrelated products), which benefits later users. In contrast, the learning in our study takes place over multiple products that are correlated with each other according to some unknown structure, but for most of our analysis we shut down the exploration motive which is the main focus of Kremer et al. (2014) and Che and Hörner (2018).

Another strand of literature explores how user data and machine learning, including recommender systems, shape market outcomes. For instance, Bergemann and Ozmen (2006) uses a two-period model to study how the adoption of a recommender system by a seller affects its pricing when it competes with a competitive fringe of sellers that do not have a recommender system. On the other hand, Calvano et al. (2020) focuses on reinforcement learning and studies how algorithmic pricing adopted by competitive firms can lead to a price supported by a tacit collusion. Biglaiser et al. (2019), de Cornière and Taylor (2020), Farboodi et al. (2019), Hagiu and Wright (2021) and Prüfer and Schottmüller (2021) consider how equilibrium outcomes between competing firms are affected when firms can improve their offerings by learning from past customer data, as is the case under a recommender system. Bergemann et al. (2021) and Ichihashi (2020a) consider the problem faced by a data intermediary and the resulting market outcome when data from a customer contains relevant information for others. Unlike these articles, we do not consider the market implications of recommender systems, and indeed abstract from any prices being charged to consumers.[2] Rather, we focus on the recommender system itself and study how it and the richness of the data it collects together create value.

At the same time, some authors have empirically investigated the economic returns to data in terms of forecast accuracy (Bajari et al. (2019), Claussen et al. (2021), Chiou and Tucker (2017)

[2]For many of the firms we have in mind, they either do not charge consumers, or if they do, they charge on a subscription basis rather than being paid based on whether the consumer follows the recommendation or not.

and Schaefer and Sapi (2021)). The most related is Schaefer and Sapi (2021). Testing the impact of data size on search quality, they highlight the importance of both the volume of previous searches on the same keyword (across-user learning) and the search histories of individual users (within-user learning). It is found that both factors contribute to higher search quality, and that they complement each other. Relative to this burgeoning literature, our paper provides the microfoundations for the learning curves underlying these types of approaches, showing how learning benefits users in terms of user surplus rather than just the accuracy of the recommendations. It measures the value of data both theoretically and empirically within a unified framework, which allows us to conduct counterfactual experiments as different aspects of the recommender system are varied.

In our empirical analysis, the Bayesian recommender system we develop is closely related to Chien and George (1999), which also takes a Bayesian approach in making predictions. Among the commonly used recommendation techniques involving collaborative filtering (Sarwar et al. (2001), Schafer et al. (2007) and Elahi et al. (2016)), content-based filtering (Pazzani and Billsus (2007) and Lops et al. (2011)) and knowledge-based filtering (Burke (2000) and Aggarwal (2016)), the probability-based Bayesian collaborative filtering approach we adopt has the advantage of tractability. Building upon a probabilistic model, we obtain closed-form representations of the predictions. While Chien and George (1999) introduces a prediction mechanism based on the users' similarities (often referred to as user-based collaborative filtering), our approach focuses on the possible correlations among items (often referred to as item-based collaborative filtering). In our approach, we partition the users into subgroups according to their historical ratings, and then try to measure similarities between the rated items and each of the target items for each subgroup.

Finally, note that throughout the paper we abstract from any privacy concerns of the firm's collection of data, which is the subject of another important literature related to user data (e.g. see Choi et al. (2019) and Acemoglu et al. (2021)).

3 Value of data

3.1 The model

A firm (which we refer to as a platform to reflect that most of the examples we have in mind are digital platforms) has $I \geq 1$ items to consider recommending to a target user and chooses an item to recommend to the user who can have either a positive or a negative experience with the item. Each of the I items is called a target item and we denote the set of target items by $\mathscr{I}$. The user has tried $C \geq 0$ items before and we call them conditioning items as the platform can customize its recommendation based on the user's experiences with the C items.

We assume that a user receives a payoff of $v_1 > 0$ from a positive experience and $v_0 \leq 0$ from a negative experience with a target item. On the other hand, the platform gets $w_1 \in \mathbb{R}$ from the user's positive experience and $w_0 \in \mathbb{R}$ from a negative experience. If the user does not try the item, both the user and the platform receive a payoff of zero. The probability of either experience with a target item is initially unknown to the platform or to the target user. Thus, the platform wants to estimate the probability for each target item using the data it has accumulated before making a recommendation. We first define the recommender system and then describe the user behavior responding to a recommendation from the system.

Data

The data X is an $M \times N$ matrix which is a collection of ratings from the $N \geq 1$ users on M items. Specifically, the data contains previous users' ratings on the target items. For analytical convenience, we let the first C rows of X contain ratings of the conditioning items, whereas the remaining I rows represent ratings of the target items. Thus, the set of conditioning items is $\{1, \cdots, C\}$ and that of target items is $\mathscr{I} = \{C+1, \cdots, M\}$. Assume each item $i \in \{1, 2, \cdots, C, C+1, \cdots, M\}$ has a finite number of possible ratings. Without loss of generality, the set of possible such ratings of item i is denoted by $\{0, 1, \cdots, n_i\}$ for $n_i \in \mathbb{N}$. Since we are interested in the binary outcomes (positive or negative) from each of the target items, we assume $n_i = 1$ for $i \in \mathscr{I}$, where 0 represents a negative experience and 1 represents a positive experience with the item. Some users may not have ratings on some of the M items. If a user has a missing rating of item i, we denote this by $\emptyset$. Thus, in the data X, a typical element $x_{ij} \in \{\emptyset, 0, 1, \cdots, n_i\}$ records user j's rating of item i. Without loss of generality, we reserve the first column of X to denote the target user's ratings. Naturally, we have $x_{i1} = \emptyset$, $\forall i \in \mathscr{I}$ given the target user's ratings of target items must be predicted.

Data X is said to be *complete* if we have $x_{ij} \neq \emptyset$, $\forall i, \forall j$ s.t. $i \leq C$ or $j \neq 1$. In other words, the platform has complete data if the only missing ratings are the target user's rating of the target items. On the other hand, if the condition for complete data is not satisfied, the data is said to be *partial*. To reduce the notational burden, we focus here on the case in which the existing user data is complete, and we explain the extension to partial data in Online Appendix A. A complete set of data can always be obtained by using at least one tester who evaluates all the items on the platform before launching it to real users.[3]

Statistical model

For a given target user, the platform makes I estimations, one for each target item. To formally define the variable which the platform wants to estimate with regard to a target item i, we begin by defining the possible outcomes that a user can have with the C conditioning items and a target item. An outcome r is a vector of length $C+1$ that records the user's ratings of the C conditioning items and a target item. The k^{th} entry of r corresponds to a rating, which is not $\emptyset$, to the k^{th} item. Let R be the collection of all such outcomes r. For example, consider a platform that conditions on the reported rating of item 1 in making predictions about item 2 and item 3, i.e., $C = 1$ and $I = 2$.

[3]For example, Netflix hires editorial analysts whose main duty is to research, tag, annotate, rate, and analyze movies and TV shows.

The main example we use throughout the paper is represented in Table 1. The set of outcomes R in this case is given by $R = \{(1,1),(1,0),(0,1),(0,0)\}$.

	user 1	user 2	user 3	user 4	user 5
Rating on item 1	1	1	0	1	1
Rating on item 2	$\emptyset$	1	0	1	0
Rating on item 3	$\emptyset$	0	1	1	0

Table 1: Main example

The target user's ratings on the C conditioning items can be represented by r', which is a subvector of $r \in R$ whose length is C. We refer to r' as a history. The set of all such histories of length C is denoted by R'. When $C = 0$, we take $R' = \emptyset$. By slightly abusing notation, we denote the vector of dimension $C+1$ created by adding 1 to the end of r' by $(r',1)$. Similarly, $(r',0)$ denotes the vector created by appending 0 to the end of r'. Any $r \in R$ one-to-one corresponds to either $(r',1)$ or $(r',0)$ for some r'. For example, when $C = 2$, $r = (1,0,1)$ is equivalent to $(r',1)$ for $r' = (1,0)$.

For each target item i, the outcome is governed by an unobserved probability vector $p^i = (p^i_r)_{r\in R}$, where p^i is drawn according to a distribution function Π^i. We assume that all Π^i are mutually independent and the support of Π^i is the interior of an $|R|-1$ simplex. This realized probability p^i is referred to as a *correlation structure* as it reveals how the user experience with the target item i and conditioning items are correlated each other. Note p^i can take any value as long as it satisfies $\sum_{r\in R} p^i_r = 1$ and $p^i_r \geq 0$, $\forall r \in R$, i.e., it belongs to the standard $|R|-1$ simplex. The true probability of event r' is denoted by $p^i_{r'}$ and it is true that $p^i_{r'} = p^i_{(r',1)} + p^i_{(r',0)}$.

Recommender system

With these specifications, the platform wants to estimate the probability that the target user has a positive experience with target item i given that she has history r'. Using the correlation structure p^i, it can be represented as

$$z_i(r') = \frac{p^i_{(r',1)}}{p^i_{r'}}. \tag{1}$$

In our main example, the estimand associated with item 2 and item 3 are respectively

$$z_2(1) = \frac{p^2_{(1,1)}}{p^2_{(1,1)} + p^2_{(1,0)}} \text{ and } z_3(1) = \frac{p^3_{(1,1)}}{p^3_{(1,1)} + p^3_{(1,0)}}.$$

The fact that the same probability $z_i(r')$ applies to all the users with the same history r' does not necessarily mean that the users share identical preferences. This is because $z_i(r')$ only represents the average of the probabilities of liking item i of the group of users with history r'. As we will see shortly, different target users can have different preferences over item i in the model. Nonetheless, any two users with the same history are treated equally by the platform in terms of its estimations and subsequent recommendations. As a more detailed history becomes available to the platform (i.e., higher C), it is possible for the platform to make more precise estimations. One of the main contributions of the current work is to study how this improvement in the precision benefits (or harms) the users.

Let $\hat{z}_i^N(r')$ be a pointwise estimator of $z_i(r')$ in (1) when the target user's history is r' and there are $N-1$ previous users. The platform also decides a threshold $\tau(r') \in [0,1]$ that applies to history r' users. When $\hat{z}_i^N(r')$ is below this threshold, the item i is deemed not suitable for the history r' target user, and therefore, the system does not recommend this item to the user. We are now ready to define a recommender system.

Definition 1 *A recommender system is a collection of functions*

$$\{\{\hat{z}_i^N(r')\}_{i\in\mathscr{I}}, \tau(r')\}_{r'\in R'}$$

where $\hat{z}_i^N(r')$ maps the data X to a unit interval and $\tau(r') \in [0,1]$ for each i and r'. Given data X and the target user's history r', target item i is recommended to the target user if and only if $\hat{z}_i^N(r')(X) \geq \max_{j\in\mathscr{I}}\{\hat{z}_j^N(r')(X), \tau(r')\}$

Note that the value of data is governed by the particular recommender system that is in use. Depending on the estimator in a recommender system, different inferences can be drawn from the same data, which, in turn, can lead users to different decisions, and so as a result, different subsequent learnings. Furthermore, if we were to take the platform's exploration motive into account, a recommender system will tend to apply a relatively lower threshold level above which an item is recommended and increase the threshold as it accumulates more data until the threshold converges to a certain level.

Inevitably, the value from a specific recommender system cannot represent the universal value that data confers through any recommender system. Nevertheless, in the limit case as $N \to \infty$, we can evaluate the value that users can enjoy from a large class of recommender systems, which we call *consistent recommender systems*. A recommender system is said to be consistent if, for each r', the estimator $\hat{z}_i^N(r')$ is statistically consistent for $z_i(r')$. Formally, a consistent recommender system can be defined as follows.

Definition 2 *A recommender system is said to be consistent if $\hat{z}_i^N(r')$ converges in probability to $z_i(r') \in [0,1]$ for each $r' \in R'$.*

The existence of a consistent recommender system is guaranteed. We develop a Bayesian recommender system which satisfies the consistency requirement in Section 4. Although consistency is an asymptotic property, it is one of the most basic requirements for a recommender system to satisfy because the violation of consistency leads to a persistent error in predictions, which is clearly undesirable. Since consistency is a weak requirement, examples of consistent recommender systems are abundant. For example, the maximum likelihood estimator or Bayesian estimators with any convergent threshold are consistent recommender systems.

In what follows we focus on consistent recommender systems, and study the asymptotic value a recommender system creates and the value that additional customization in recommendations offers to users. Although the user value from any consistent recommender system converges to the asymptotic value as more data is available to the system, the value when only finite data is available depends on the specific recommender system in use.[4] To maintain the generality of our results even in finite data cases, we focus on a subclass of consistent recommender systems that satisfy a simple condition which we introduce in Section 3.5. Examples of such recommender system include the maximum likelihood estimator and a Bayesian recommender system which we introduce in Section 4.2. The latter is the system we use for our empirical analysis.

User behavior

For the target user, the true probability of a positive experience with a target item is denoted by q^i and it is drawn from a distribution function Q^i over the unit interval independently across the history r' users. A natural restriction we impose is the mean of Q^i is $p^i_{(r',1)}/p^i_{r'}$ as the ratio is the average probability of positive experiences with item i of the history r' user group. Thus, Q^i is dependent on the realized $p^i_{(r',1)}/p^i_{r'}$ and one can think of Q^i as an individual noise term that captures heterogeneity among r' history users. While both Π^i and Q^i are assumed to be common knowledge, the realizations p^i and q^i are not observable to users. As we focus on the asymptotic value, we consider the behavior of the target user who assumes the system has access to an asymptotically large database. Additionally, the user knows the recommended item has the highest average probability of positive experience for the history r' user group and it is above the threshold level $\tau(r')$.

Under this prior belief and knowledge, if an item is ever recommended to the target user, the user updates his belief using $\{\Pi^i\}_{i\in\mathscr{I}}$, $\{Q^i\}_{i\in\mathscr{I}}$ and $\tau(r')$. Precisely, for the recommended item i^*, the updated belief is a compound distribution of two distributions: (1) the distribution of the largest value above $\tau(r')$ among the realizations from $\{\Pi^i\}_{i\in\mathscr{I}}$ and (2) Q^{i^*}.[5] This compound

[4] The value that the Bayesian recommender system characterized in Section 4 offers to users when it learns from a finite dataset is given in Online Appendix B.

[5] For instance, if Π^i is a Dirichlet distribution with unit parameter for all i in $\mathscr{I}$, $p^i_{(r',1)}/p^i_{r'}$ follows $Beta(1,1)$. Accordingly, the distribution associated with the recommended item is the highest order statistic of I uniform distributions. At the individual level, let q^i be a draw from a Bernoulli trial: 1 with probability $p^i_{(r',1)}/p^i_{r'}$ and 0 otherwise. For simplicity, let $\tau(r') = 0$. In this case, the resulting $\bar{H}^{i^*}$ is a Bernoulli distribution with weights $I/(I+1)$ on 1 and $1/(I+1)$ on 0. Here, $I/(I+1)$ is the mean of the largest order statistic of I independent draws from the uniform distribution. Similarly, in case of a continuous distribution at the user level belief, we can find such a compound distribution.

distribution for the recommended item and its mean respectively is denoted by $\bar{H}^{i^*}$ and $\bar{\mu}^{i^*}$. On the other hand, for the non-recommended item $i \neq i^*$, the distribution is a compound distribution of (1) the distribution of the non-largest value among realizations from $\{\Pi^i\}_{i \in \mathscr{I}}$ when the largest value is above $\tau(r')$ and (2) Q^i. This distribution is denoted by $\bar{H}^i$ and its mean is denoted by $\bar{\mu}^i$.[6] Note that we have $\bar{\mu}^{i^*} \geq \bar{\mu}^i \ \forall i \in \mathscr{I}$ as each Q^i has mean $p^i_{(r',1)}/p^i_{r'}$. In contrast, if no item is recommended, the user updates the belief for i using a truncation of the original distribution. Because each $p^i_{(r',1)}/p^i_{r'}$ is below $\tau(r')$, the updated belief for i is a compound distribution of (1) the original distribution Π^i truncated at and above $\tau(r')$ and (2) Q^i. By $\underline{H}^i$ and $\underline{\mu}^i$, we respectively denote the distribution of i when no item is recommended. The inequality $\bar{\mu}^{i^*} > \max_{i \in \mathscr{I}}\{\underline{\mu}^i\}$ holds true as $\bar{\mu}^{i^*} \geq \tau(r')$ and $\underline{\mu}^i < \tau(r')$.

When the correlation structure is known to the platform, it is optimal for the user if the item with the highest probability is recommended *if and only if* the probability of liking the item is above $\frac{-v_0}{v_1 - v_0}$. We denote this user-optimal threshold level by τ^u. Similarly, provided the correlation structure is known to the platform, the maximum of the platform's surplus can be achieved by the threshold $\tau^p = \max\{\frac{-w_0}{w_1 - w_0}, 0\}$. By taking the maximum, we ensure τ^p stays positive. We rule out two trivial cases. First, if $\bar{\mu}^{i^*} < \tau^u$, the user never tries an item regardless of the recommendation. If this is true, the value a recommender system creates is zero. Second, if $\underline{\mu}^i \geq \tau^u$ for some i, the user tries item i even if no item is recommended to try. However, in most real-world applications, platforms are biased towards usage and exploration, and therefore, the threshold adopted by the platform is lower than the user-optimal threshold level. Thus, we focus on the case $\tau(r') \leq \tau^u$ which rules out the second case. In the remaining case, which is the focus of this study, we have (1) $\bar{\mu}^{i^*} \geq \tau^u$ and (2) $\underline{\mu}^i < \tau^u$. It is easy to see that in this case the target user always follows the recommendation from the platform.[7]

3.2 Value of data and its decomposition

Focusing on a consistent recommender system, we start by offering a characterization and decomposition of the value offered by data. Specifically, value is created for users in two steps—the system learns the correlation structure p from ratings left by previous users and it then personalizes predictions based on the target user's history r'. We can compare this full learning, which combines both learning across different users and within the same user, to the case without any customization, so that the system has to recommend items to the target user without any information on the target user. We call this latter case, a *generic recommendation* since it involves no customization to the user. Formally, using our terminology, it can be defined as follows. Let z^G_i be the true probability that a user has a positive experience with target item i, i.e., $z^G_i = \sum_{r' \in R'} p^i_{(r',1)}$.

[6]Continuing on the example of the previous footnote, the average probability density function of $p^i_{(r',1)}/p^i_{r'}$ when i is not recommended is $\frac{I}{I-1}(1 - z^{n-1})$. Using this, we can derive that the resulting $\bar{H}^i$ is a Bernoulli distribution with weights $\frac{I}{I-1}\left(\frac{1}{2} - \frac{1}{I+1}\right)$ on 1.

[7]The example with the utility specification $(v_1, v_0) = (1, -1)$ belongs to this case.

Definition 3 *A generic recommender system is a recommender system* $\{\{\hat{z}_i^G\}_{i\in\mathscr{I}}, \tau^G\}$ *such that* $\hat{z}_i^G$ *is a consistent estimator for* z_i^G.

Both the estimator and the threshold of a generic recommender system do not depend on a particular history of the target user. Generic recommendations are closely related to the average rating mechanism commonly found in review sites such as those used on Amazon's marketplace, Apple's Appstore, eBay, Yelp, and so on, and widely studied in the literature (see Dellarocas et al. (2006) for a survey on such rating systems). In this simple alternative to a recommender system, users' ratings are averaged, and users can use this average score to determine whether to "use" the item. The direct relationship between a generic recommender system and the average rating mechanism is studied in Section 4.2.

We decompose the learning from a recommender system into its generic and customized components. Note that we focus on the ex-ante value of data. The actual value created by data is affected by the realization of data and a target user's history. For example, because more users tried item 1, the data in the example of Table 1 is generally more useful to the users who already had a positive rating of item 1 than those users who had a negative rating of item 1. Thus, in evaluating the value of data, we first find the ex-post value in terms of a particular realization of data and a target user's history, and then we quantify the value from the ex-ante perspective evaluated with respect to the true probability measure p. That is, to find the value of data, we take a weighted average over the values created from all possible realizations of data, with the weights being the true underlying probability measures.

Lemma 1 characterizes the steady-state utility that the target user and the platform can expect from a consistent recommender system and a generic recommender system when there is only one item available to the user.

Lemma 1 **(**$I = 1$ case**)** *Suppose there is only one item available to the target user, i.e.,* $i = C + 1$.

1. *For any consistent mechanism, the expected utility to the target user and to the platform respectively converges to*

$$\sum_{r'\in R'} \mathbf{1}\{p^i_{(r',1)} \geq \tau(r')p^i_{r'}\}(v_1 p^i_{(r',1)} + v_0 p^i_{(r',0)}) \text{ and}$$
$$\sum_{r'\in R'} \mathbf{1}\{p^i_{(r',1)} \geq \tau(r')p^i_{r'}\}(w_1 p^i_{(r',1)} + w_0 p^i_{(r',0)}).$$

2. *The expected utility to the target user and to the platform from a generic recommendation*

respectively converges to

$$\mathbf{1}\{\sum_{r'\in R'} p^i_{(r',1)} \geq \tau^G\} \sum_{r'\in R'} (v_1 p^i_{(r',1)} + v_0 p^i_{(r',0)}) \text{ and}$$
$$\mathbf{1}\{\sum_{r'\in R'} p^i_{(r',1)} \geq \tau^G\} \sum_{r'\in R'} (w_1 p^i_{(r',1)} + w_0 p^i_{(r',0)}).$$

The derivation of the above expressions proceeds as follows. When the system completes its learning about p, the target item is recommended to a user with history r' if and only if $\frac{p^i_{(r',1)}}{p^i_{r'}} \geq \tau(r')$. Once the item is recommended and tried by the user, the user receives $v_1 \frac{p^i_{(r',1)}}{p^i_{r'}} + v_0 \frac{p^i_{(r',0)}}{p^i_{r'}}$. The ex-ante utility is derived taking into account that the target user has a history r' with probability $p^i_{r'}$. The same logic can be applied in deriving the other utility characterizations. The key difference in the expected utilities between the two systems directly comes from the difference in their estimators. A consistent mechanism for history r' target user seeks to estimate the conditional probability that the user likes the item, i.e., $\frac{p^i_{(r',1)}}{p^i_{r'}}$. This conditional probability averages out to the total probability $\sum_{r'\in R'} \frac{p^i_{(r',1)}}{p^i_{r'}} \cdot p^i_{r'}$, which is the estimand of a generic recommender system that does not offer customization.

In Proposition 1 we evaluate the expected utility of a target user when the system adopts the threshold level that is optimal for users and decompose the value into two parts: the value from learning the correlation structure and the value from customization.

Proposition 1 *1. For the user-optimal threshold, in any consistent recommender system, as $N \to \infty$, the expected utility to a user converges to*

$$\sum_{r'\in R'} \max_{i\in\mathscr{I}} \{v_1 p^i_{(r',1)} + v_0 p^i_{(r',0)}, 0\}.$$

2. For the user-optimal threshold, the expected utility created from a generic recommendation in the limit as $N \to \infty$ is

$$\max_{i\in\mathscr{I}} \{\sum_{r'\in R'} (v_1 p^i_{(r',1)} + v_0 p^i_{(r',0)}), 0\}.$$

Customization can be regarded as a finer segmentation of users according to observable features. To see why a finer segmentation necessarily leads to a higher user welfare under the user-optimal threshold level, consider a history r' group of users whose average probability of liking a (unique) target item is given by z. Under the user-optimal threshold, the average utility of the history r' users is $\max\{v_1 z + v_0(1-z), 0\}$. If a further segmentation of r' users into subgroups r'_a and r'_b is available to the system, the system makes two separate recommendations according to the two new average probabilities, z_a and z_b, of the two subgroups. Let p_a portion of r' users is assigned to r'_a subgroup. The expected utility of users under segmentation is then $p_a \max\{v_1 z_a + v_0(1 - z_a), 0\} + (1-p_a)\max\{v_1 z_b + v_0(1-z_b), 0\}$, where $z = p_a z_a + (1-p_a) z_b$. The customization reveals the subgroup that is on-average better off trying (or not trying) the item, which in turn increases the overall expected utility. In Figure 1, the expected utility without customization is zero while with customization it is $(1-p_a)(v_1 z_b + v_0(1-z_b))$.

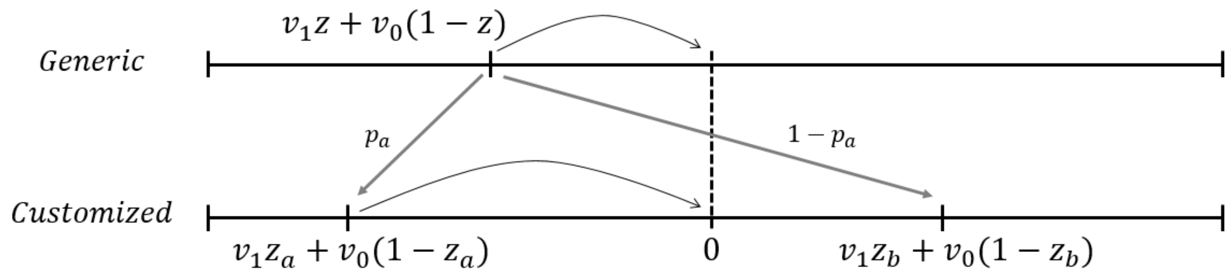

Figure 1: Customization under the user-optimal threshold

The difference between the two expressions in the proposition is due to the different degree of customization each recommender system offers. A consistent recommender system makes history-dependent recommendations by selecting the best item for each group of history r' users. Instead, a generic recommender system only selects the item on average best for every user. Therefore, the value from generic recommendations in Proposition 1 can be attributed to pure across-user learning. On the other hand, the difference between the value from consistent recommendations and the value from generic recommendations can be attributed to the customization benefits that the consistent recommendation provides. Note that, under the user-optimal threshold, the expected utility of the target user created from a consistent recommender system is always positive, and it is at least as great as the expected utility generated from the generic recommendation, which is also always positive. That is, both the pure across-user learning and the customization component always add value to users when the recommender system is user-centric.

Note that, although the two systems differ in their degrees of customizations, they share the same value creation process: a recommender system selects the most suitable item and at the same time, screens items not suitable for the user. Obviously, the more items in the target item pool, the more value is created from both of the systems. On the other hand, the benefit from screening is maximized when the threshold is properly chosen for the user.

Example 1: The main example

To illustrate and better understand how a recommender system adds value through learning and customization, consider our main example in Table 1 and assume $(v_1, v_0) = (1, -1)$. In the example, the platform makes a recommendation between item 2 and item 3, or none based on the target user's rating on item 1. According to Proposition 1, when the threshold level is τ^u, the expected utility of a user under a consistent mechanism in the limit as $N \to \infty$ is given by

$$\underbrace{\max\{p^2_{(1,1)} - p^2_{(1,0)},\ p^3_{(1,1)} - p^3_{(1,0)},\ 0\}}_{Recommendation\ for\ history\ 1\ user} + \underbrace{\max\{p^2_{(0,1)} - p^2_{(0,0)},\ p^3_{(0,1)} - p^3_{(0,0)},\ 0\}}_{Recommendation\ for\ history\ 0\ user}. \tag{2}$$

Depending on the user's history, different items can be recommended. For example, if $p^2_{(1,1)} - p^2_{(1,0)} > \max\{p^3_{(1,1)} - p^3_{(1,0)}, 0\}$ and $0 > \max\{p^2_{(0,1)} - p^2_{(0,0)}, p^3_{(0,1)} - p^3_{(0,0)}\}$, item 2 is recommended to users who had a positive experience with item 1 but no item is recommended to the user if she had a negative experience with item 1. On the other hand, under a generic recommender system, which only makes use of across-user learning, users receive the same recommendation regardless of their history. The value from a generic recommendation with the same threshold being applied is given by

$$\underbrace{\max\{p^2_{(1,1)} + p^2_{(0,1)} - p^2_{(1,0)} - p^2_{(0,0)},\ p^3_{(1,1)} + p^3_{(0,1)} - p^3_{(1,0)} - p^3_{(0,0)}, 0\}}_{No\ customization\ in\ recommendations}. \tag{3}$$

It is clear that the value from customization is always weakly positive. On the other hand, it is strictly positive *if and only if* users who had different experiences with item 1 receive different recommendations.

Example 2: Degree of correlation and value from customization

The correlation between a target item and conditioning items on which the system customizes its recommendation plays an important role in determining the value from customization. To illustrate, consider the case in which there are only one target item and one conditioning item: $I = 1$ and $C = 1$. A consistent recommender system conditions its recommendation of the target item (item 2) on the user's experience with item 1. To see how correlation is related to the value a recommender system creates, we impose a structural assumption on p^2. Consider a specific correlation structure in which the correlation is captured by a parameter γ: $p^2_{(1,1)} = \gamma p^2_{(1,0)}$ and $p^2_{(0,0)} = \gamma p^2_{(0,1)}$. So, if a user has a positive experience with item 1, she is γ times more likely to have a positive experience with item 2 than a negative experience. Similarly, if a user has a negative experience with item 1, she is γ times more likely to have a negative experience with item 2 than a positive experience. Denote the probability of a positive experience with item 2 by t, i.e., $t = p^2_{(1,1)} + p^2_{(0,1)}$. Under this specific structure, the value of item 2 in terms of t under our recommender system with its customized recommendations is

$$\max\{t - \frac{1}{\gamma+1}, 0\} + \max\{t - \frac{\gamma}{\gamma+1}, 0\}.$$

The value from the generic recommendation is $\max\{t - (1-t), 0\}$, while the value without any recommendation is $t-(1-t)$ if the user tries the item. Figure 2 plots the value of item 2 in terms of its success probability t when $\gamma = 2$ and when $\gamma = 10$. As we can see, the value from customization increases as we increase the correlation parameter γ.

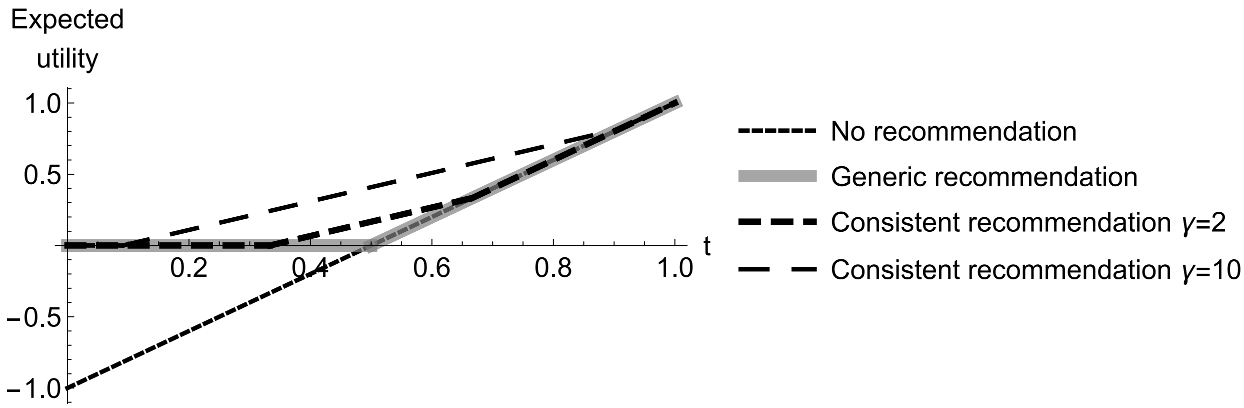

Figure 2: Value of customization

Before turning to study the marginal benefit from customization and an extra data point[8], some remarks are worth making with respect to Proposition 1. First, the utility representation of the proposition can be extended to accommodate the situations in which the user has multiple demands and the platform recommends multiple items in response to the demands. Suppose that the user is willing to try at most D items and the user utility is defined to be the sum of the utilities from each item the user tries. By $v_1 p^{(k)}_{(r',1)} + v_0 p^{(k)}_{(r',0)}$, we denote the k^{th} largest[9] value among $\{v_1 p^{i}_{(r',1)} + v_0 p^{i}_{(r',0)}\}_{i \in \mathscr{I}}$ for each history $r' \in R'$. When the platform recommends at most D items whose predicted probability of liking an item is above the user-optimal threshold level, the user utility from the recommender systems can be represented as follows[10]:

Corollary 1 (Corollary to Proposition 1) *Suppose the recommender system makes at most $D \leq I$ recommendations to a user who is willing to try the items.*

1. *For the user-optimal threshold, in any consistent recommender system, as $N \to \infty$, the expected utility converges to*

$$\sum_{r' \in R'} \sum_{k \in \{1, \cdots, D\}} \max\{v_1 p^{(k)}_{(r',1)} + v_0 p^{(k)}_{(r',0)}, 0\}.$$

2. *For the user-optimal threshold, the expected utility created from a generic recommendation in the limit as $N \to \infty$ is*

$$\sum_{k \in \{1, \cdots, D\}} \max\Big\{ \sum_{r' \in R'} (v_1 p^{(k)}_{(r',1)} + v_0 p^{(k)}_{(r',0)}), 0\Big\}.$$

[8]When data is complete, each set of ratings from a previous user corresponds to an outcome. We call this set of ratings over C conditioning items and I target items a data point.

[9]In case of a tie, the lower index i in $\mathscr{I}$ takes the lower value for k.

[10]A minor and straightforward modification to the assumption on the user behavior is required to ensure the user tries the recommended items.

Second, it is worth mentioning that our ex-ante approach makes it possible to find the achievable asymptotic lower bound of prediction error of a consistent recommender system in a simple closed-form. It can be shown that this lower bound of the error is weakly decreasing in the degree of customization. Specifically, suppose that a consistent recommender system $\{\{\hat{z}_i^N(r')\}_{i\in\mathscr{I}}, \tau(r')\}_{r'\in R'}$ is in use. Let ϵ_i^m be the number of wrong predictions, both the false positives (negative experience from a recommended item) and false negatives (unrealized positive experience when no item is recommended), out of the total of m predictions made about the target items. The prediction error is defined to be $\epsilon_i^m(X)/m$. In the limit, using the same logic we used in the derivation of Lemma 1, the prediction error of a consistent recommender system with the threshold level $\{\tau(r')\}_{r'\in R'}$ when there is only one target item ($i = C+1$) converges to

$$\sum_{r'\in R'} p^i_{r'}\left(\frac{p^i_{(r',0)}}{p^i_{r'}}\mathbf{1}\left\{\frac{p^i_{(r',1)}}{p^i_{r'}} \geq \tau(r')\right\} + \frac{p^i_{(r',1)}}{p^i_{r'}}\mathbf{1}\left\{\frac{p^i_{(r',1)}}{p^i_{r'}} < \tau(r')\right\}\right). \tag{4}$$

That is, history r' occurs with probability $p^i_{r'}$, a false positive recommendation is made with probability $\frac{p^i_{(r',0)}}{p^i_{r'}}$ once the item is recommended to history r' target user, and a false negative event happens with probability $\frac{p^i_{(r',0)}}{p^i_{r'}}$ when the item is not recommended to the user. The lower bound of this asymptotic error can be achieved when the threshold is user-optimal, and it is characterized as follows for the general $I \geq 1$ cases.

Corollary 2 (Corollary to Proposition 1) *For $C \geq 0$ and $I \geq 1$,*

1. *The asymptotic lower bound of prediction error in a consistent recommender system is*

$$L_C = \sum_{r'\in R'} \min_{i\in\mathscr{I}} \min\{p^i_{(r',1)}, p^i_{(r',0)}\}. \tag{5}$$

2. *L_C weakly decreases in C.*

From the second point of the corollary, it is clear that the lower bound of a generic recommender system ($C = 0$) is always weakly higher than that of a consistent recommender system. Additionally, as can be directly seen from (5), the lower bound is also weakly decreasing as more items are added to the target item pool, i.e., L_C decreases in I.

Lastly, with a slight modification, Proposition 1 can be applied to evaluate the platform surplus under the platform-optimal threshold level, or the social surplus when the socially-optimal threshold level is adopted by the system. For the platform-optimal threshold, the limit of the platform's expected surplus from a consistent recommender system converges to $\sum_{r'\in R'} \max_{i\in\mathscr{I}}\{w_1 p^i_{(r',1)} + w_0 p^i_{(r',0)}, 0\}$ whereas that from a generic recommender system converges to $\max_{i\in\mathscr{I}}\big\{\sum_{r'\in R'}(w_1 p^i_{(r',1)} + w_0 p^i_{(r',0)}), 0\big\}$. Moreover, if we define the social surplus to be the sum of the user surplus and the platform surplus, the limit social surplus from the two recommender systems are respectively

$\sum_{r' \in R'} \max_{i \in \mathscr{I}} \{(v_1 + w_1) p^i_{(r',1)} + (v_0 + w_0) p^i_{(r',0)}, 0\}$ whereas that from a generic recommender system converges to $\max_{i \in \mathscr{I}} \left\{ \sum_{r' \in R'} ((v_1 + w_1) p^i_{(r',1)} + (v_0 + w_0) p^i_{(r',0)}), 0 \right\}$ when the socially optimal threshold level is being applied, which is $\frac{-v_0 - w_0}{v_1 - v_0 + w_1 - w_0}$.

3.3 Harmful customization

While Proposition 1 shows customization always (weakly) benefits users under the user-optimal threshold, it can also be shown that such a user-optimal threshold is the only threshold level for which customization benefits the user regardless of the correlation structure and user history.

Proposition 2 *For any $\tau(r') \neq \tau^u$ and for any $I \geq 1$, there exists a collection of I correlation structures $q = \{q^i\}_{i \in \mathscr{I}}$ such that the history r' target user is strictly worse-off from customization when the true correlation structures are q.*

The same argument holds for the platform value or the social value. The customization in predictions is advantageous to the target user, the platform or the society only when the threshold level is properly chosen. Put differently, when $\tau^u \neq \tau^p$, either the user or the platform will be strictly worse off under some correlation structures when customization is taken place. Thus, customization has the scope to hurt users and total welfare if the platforms' interests cause its threshold level to diverge from the user-centric or welfare-centric thresholds. We revisit this issue in Section 3.4.

The value representation in Proposition 1 can be generalized to accommodate an arbitrary threshold level, τ^p, adopted by the platform. Under τ^p, the item with the highest estimated probability of positive experience is recommended only if it is above τ^p. The resulting expected utility to the target user can be represented as follows.

Corollary 3 (Corollary to Proposition 2) *Suppose the recommender system applies a threshold level τ^p. In any consistent recommender system with τ^p, as $N \to \infty$, the expected utility generated by the target item converges to*

$$\sum_{r' \in R'} \mathbf{1}\big\{ \max_{i \in \mathscr{I}} v_1 p^i_{(r',1)} + v_0 p^i_{(r',0)} \geq p_{r'}(v_1 \tau^p + v_0(1 - \tau^p))\big\} \big(\max_{i \in \mathscr{I}} v_1 p^i_{(r',1)} + v_0 p^i_{(r',0)} \big).$$

3.4 Marginal value of customization

As highlighted in the previous section, recommender systems learn about target items not only from the ratings left on the target items by other users, but also from the ratings left from the target user on other items so as to better customize the recommendation. Hence, the quality of a recommendation and the resulting user surplus depend also on the degree of customization the platform provides. In this section, we study how customization affects user value through the recommender system. Specifically, we are interested in the effect of changes in the degree of customization to the value generated by each target item under a recommender system. To do so, we isolate the target item by considering the case $I = 1$, and investigate the sources behind the value creation process of a recommender system.

Suppose that a consistent recommender system takes one more conditioning item into account in making a recommendation on the target item. We denote the new item and the target item respectively by item $C+1$ and item $C+2$. For simplicity, we assume that item $C+1$ can have binary ratings. In this case, since the recommender system can condition its predictions on one more item, the predictive accuracy of the target item improves. However, the improvement in accuracy does not necessarily lead to higher user welfare because of the conflict in interests between the user and the platform which is captured by the threshold level. Here, we study this marginal benefit or harm of customization. Since there is only one target item, we save notation by using p for p^{C+2}, the correlation structure associated with the $C+1$ conditioning items and the target item.

Definition 4 *For $r' \in R'$ over items $\{1, \cdots, C\}$, item $C+1$ and item $C+2$ are positively correlated conditional on r' if $v_1 p_{(r',1,1)} + v_0 p_{(r',1,0)} \geq 0$ and $v_1 p_{(r',0,1)} + v_0 p_{(r',0,0)} \leq 0$, with at least one inequality holding strictly. If both inequalities are strict, we say that they are strictly positively correlated conditional on r'.*

When the target item and item $C+1$ are positively correlated conditional on r', users whose history is r' over the other C conditioning items are more likely to have net positive utility from the target item if they liked item $C+1$. On the other hand, if they did not like the item, it is more likely that they have net negative utility from the target item. When $(v_1, v_0) = (1, -1)$, the condition is satisfied if we have $p_{(r',1,1)} \geq p_{(r',1,0)}$ and $p_{(r',0,1)} \leq p_{(r',0,0)}$ with at least one inequality holding strictly. In a parallel way we can define a (strict) negative correlation between item $C+1$ and the target item. If item $C+1$ and the target item are positively or negatively correlated conditional on $r' \in R'$, they are said to be correlated conditional on r'. For strict inequalities, they are said to be strictly correlated conditional on r'.

We focus on a universal threshold level $\tau = \tau(r'),\ \forall r'$. For each r', there are two trivial cases in which an extra degree in customization yields only zero marginal customization effect: when the threshold level is too high or too low as stated below.

$$\tau \leq \min\left\{\frac{p_{(r',1,1)}}{p_{(r',1)}}, \frac{p_{(r',0,1)}}{p_{(r',0)}}\right\} \text{ or } \tau \geq \max\left\{\frac{p_{(r',1,1)}}{p_{(r',1)}}, \frac{p_{(r',0,1)}}{p_{(r',0)}}\right\}. \tag{6}$$

If the former is the case, the item is recommended to the user regardless of the history and the introduction of the extra customization. This follows from the following equality.

$$\frac{p_{(r',1,1)} + p_{(r',0,1)}}{p_{r'}} = \frac{p_{(r',1)}}{p_{r'}} \frac{p_{(r',1,1)}}{p_{(r',1)}} + \frac{p_{(r',0)}}{p_{r'}} \frac{p_{(r',0,1)}}{p_{(r',0)}}.$$

Similarly, in the latter case, the item is not recommended regardless of the introduction of the extra degree in customization. As a result, the extra customization cannot affect user welfare if (6) holds. In what follows, we focus on the remaining case, i.e.,

$$\min\left\{\frac{p_{(r',1,1)}}{p_{(r',1)}}, \frac{p_{(r',0,1)}}{p_{(r',0)}}\right\} < \tau < \max\left\{\frac{p_{(r',1,1)}}{p_{(r',1)}}, \frac{p_{(r',0,1)}}{p_{(r',0)}}\right\}.$$

In the next proposition, we find that a strict correlation between the newly added item and the target item is a necessary and sufficient condition under which an extra degree in customization is (on average) beneficial to users regardless of the misalignment of interests between the platform and users.

Proposition 3 *The marginal customization strictly benefits r' user for any threshold level if and only if item $C+1$ and the target item (item $C+2$) are strictly correlated conditional on r'.*

When the newly added item $(C+1)$ is strictly correlated to the target item for r' users, the system can use the item to segment r' users into subgroups according to their experiences with the item. By the definition of correlation, one subgroup should have on-average a positive experience with the target item while the other group does not. Thus, screening the latter subgroup is a welfare-improving segmentation for all threshold levels.

Next, we consider the expected utility of all users. From Proposition 3, it is evident that the expected utility is strictly positive if the newly added item and the target item are correlated conditional on all histories. However, if the newly added item is not correlated with the target item for some users with particular histories, the expected utility which is represented in Corollary 3 can take a negative value. The following corollary of Proposition 3 finds that this cannot happen under the user-optimal threshold level. The correlation of the two items for some history of users is a necessary and sufficient condition for the marginal benefit to be strictly positive.

Corollary 4 *Under τ^u, the benefit from an extra degree in customization is*

1. *always (weakly) positive and given by*

$$\sum_{r'\in R'}\big[\max\{v_1 p_{(r',1,1)} + v_0 p_{(r',1,0)}, 0\} + \max\{v_1 p_{(r',0,1)} + v_0 p_{(r',0,0)}, 0\}$$
$$- \max\{v_1 p_{(r',1,1)} + v_0 p_{(r',1,0)} + v_1 p_{(r',0,1)} + v_0 p_{(r',0,0)}, 0\}\big]$$

2. *strictly positive if and only if $\exists r' \in R'$ s.t. item $C+1$ and the target item are correlated conditional on r'.*

One may wonder how the marginal value of having one more item to customize on is affected as the total number of items used for prediction increases. This can go in either direction. The marginal value depends largely on the underlying correlation structure. To see this, let $(v_1, v_0) =$

$(1,-1)$ and consider a platform that has three items under the user-optimal threshold level. The target item is fixed at item 3 and we compare user surplus under two correlation structures. In both cases, we will find the expected user utility using Proposition 1 when (1) the platform does not customize the recommendation, (2) it conditions its prediction on only one other item, and (3) it conditions its prediction on both items. The two correlation structures are given as follows:

$$q^1 = \left(\frac{1}{6}, \frac{1}{12}, \frac{1}{12}, \frac{1}{6}, \frac{1}{12}, \frac{1}{6}, \frac{1}{6}, \frac{1}{12}\right),\ q^2 = \left(\frac{1}{6}, \frac{1}{12}, \frac{1}{6}, \frac{1}{12}, \frac{1}{12}, \frac{1}{6}, \frac{1}{12}, \frac{1}{6}\right).$$

When the platform initially does not customize its prediction, the expected utility from item 3 in both cases is zero as

$$q^1_{(1,1,1)} + q^1_{(1,0,1)} + q^1_{(0,1,1)} + q^1_{(0,0,1)} = q^2_{(1,1,0)} + q^2_{(1,0,0)} + q^2_{(0,1,0)} + q^2_{(0,0,0)} = \frac{1}{2}$$

On the other hand, if the platform conditions on both items in making a prediction about item 3, the expected utility to a user is $\frac{1}{6}$ in both cases. Now, when the platform conditions only on item 1, the user surplus under q^1 is

$$\max\{q^1_{(1,1,1)} + q^1_{(1,0,1)} - q^1_{(1,1,0)} - q^1_{(1,0,0)}, 0\} + \max\{q^1_{(0,1,1)} + q^1_{(0,0,1)} - q^1_{(0,1,0)} - q^1_{(0,0,0)}, 0\}$$

which is zero. However, under q^2, the resulting user surplus is $\frac{1}{6}$. So, under q^1, the marginal increment of expected utility in terms of the degree of customization is increasing while it is decreasing under q^2. The same result holds when we condition on item 2 only.

Suppose now that the system uses the platform-optimal threshold level $\tau = \tau^p = \frac{-w_0}{w_1 - w_0}$. We assume that $w_1 > 0$, $w_0 < 0$ and $\tau^p \neq \tau^u$. As is pointed out in Proposition 2 customization can harm user welfare under some realizations of the correlation structure, p. We identify two situations under which an extra degree in customization harms history r' user's surplus[11].

Proposition 4 *Let τ^p be the platform-optimal threshold. A history r' target user is strictly worse off from an additional degree in customization if and only if one of the following is true*

1. $\frac{p_{(r',1,1)} + p_{(r',0,1)}}{p_{(r',1,0)} + p_{(r',0,0)}} < \frac{-w_0}{w_1} \leq \frac{p_{(r',e,1)}}{p_{(r',e,0)}} < \frac{-v_0}{v_1}$, $e = 0$ *or* 1
2. $\frac{p_{(r',1,1)} + p_{(r',0,1)}}{p_{(r',1,0)} + p_{(r',0,0)}} \geq \frac{-w_0}{w_1} > \frac{p_{(r',e,1)}}{p_{(r',e,0)}} > \frac{-v_0}{v_1}$, $e = 0$ *or* 1.

When an extra degree of customization is introduced to the system, it may generate a false positive or a false negative for the user if the threshold is not optimally chosen from the user's perspective. A false positive is a situation under which the target item is recommended to a user under the platform-optimal threshold but which would not have been recommended to the user under the user-optimal threshold. On the other hand, a false negative refers to the situation under

[11] To provide a more general implication, in this proposition, we discard the assumption $\tau^p < \tau^u$ we imposed before.

which a user is not recommended the item under the platform-optimal threshold even though the item would have been recommended under the user-optimal threshold. The first bullet point in the proposition is related to the false positive situation, whereas the second is related to the false negative situation. To see this, let $r' = \emptyset$ and $e = 1$ for simplicity and suppose that the first bullet point is satisfied. Firstly, $\frac{p_{(1,1)}+p_{(0,1)}}{p_{(1,0)}+p_{(0,0)}} < \frac{-w_0}{w_1}$ implies that the target item (item 2) will not be recommended to the user if the recommendation is not customized to the user's history. The overall probability of a positive experience with item 2 is lower than the system's threshold. However, when the recommendation is customized, item 2 will be recommended to the user who had a positive experience with item 1. This is captured by the second inequality $\frac{-w_0}{w_1} \leq \frac{p_{(1,1)}}{p_{(1,0)}}$. However, this recommendation is harmful to the user as the probability of a positive experience to the user is still lower than the threshold that is optimal for the user, i.e., $\frac{-v_0}{v_1} > \frac{p_{(1,1)}}{p_{(1,0)}}$. The second bullet point can also be interpreted using similar reasoning.

3.5 Value of an additional user

Thus far we have focused on the asymptotic value that a recommender system creates. In this section, we study the value from finite data to study how a marginal data point (i.e. on another user) adds value to other users. Contrary to the asymptotic value, the value from finite data varies depending on the estimator that a recommender system adopts as different consistent estimators can induce different predictions from finite data. Thus, we impose a minimal structure that a consistent estimator should satisfy, and study the value from marginal data and how it behaves as the system accumulates more data as it adds users.

Bajari et al. (2019) and Acemoglu et al. (2021) show theoretically that prediction error decreases with a diminishing rate with respect to the size of the data used in prediction. In our benchmark setting we are able to find a similar result but for user surplus, which we find is monotone, with its increment diminishing in the size of accumulated data. We will briefly discuss how the learning curve can be convex or S-shaped in Section 6. These findings, on whether data is sub-additive or super-additive, have implications for different market outcomes as studied in Gu et al. (2021), Hagiu and Wright (2021) and Ichihashi (2020b).

To keep the analysis as simple as possible, we consider the case with only two items when $(v_1, v_0) = (1, -1)$, and then look at the choice problem of the target user when the platform has learned the correlation structure from the previous $N-1$ users and the target user's experience with the first item. Under this simplification, the system elicits the target user's preference from her ratings of the first item ($i = 1$) and makes a prediction about her experience with the target item ($i = 2$).

Let the data X be summarized by $y = (y_{(1,1)}, y_{(1,0)}, y_{(0,1)}, y_{(1,1)})$. Suppose y_r records the occurrences of outcome $r \in R = \{(1,1), (1,0), (0,1), (0,0)\}$ that appear in X. Suppose the target user has a positive experience with item 1 and there are N_1 previous users who had the same experience as the user with item 1. The same analysis applies in the case in which the user has a negative experience with item 1. The true but unknown probability of a positive experience with item 2 is $\frac{p_{(1,1)}}{p_{(1,1)}+p_{(1,0)}}$, which we denote by s. Under this setup, the target user's history can be denoted by $1 \in R' = \{1, 0\}$. We focus on the recommender systems that satisfy the following condition.

Definition 5 *A recommender system is said to be unbiased if the target item is recommended to history r' target user if and only if $y_{(r',1)} \geq y_{(r',0)}$.*

In any unbiased recommender system, the target item is recommended to try if and only if the target user is more likely to have a positive experience with the target item in the sense that the previous users who share the same history with the target item have reported more positive ratings than negative ratings.

For a given s, according to this recommendation rule, the item is recommended to the user with probability

$$P[y_{(1,1)} \geq y_{(1,0)}] = \sum_{k=\lceil (y_{(1,1)}+y_{(1,0)})/2 \rceil}^{y_{(1,1)}+y_{(1,0)}} \binom{y_{(1,1)}+y_{(1,0)}}{k} s^k (1-s)^{y_{(1,1)}+y_{(1,0)}-k}, \tag{7}$$

where $\lceil x \rceil$ denotes the ceiling function, i.e., the least integer greater than or equal to x. To avoid a mathematical complexity that is involved with binomial probability and the ceiling function, assume that $y_{(1,1)} + y_{(1,0)}$ is only an odd number by considering the case of $y_{(1,1)} + y_{(1,0)} = 2m - 1$, $m \in \mathbb{N}$. The same results apply to the even-number cases. Let $v(m|s)$ denote the expected value of the target user in terms of the number of previous users $2m - 1$ and the true (and unknown) conditional probability of a positive experience s. Formally, we have

$$v(m|s) = sP[y_{(1,1)} \geq y_{(1,0)}] - (1-s)P[y_{(1,1)} \geq y_{(1,0)}].$$

Note that $v(\cdot|\cdot)$ can take a negative value. For example, for any $s \in (0, \frac{1}{2})$, it is possible that the occurrences of $y_{(1,1)}$ exceed that of $y_{(1,0)}$. Therefore, the possibility of a wrong recommendation always exists even though the platform is committed to giving the recommendation that it expects to be best for the target user. Although the addition of a data point from an additional user always leads to a more informative data structure, it can be shown that the incremental value diminishes. The following results characterize the nature of data under the recommender system.

Proposition 5 *Let $v(m|s)$ denote the value to the target user and $\Delta v(m|s)$ be the marginal value that the target user contributes to the next user.*

1. *The value is positive and increases in the number of previous users. i.e., $v(m|s)$ increases in m, $\forall s$.*

2. *The marginal value diminishes. i.e., $\Delta v(m|s)$ decreases in m, $\forall s$.*

Note that the increment is strictly positive for any $s \neq \frac{1}{2}$. As a result, when a user is uncertain about s, she expects a strictly higher payoff from a platform using a recommender system with a larger dataset unless she holds a degenerate prior belief (i.e. $s = \frac{1}{2}$). An immediate corollary follows from this fact.

Corollary 5 *Suppose the target user believes s is distributed according to a CDF S, which does not have a mass of one at $s = \frac{1}{2}$. Then $E_s[v(m+1|s)] > E_s[v(m|s)]$, $\forall m$.*

Corollary 5 says that, except for a degenerate case, the recommender system generates a positive data network effect in which users' expected utility increases as more other users are added.

The reason behind the submodularity of data has a clear connection to how a Bayesian posterior is formulated. A Bayesian posterior is a compromise between a prior and data. As we increase the data size, the relative contribution of each data point to the posterior becomes smaller. Thus, any particular data point that has the same history as the target user's becomes less influential in forming the expected utility of the target user.

4 Bayesian model for a recommender system

In order to quantify the results implied by our theory with data, in this section we construct a Bayesian model of a recommender system that learns the correlation structure and makes customized predictions based on the target user's history. It is a consistent recommender system and satisfies the recommendation rule that we imposed in Section 3.5 provided the prior distribution is properly chosen. We use this Bayesian model in Section 5 for our counterfactual experiments. As before, we focus on *complete* data, although we do not require that the target user has given ratings for all the C items. The extension of our Bayesian recommender system to handle *incomplete* data is explained in Online Appendix A.

4.1 Bayesian model of a recommender system

Let $C+1$ be the target item. The platform's objective is to learn p^{C+1}, the correlation structure, so as to make predictions about a target user's preference for items that the user has not yet tried. Again, to save notation, we write p for p^{C+1}. We use Bayesian parametric inference to model the learning. To specify the Bayesian model, we first set our prior distribution for p, and then update the prior distribution using the collected data. The posterior distribution for p immediately follows from the prior distribution and the likelihood function that generates the data. From the posterior distribution, we take our point estimator as the posterior mean.

Bayesian learning: The prior

Initially, p is known only to the extent of a prior belief, which captures the platform's knowledge about p. This includes any information about the items' intrinsic values and relationships between the values of items. It is expressed in our model through the Dirichlet distribution:

$$q^0 = (q_r^0)_{r\in R} \sim Dir(\alpha^0),$$

where $\alpha^0 = (\alpha_r^0)_{r\in R}$ represents the concentration parameters. Note q^0 itself is a random vector, and there is no restriction imposed on the concentration parameters α_r^0, $r \in R$ as long as they are positive scalars. The Dirichlet distribution is a generalization of the Beta distribution to the multivariate case. The shape of the distribution is determined by the concentration parameters α^0, and different concentration parameters can be used to accommodate different prior information. For example, $\alpha_r = 1,\ \forall r \in R$, corresponds to the uniform prior. Jeffreys prior, a commonly used non-informative prior, also can be accommodated by letting $\alpha_r = \frac{1}{2},\ \forall r \in R$. In our various examples throughout the paper, and in the empirical setting of Section 5, we focus on cases in which all items have binary ratings. In such cases, the prior distribution is represented as

$$q^0 = (q^0_{(1,\cdots,1)}, \cdots, q^0_{(0,\cdots,0)}) \sim Dir(\alpha^0_{(1,\cdots,1)}, \cdots, \alpha^0_{(0,\cdots,0)}).$$

The Dirichlet distribution is a widely accepted way to describe prior knowledge in settings like this. Most importantly from our perspective, it is a conjugate prior for the multinomial distribution (see Diaconis et al. (1979)), so it is analytically and computationally tractable. Second, from a theoretical standpoint, the neutrality property of the Dirichlet distribution has an implication for the reliability of user-generated data. As shown in Lee (2021), when the platform uses the Dirichlet distribution as a prior to learn the correlation structure, users cannot manipulate the platform's prediction system by providing false information about their past experiences. Other justifications for the selection of the Dirichlet prior can also be found in Mosimann (1962), Rothschild (1974), Böge and Möcks (1986) and Diniz et al. (2016).

Bayesian learning: Data and likelihood function

For X, let X' refer to the collection of all ratings excluding user 1's history and other items irrelevant in making a prediction, i.e., item $C+2$ to item $C+I$. That is, X' is the $(C+1)\times(N-1)$ submatrix of X which is obtained from X by removing the $C+2$ to $C+I$ rows of the first column. When the condition for complete data is satisfied, data from the $N-1$ previous users is simply a collection of outcomes each of which is independently generated according to the unknown probability vector p. We use y_r to denote the occurrences of r in X'. Let $y = (y_r)_{r\in R}$. The likelihood function associated with the data is

$$y|q^0 \sim Multinomial(N-1, q^0).$$

Bayesian learning: The posterior

Finally, it can be verified that the posterior distribution induced from the prior distribution and the likelihood is a Dirichlet distribution with a concentration parameter $y + \alpha^0$:

$$q^1 \overset{d}{=} q^0 | X' \sim Dir(y + \alpha^0).$$

In our main example, Table 1, we have $X' = \{(1,1),(0,0),(1,1),(1,0)\}$ and $y = (2,1,0,1)$. Furthermore, if we take the uniform prior distribution, the resulting posterior distribution is $Dir(3,2,1,2)$.

4.2 Predicting what users like

From the posterior in the previous section, we are now ready to make a prediction, which we define as the conditional probability of the target user liking the target item i given the data. Let x denote the target user's ratings of items, i.e., x is the first column of X. For x, we define two disjoint sets $R_i^+(x)$ and $R_i^-(x)$ as follows:

$$\begin{aligned} R_i^+(x) =& \{r \in R | r_i = 1 \ and \ r_k = x_k, \ \forall k \ s.t. \ x_k \neq \emptyset\} \\ R_i^-(x) =& \{r \in R | r_i = 0 \ and \ r_k = x_k, \ \forall k \ s.t. \ x_k \neq \emptyset\}. \end{aligned}$$

That is, $R_i^+(x)$ is the subset of R that satisfies (1) the rating of the target item i is positive, and (2) the rating of the conditioning item k, $k \neq i$, is the same as the target user's rating of the conditioning item k if the user has left a rating on it. The sets are specific to the target user's ratings and also to the target item. For example, in the case of Table 1, when user 1 is the target user, we have

$$R_2^+((1,\emptyset)) = \{(1,1)\} \ and \ R_2^-((1,\emptyset)) = \{(1,0)\}.$$

On the other hand, if a new target user, user 6, participates in the platform and when item 1 is the target item, the sets now become

$$R_1^+((\emptyset,\emptyset)) = \{(1,1),(1,0)\} \ and \ R_1^-((\emptyset,\emptyset)) = \{(0,1),(0,0)\}.$$

Let the posterior belief q^1 follow $Dir(\alpha^1)$ after learning from data X', where $\alpha^1 = y + \alpha^0$ and y is derived from the data X'. The true probability that the target user likes the target item i is denoted by

$$z_i(x) = \frac{\sum_{r \in R_i^+(x)} p_r}{\sum_{r \in R_i^+(x) \cup R_i^-(x)} p_r}.$$

Let Y_i be a random variable which takes value 1 with probability $z_i(x)$ and 0 otherwise (i.e. $Y_i \sim Bernoulli(z_i(x))$). Using q^1 and x, we define the pointwise estimator $\hat{z}_i^N(x)$ of $z_i(x)$ by its conditional probability

$$\hat{z}_i^N(x) \equiv P[X_i = 1|X] = \frac{P[Y_i = 1, x|X']}{P[x|X']}.$$

We denote the associated posterior predictive distribution by $\hat{\mathbf{z}}_\mathbf{i}^\mathbf{N}(x)$. Note that an equivalent representation for the estimator is the expected value of Y_i conditional on X.

The following proposition characterizes the estimator and the predictive distribution for the parameters, showing they have concise expressions in terms of the concentration parameters. Moreover, the process of Bayesian learning and prediction is *computationally efficient*, in the sense that the new information can be updated by counting.

Proposition 6 *With complete data, the recommender system implies:*

1. *The probability of the target user liking item i is*

$$\hat{z}_i^N(x) = \frac{\sum_{r \in R_i^+(x)} \alpha_r^1}{\sum_{r \in R_i^+(x) \cup R_i^-(x)} \alpha_r^1}.$$

2. *The associated distribution for the target user liking item i is*

$$\hat{\mathbf{z}}_\mathbf{i}^\mathbf{N}(x) \sim Beta\left(\sum_{r \in R_i^+(x)} \alpha_r^1, \sum_{r \in R_i^-(x)} \alpha_r^1\right).$$

3. *For any parameter α^0 for the prior distribution and for all x, $\hat{\mathbf{z}}_\mathbf{i}^\mathbf{N}(x) \to z_i(x)$, in mean-square. Hence, $\hat{z}_i^N(x) \to z_i(x)$.*

The proposition implies asymptotic learning occurs in our model, so that the prediction becomes more and more precise as we use more data. The Bayesian approach we present is consistent and robust in the sense that the predictions converge to the true probabilities regardless of the choice of the prior distribution within the Dirichlet family. Furthermore, for properly chosen concentration parameters α^0, such as parameters of the uniform prior or the Jefferey prior, they also satisfy the *unbiasedness* condition presented in Definition 5.

The Bayesian model we propose can immediately accommodate generic recommendations. A generic recommendation ignores the ratings left by the target user when making a prediction, even if the target user has left some ratings on items other than the target item. That is, the system takes $x = \emptyset$ in processing the prediction. Denote by $\hat{z}_i^G$, the generic prediction, i.e.,

$$\hat{z}_i^G = \frac{\sum_{r \in R_i^+(\emptyset)} \alpha_r^1}{\sum_{r \in R} \alpha_r^1}.$$

As mentioned previously, the generic recommendation is closely related to the average rating mechanism that displays average ratings to users. Let $\hat{x}_i^{AR}$ denote a prediction that the average rating mechanism provides. When the data from the previous users can be summarized by a multinomial outcome y, it is defined as

$$\hat{z}_i^{AR} = \frac{\sum_{r \in R_i^+(\emptyset)} y_r}{N-1}.$$

The prediction mechanism we propose is general in the sense that $\hat{z}_i^{AR}$ can be obtained from an affine transformation of $\hat{z}_i^G$. Formally, we can show the following result:

Proposition 7 *1. For any data X, $\hat{z}_i^{AR}$ can be recovered from $\hat{z}_i^G$ via an affine transformation.*

2. $\hat{z}_i^G$ and $\hat{z}_i^{AR}$ are asymptotically equivalent.

They are asymptotically equivalent in the sense that the two estimators converge to each other as the recommender systems accumulate more data about the target item. The two approaches differ only by terms that are determined by the selection of the prior parameter, α^0. To see this, we can rewrite $\hat{z}_i^G$ as

$$\hat{z}_i^G = \frac{\sum_{r \in R_i^+(\emptyset)} y_r + \sum_{r \in R_i^+(\emptyset)} \alpha_r^0}{N - 1 + \sum_{r \in R} \alpha_r^0}.$$

The two estimates are related according to the following equation:

$$\hat{z}_i^G = \underbrace{\frac{\sum_{r \in R} \alpha_r^0}{\sum_{r \in R} \alpha_r^0 + \sum_{r \in R} y_r}}_{weights\ on\ prior} \underbrace{\frac{\sum_{r \in R_i^+(\emptyset)} \alpha_r^0}{\sum_{r \in R} \alpha_r^0}}_{prior\ point\ estimate} + \underbrace{\frac{\sum_{r \in R} y_r}{\sum_{r \in R} \alpha_r^0 + \sum_{r \in R} y_r}}_{weights\ on\ observations} \hat{z}_i^{AR}. \tag{8}$$

From (8), it can be easily checked that the two estimates will be very similar when we have a reasonable amount of data. In a small sample situation, the difference between the two estimates remains small if we use a flat or uninformative prior.

4.3 Example

To illustrate how the recommender system works, we revisit the example in Table 1. We make predictions for user 1 using the Bayesian recommender system and the generic recommender system. The set of outcomes in this example is given by $R = \{(1,1),(1,0),(0,1),(0,0)\}$. Recall we have $X' = \{(1,1),(0,0),(1,1),(1,0)\}$ and $x = \{1\}$. Consider the prediction of whether the target user will like item 2. The resulting posterior distribution is

$$q^1 \sim Dir(\alpha^0_{(1,1)} + 2, \alpha^0_{(1,0)} + 1, \alpha^0_{(0,1)}, \alpha^0_{(0,0)} + 1).$$

Here, since we have $x = \{1\}$, $R_2^+(x) = \{(1,1)\}$ and $R_2^-(x) = \{(1,0)\}$. Thus, the prediction can be calculated as

$$\hat{z}_2^5(1) = \frac{\alpha^0_{(1,1)} + 2}{\alpha^0_{(1,1)} + 2 + \alpha^0_{(1,0)} + 1}.$$

For example, if we assume the uniform prior for the initial Dirichlet distribution, we have $\hat{z}_2^5 = \frac{3}{5}$.

On the other hand, under the generic recommender system, we have $x = \{\emptyset\}$.

$$\hat{z}_2^G = \frac{\alpha^0_{(1,1)} + 2 + \alpha^0_{(0,1)}}{\alpha^0_{(1,1)} + 2 + \alpha^0_{(1,0)} + 1 + \alpha^0_{(0,1)} + \alpha^0_{(0,0)} + 1} = \frac{1}{2}.$$

Lastly, under the average rating mechanism, it can be easily shown that $\hat{z}_2^{AR} = \frac{1}{2}$.

5 Evidence from data

In this section we estimate the value a recommender system creates and provide its decomposition and marginal value according to our theoretical findings using the Jester dataset released by AUTOLAB.[12]

The Jester dataset contains anonymous ratings of 100 jokes from 73,421 users, collected over the period from April 1999 to May 2003. Participants choose their ratings via a rating bar over the interval $[-10, 10]$. The dataset contains their recorded ratings, which are rounded to two decimal places. To fit the data into our environment, we convert the data to a binary rating: positive ratings and negative ratings.[13] We call this the Jester binary dataset. Not all users rate all jokes, with around 40% of the ratings out of the total 7,342,100 being missing. Although the sparsity is 40%, there are 14,116 users who have completed rating on all 100 jokes.

Throughout the empirical analysis, we assume our theoretical framework holds, and take the Bayesian learning and prediction model in the previous section as the estimator of the recommender system.[14] There are three key parameters, C, I, τ, which define such a recommender system (recall these are the number of target items, the threshold, and the number of non-target items). As a benchmark setting, we take $M = 10$ and assume $C = 9$ and $I = 1$ with $\tau = \frac{1}{2}$, as well as a fixed way of running simulations, as we will now explain.

For a given counterfactual experiment, we run 1,000 simulations. In each simulation, we randomly select 10,000 users (the training group) from the 14,116 users who have completed rating on all 100 jokes. We then randomly select another 10,000 users (the test group) from the remaining 63,421 users. In this benchmark setting, we assume $I = 1$ and choose nine conditioning items (i.e.

[12] Goldberg et al. (2001)

[13] There are 4,116 zero ratings and they are converted to negative ratings, reflecting that out of the total of 4,136,360 ratings, 2,418,393 are positive ratings and 1,717,967 are negative ratings, with a zero rating being below both the average and the median rating.

[14] We will use the uniform distribution for the prior distribution.

jokes) at random, i.e., $C = 9$.[15] The Bayesian mechanism learns the correlation structure about the ten items using the training group's data. As we will see, ten items turns out to be sufficient to learn the value that the recommender system creates.[16] After learning the correlation structure from the training group, we make predictions about the test group's experiences with the target item. We assume $(v_1, v_0) = (1, -1)$ and apply the user-optimal threshold $\tau^u = \frac{1}{2}$. Thus, a try-recommendation is made to the target user only if the target item is more likely to induce a positive experience with the user (i.e., if the predicted probability of a positive rating is above 1/2). The simulation assumes the recommended item is tried. And we take the user's actual rating (either one or negative one) as the resulting utility to the user. If the predicted rating is strictly below 1/2, the item is assumed to be not recommended to the user and so not tried. The resulting utility is recorded as zero in this case. We consider 1,000 such simulations, each time with a different random selection of the training group and test group, and a different random selection of the target item and nine other items.

Two remarks are in order in regards to the target user's data for this benchmark approach. First, for any $C > 0$, since target users often do not have ratings for all C items, the actual degree of customization in effect is smaller than C. Second, when a target user does not have a rating on the recommended item, we ignore that user in calculating the average of users' utility. To examine any selection bias that might arise from this screening, we have run simulations using the complete subset of the whole data and obtained essentially the same results. The results from a complete dataset is presented in the Online Appendix C.

5.1 Value of data and its decomposition: Evidence

To quantify our theoretical results from Section 3.2, we first investigate how data increases user welfare through the recommender system for the Jester binary dataset.

Our theoretical findings imply that recommender systems add value via three key functions: customization, selection and screening. To measure the value created from each of the functions, we run the benchmark case but vary C, I and τ, each from their benchmark levels ($C = 9,\ I = 1,\ \tau = \frac{1}{2}$). To do so we consider the alternative values $C = 0$, $I = 3$ and $\tau = 0$, respectively, and consider the set of all combinations of parameters implied by these alternatives: $\bar{C} \times \bar{I} \times \bar{\tau}$, where we take $\bar{C} = \{0, 9\}$, $\bar{I} = \{1, 3\}$ and $\bar{\tau} = \{0, \frac{1}{2}\}$. Other things equal, changes in C, I and τ will respectively capture the value created from *customization*, *selection* and *screening*. We refer to a recommendation under $C = 0$ as a generic recommendation (generic RS) and that under $C = 9$ as a customized recommendation (customized RS).

Table 2 presents the estimated average utility that users receive from recommender systems

[15] All the random samplings within each simulation are done without replacement.

[16] According to our theoretical results in Section 3.4, the user surplus (weakly) increases in the number of items that the prediction mechanism conditions on. However, more items require more data points to ensure asymptotic learning. With our dataset, it can be checked that ten items is sufficient to approximate maximal learning.

		Average Utility	*Standard Error*	*Min / Max*
$I = 1$, $\tau = 0$	*Generic RS* ($C = 0$)	0.148	(0.0077)	-0.567/0.616
	Customized RS ($C = 9$)	0.148	(0.0077)	-0.567/0.616
$I = 1$, $\tau = 1/2$ *(scr only)*	*Generic RS*	0.193	(0.0057)	-0.159/0.616
	Customized RS	0.258	(0.0049)	-0.051/0.622
$I = 3$, $\tau = 0$ *(sel only)*	*Generic RS*	0.342	(0.0048)	-0.284/0.617
	Customized RS	0.357	(0.0049)	-0.295/0.651
$I = 3$, $\tau = 1/2$ *(scr and sel)*	*Generic RS*	0.344	(0.0046)	-0.133/0.617
	Customized RS	0.371	(0.0042)	-0.013/0.634

Table 2: User surplus from recommender systems
scr: screening, sel: selection
Standard Error=sample standard deviation/$\sqrt{number\ of\ simulations}$

with different parameter settings. Without selection and screening, i.e., $I = 1$ and $\tau = 0$, the simulation results show that the average utility of users is 0.148 without any customization. This corresponds to the utility users should expect when there is no recommender system available and instead they try each of the randomly selected items. On top of this baseline value, screening ($\tau = 1/2$) adds 0.045 additional average utility when there is no customization ($C = 0$) and 0.110 under customization ($C = 9$). This is for the baseline case without selection ($I = 1$). This shows, without selection, out of the total 0.110 increase in average utility from learning, generic learning contributes 0.045 to the increase in average utility, where as 0.065 is created from customization, suggesting more value comes from within-user customization than from across-user learning about the target item.[17] On the other hand, adding selection without screening (so focusing on the case with $\tau = 0$ but comparing $I = 3$ with $I = 0$), adds 0.194 without customization and 0.209 with customization. The additional value from customization is less pronounced when there are multiple jokes to select from. This reflects that in the jokes database, even with a small number of jokes, there is a high likelihood of there being at least one joke which most people like.[18] This is consistent with our Proposition 3, since items that have positive (or negative) ratings regardless of history cannot add significant amounts of value through customization. The recommender system adds even more value to users when selection, screening and customization coexist. When $I = 3$, $\tau = \frac{1}{2}$, the consistent recommender system adds 0.223 more value to users compare to the situation without a recommender system.

To illustrate Corollary 1 with our data, we analyze the average utility of users who have multi-

[17]However, note that customization creates no value unless it is combined with across-user learning since without data from multiple users there is no way to learn the correlation structure required for customization. So in this sense, the value created from customization augments the value created from having data on many users and can really be thought of as value created from combining customization with across-user learning.

[18]More than 50% of all 100 jokes have above 60% positive ratings.

unit demand. In this case, under $I = 3$, we consider two cases — that the recommender system recommends at most two items or that it recommends all items above the threshold. These correspond to the situation that users have multi-unit demand for two items, or in the latter case, for all three items. As we normalize utility from not consuming an item to zero, each user's final utility is an average of their ratings (-1 or 1) from the recommended items and zero from the outside option when they are not recommended any item. The resulting average utility is depicted in Table 3. Overall, the resulting average utility behaves similarly to the case in which users are only interested in consuming a single item, with customization adding positive value to users, consistent with the corresponding theoretical finding.

		Average Utility	*Standard Error*	*min / max*
$I = 3,\ \tau = 1/2$ *up to 2 units*	*Generic RS*	0.262	(0.0046)	-0.104/0.584
	Customized RS	0.308	(0.0039)	-0.001/0.598
$I = 3,\ \tau = 1/2$ *up to 3 units*	*Generic RS*	0.257	(0.0051)	-0.394/0.584
	Customized RS	0.259	(0.0052)	-0.396/0.580

Table 3: User surplus with multi-unit demand

5.2 Marginal value of learning

Our empirical findings in the above section confirm that each of the three functions of a recommender system add significant value to users. In this section we evaluate the impact of marginal changes in each of the three functions. Furthermore, we assess the marginal value of having data from an additional user.

Marginal value of customization

Proposition 3 predicts additional customization increases user welfare under the user-optimal threshold level. To measure how much user surplus increases from additional customization, we run simulations following our benchmark setting ($I = 1,\ \tau = \frac{1}{2}$) but changing the degree of customization C from one to nine. Thus, for a given simulation with 10,000 training group users, 10,000 test group users, one target item and nine conditioning items, all randomly drawn, we first let the Bayesian recommender system learn from the training group users' ratings on the target item only, and make predictions about the test group users. This corresponds to the situation in which zero degrees of customization have taken place ($C = 0$ case). Next, we increase C by one ($C = 1$), letting the system learn from the same training group user data about the two-item correlation structure associated with the target item and one of the items from the nine-item selection before it makes customized predictions about the test group. We repeat this until $C = 9$. In short, we increase the degree of customization in recommendations from zero to nine and compare the resulting average utilities, keeping everything else fixed according to the benchmark setting. The final user utility is calculated as the average utility over 1,000 such simulations.

Figure 3 summarizes the average utility of users from different degrees of customization and the corresponding 95% confidence intervals from running the simulation 1,000 times. For comparison purposes, the average utility users receive in the absence of a recommender system is also presented in the plot ("w/o RS" — the average utility under $I = 1$ and $\tau = 0$).

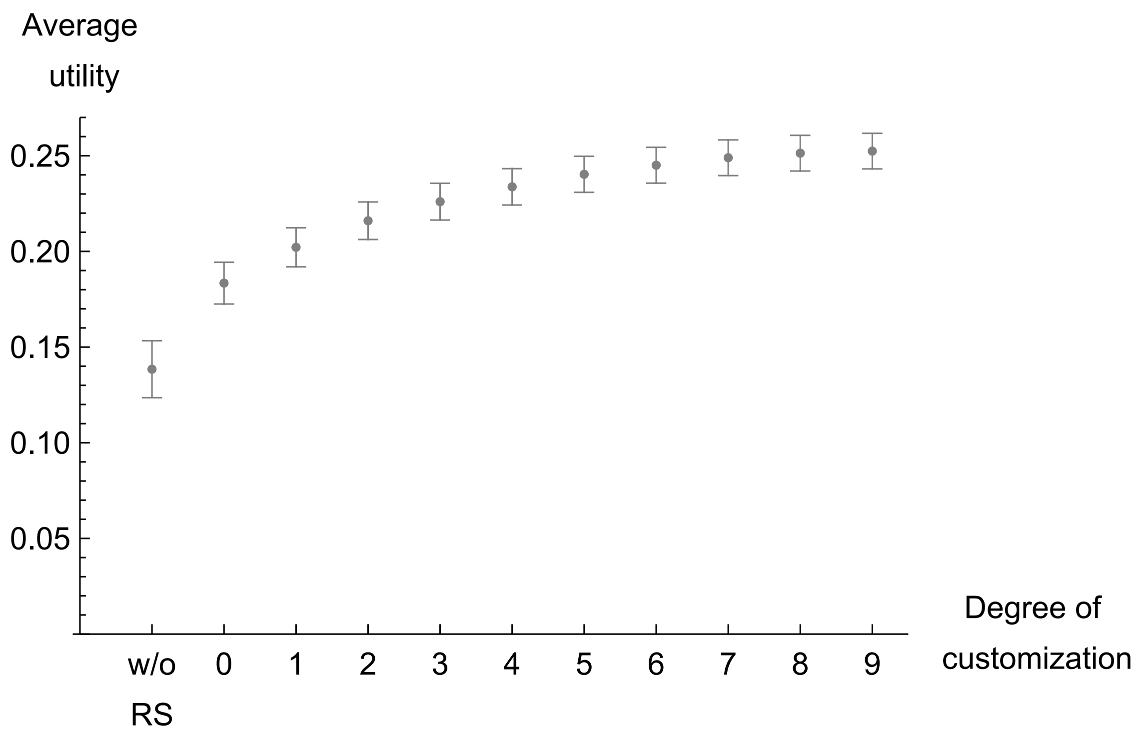

Figure 3: Average utility for different degrees of customization

Consistent with our theoretical predictions in Section 3.4, the user value increases as the recommender system provides more customized predictions. It is also observed that the average utility exhibits a diminishing return to customization when averaged over the 1,000 simulations. However, for any particular set of items, we know from Section 3.4, that the increment can sometimes increase rather than decrease in the degree of customization. We confirmed this is true in our data by considering a single draw of ten items and inspecting how the increment of average utility changes in the degree of customization. The details are presented in Online Appendix D.

Marginal value of selection

We turn to measuring the marginal value to users of additional target items. As predicted in Proposition 1, wider availability of selection makes it easier for the system to find a better item for users. To quantify the marginal value, based on our benchmark setting, we change the number of target items (I) from one to fifteen, and evaluate the resulting average utility of users. Figure 4 depicts the average utility of users in terms of I under $C = 9$ and $\tau = \frac{1}{2}$ and the corresponding 95% confidence intervals.

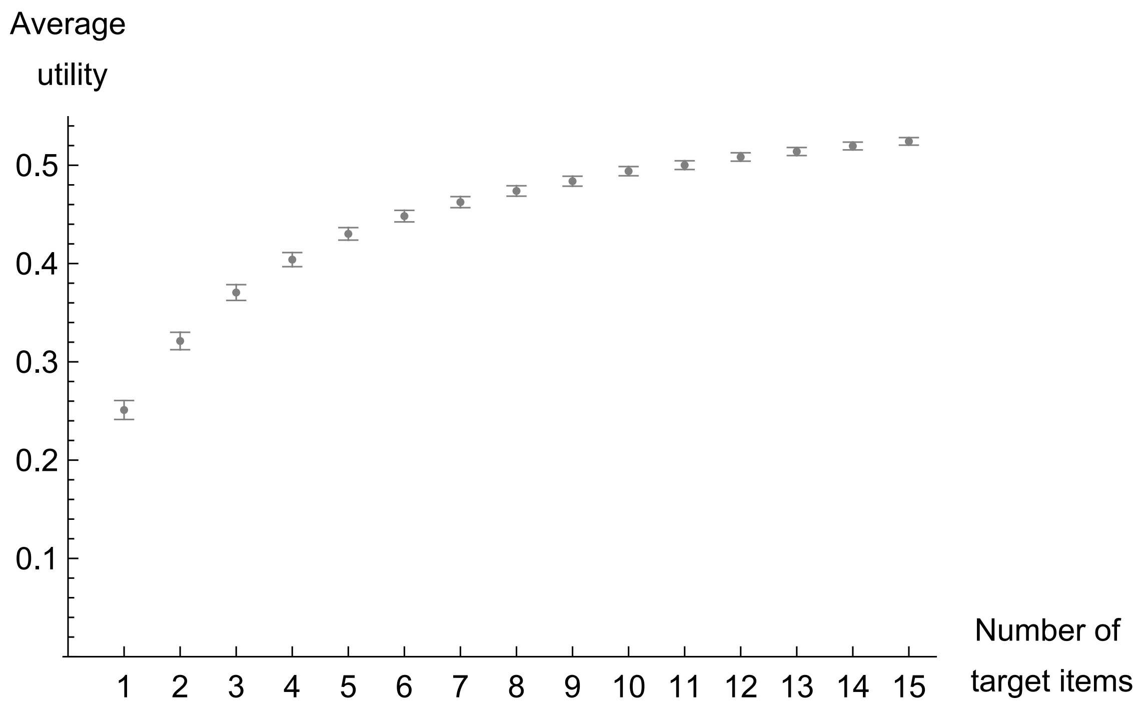

Figure 4: Average utility in terms of the number of target items

Figure 4 also shows that the marginal benefit of having an additional target item to choose from diminishes as more items become available for selection. This is related to the diminishing marginal increment property of the mean of the largest order statistic. When I items are drawn at random according to a fixed distribution, it can be shown that the mean of the largest order statistic increases in the number of items, I, at a diminishing rate. Although we do not explicitly assume any distribution behind the item's selections, we believe a similar logic also applies to our case.

Marginal value of screening

In the model, the recommender system adjusts its level of screening by changing the threshold, which in our framework is the only channel through which a possible misalignment of interest between the platform and the users can arise. We explore what happens when the platform-optimal threshold differs from that which is best for a user ($\tau^p \neq \tau^u$) at each degree of customization. We evaluate the value to users under different platform-optimal thresholds, $\tau^p = 0.1,\ 0.2, \cdots, 0.9$. For all other parameters, we stick to our benchmark setting. Note a low value of τ^p could capture a platform that is biased towards usage because it is compensated based on users trying the item rather than whether they like it or not, while a high value of τ^p could capture an overly conservative platform that does not want the user to try the item unless it is very confident the user will like it.

For each threshold τ^p, we repeat 1,000 simulations exactly as in our benchmark setting, changing the degree of customization from zero to nine. We find that the reduction in average utility associated with the misalignment becomes more significant and the customization becomes less valuable as the level of misalignment increases. The simulation results are summarized in Figure 5.[19] At the maximum degree of customization we test, $C = 9$, the average utility is 0.148, 0.191, 0.204, 0.031 when $\tau^p = 0.1,\ 0.3,\ 0.7,\ 0.9$, which are 0.110, 0.057, 0.053, 0.221 lower than the average utility the optimal threshold $\tau^p = 0.5$ creates. Overall, we observe around 10.9% of reduction in average utility if we lower τ^p by 0.1 from the optimal level $\tau^p = 0.5$. On the other hand, if we increase τ^p by the same margin, there is around a 22.0% reduction in average utility.[20]

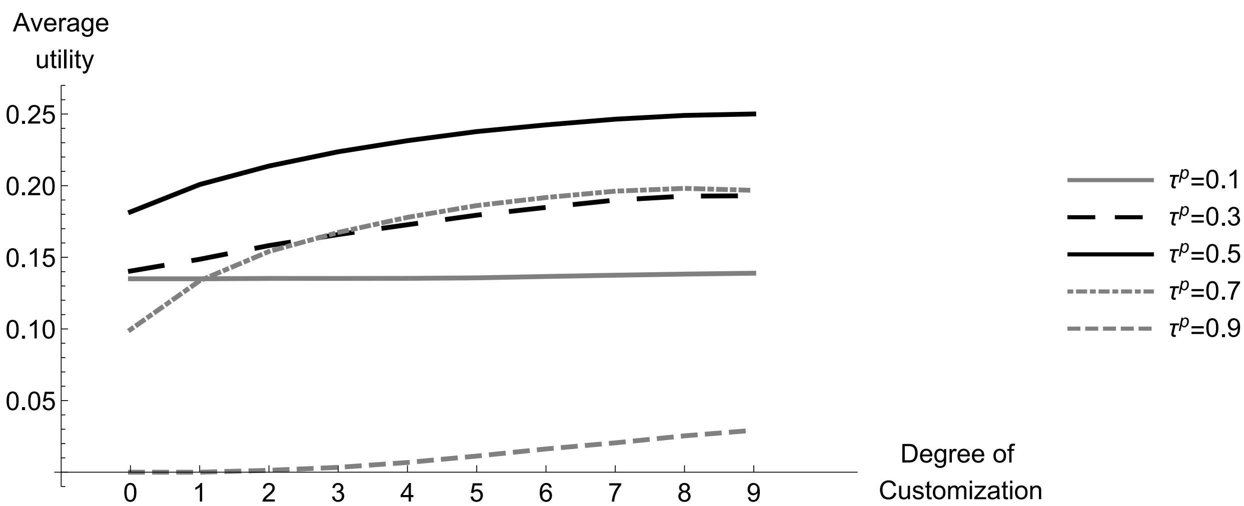

Figure 5: Average utility under different threshold levels τ^p

Additionally, a higher degree in customization leads to a higher average utility of users, as can be seen from the increasing nature of the curves in Figure 5. However, the increased utility from customization is less when the platform's threshold is misaligned with users. Especially, when $\tau^p < 1/2$, the gain from customization is much more limited due to the misalignment: Customization contributes 0.06 to the increase in average utility when $\tau^p = 1/2$, whereas it only contributes 0.003 to the increase when $\tau^p = 0.1$ and 0.048 when $\tau^p = 0.3$. Although our theoretical analysis suggests that when τ^p is far away from the user-optimal level of 1/2 it is possible for customization to harm consumers, that situation doesn't arise in our data.

[19]To ensure the figure is readable, the statistical significance of the simulation results is relegated to Online Appendix E.

[20]The findings that the average utility stays positive even with $\tau^p = 0.1$ and that there is a larger utility loss from a high threshold than a low threshold reflect that the overall average utility from our rating data is positive. Since we normalize the outside option of not trying a joke to zero, the average utility from trying all of any set of jokes recommended will tend to be strictly positive.

Marginal value of additional users

We study how the average utility of users changes as the number of previous users that the platform can learn from increases. For one simulation, we start as usual by randomly selecting the training group (10,000 users) and test group (10,000 users) and ten items according to our benchmark setting, i.e., $(C, I, \tau) = (9, 1, 1/2)$. Then for each given simulation, we run 1,000 rounds of predictions as follows. We start by taking only one user from the training group, to form the training *set*. We run our usual prediction exercise but with the training group replaced by this training set. In each subsequent round, we pick ten new users at random from the training group and add that user data to the existing training set to form a larger training set. We repeat our usual prediction exercise with the corresponding training set in each round. This is repeated until all 10,000 users have been added. This way we can measure how the size of the training set affects the resulting user welfare from the test group. Note as usual, the test group and the set of ten items remain fixed for a given simulation. All other aspects remain the same as in our benchmark setting, and we run 1,000 such simulations.

Figure 6 depicts average user utility in terms of the number of previous user data points, with the corresponding 95% confidence interval being represented by the shaded area. After all 10,000 users' data has been used, the average target user utility is around 0.260, which is around 78.87% higher than the average target user utility when no learning has taken place (i.e., when only one user's data has been used to train the model). Furthermore, it is clear that the marginal increment in the average user utility diminishes as we increase the data size as is predicted in Section 3.5.

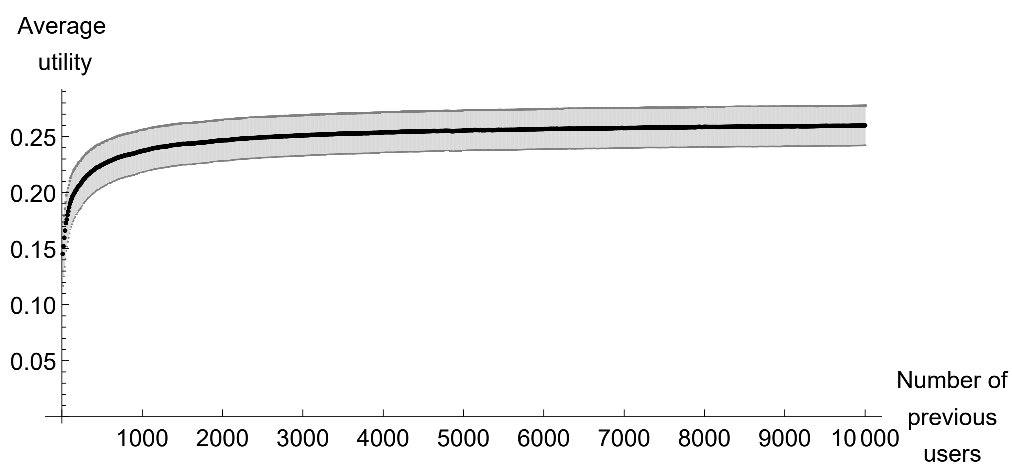

Figure 6: Average utility in terms of number of previous users

5.3 Data complementarity

Lastly, we assess whether the number of previous users and the degree of customization are complements or substitutes in creating user value. We use our benchmark setting with $(I, \tau) = (1, 1/2)$ to run 1,000 simulations where we adjust the degree of customization C and the number of users used in the training group N from their usual values of $C = 9$ and $N = 10,000$. Let $V(C, N)$ be the average utility generated by the Bayesian recommender system that learns from N previous users' data about $C + 1$ number of items. To investigate how the two relevant dimensions interact, we define and measure the following discrete version of the cross partial of V with respect to C and N:

$$\Delta_{C,N} = \big(V(C + \delta_C, N + \delta_N) - V(C + \delta_C, N)\big) - \big(V(C, N + \delta_N) - V(C, N)\big).$$

Here δ_C and δ_N are increments in the degree of customization and in the number of previous users. Depending on whether the cross partial difference is positive or negative, it can be evidence that the two dimensions are complements or substitutes. For instance, if $\Delta_{C,N} > 0$, this means that over the relevant range, using more items to customize predictions on makes having additional user data even more valuable. Moreover, the magnitude of the cross effects at each point (C, N) can be captured by measuring the absolute value of $\Delta_{C,N}$. In our analysis, we take $\delta_C = 1$ and $\delta_N = 500$. Figure 4 presents the value of our measure of the cross-partial derivative $\Delta_{C,N}$ for $C \in \{0, 1, \cdots, 9\}$ and $N \in \{0, 500, 1000, \cdots, 10000\}$. Given we take $V(C, N)$ as the average utility for each combination of (C, N), the significance level is not presented in this figure. The area depicted in dark blue represents combinations of (C, N) which generate positive values of the cross-partial derivative (above 0.0005), so in which the two dimensions are complements. On the other hand, the area in white represents combinations of (C, N) which generate negative values of the cross-partial derivative (below -0.0005), so in which the two dimensions are substitutes. Lastly, the area with light gray represents combinations of (C, N) around zero, within the error bound of $(-0.0005, 0.0005)$.

The two dimensions exhibit strong complementarity when both the degree of customization and number of users in the training group are low, as shown by the peak in Figure 7 when C and N are both low. Adding data from additional users provides more value when there is more customization starting from the point where both are relatively low. On the other hand, when the degree of customization is high and the training group size is very small, the two dimensions exhibit substitutability. More specifically, at $N = 0$, the cross marginal effect is $\Delta_{6,0} = -0.0011$, $\Delta_{7,0} = -0.0026$, and $\Delta_{8,0} = -0.0035$.[21] This reflects the lack of degrees of freedom when N is small and C is large. When C increases, the number of parameters in the correlation structure increases, and hence, the recommender system requires a larger number of users to achieve maximal learning (although the average utility achievable from maximal learning is greater). Therefore, when the amount of users is insufficient, a moderate degree of customization can induce higher average utility than a high degree of customization. In fact, when $N = 500$, the average utility is maximized at $C = 6$. As we increase N, the optimal C also increases, until the average utility is eventually maximized by $C = 9$ for all $N \geq 4000$. Finally, we note that regardless of the degree of

[21] A similar result holds if we increase the size of the training set from $N = 1$ to $N = 500$.

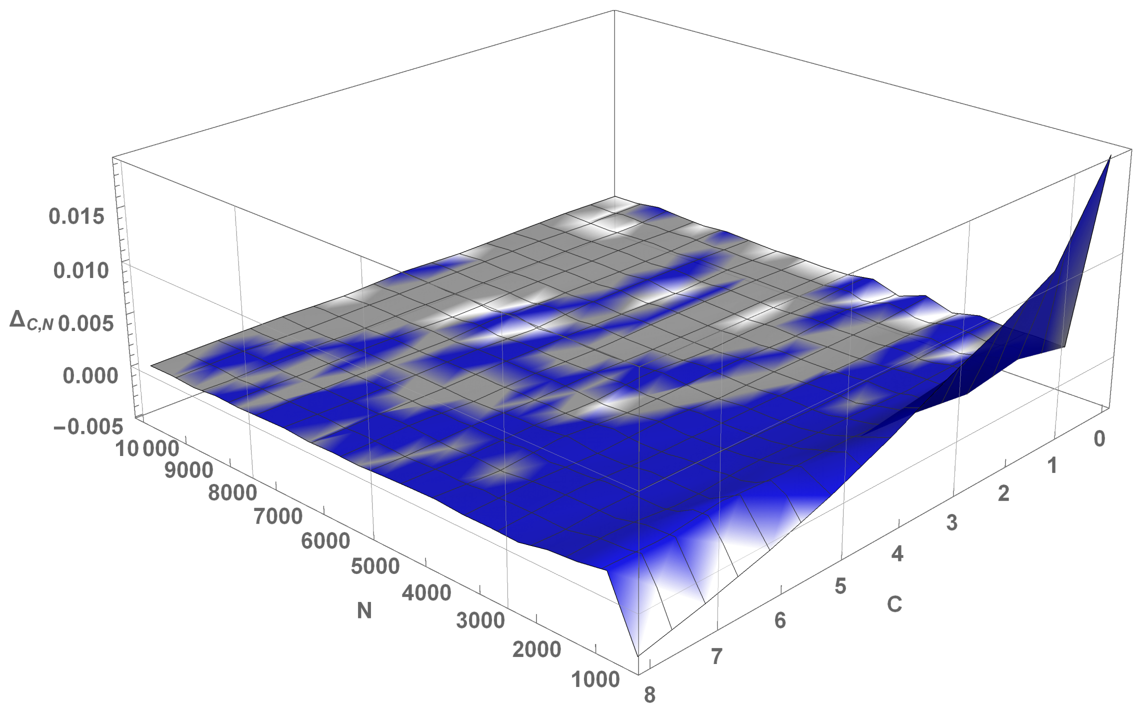

Figure 7: Complementarity between the two dimensions

customization, the cross effects largely disappear once there are enough users to learn from. In our exercise, once $N \geq 5000$, the magnitude of any cross marginal effect is less than 0.0005.

6 Extensions

This section briefly explores three extensions of our baseline model. In Section 6.1 we discuss the platform's and users' exploration motive. In Section 6.2 we provide a micro-foundation for user behaviors in providing ratings and deciding whether to follow the platform's recommendation. In Section 6.3 we explore how the marginal return to data can sometimes be increasing using simulated data.

6.1 Within-user exploration

Throughout the paper we have ignored the motive the platform and users have to explore. From a dynamic standpoint, there are two channels by which exploration can potentially take place. First, although the target item is not suitable for the target user, the platform may want the user to try the item to improve its predictions for future users. This is closely related to the standard exploration vs. exploitation issue that is covered in the literature of the multi-armed bandit.[22] For

[22]There is a growing literature on the multi-armed bandit problem in which rewards from the arms are correlated with each other (Gupta et al. (2019), Pandey et al. (2007) and Srivastava et al. (2015)), which is similar to exploration in the correlated items setting we are dealing with. However, the optimal policy for a multi-armed bandit with a general correlation structure, which corresponds to our setting, remains unknown.

a large part of our theoretical results, we focused on limit results, where across-user exploration by the platform was not an issue since the platform was already assumed to have learnt the true correlation structure between items and was only deciding whether to recommend one or more item to a target user.

However, for certain results including our empirical findings, exploration could provide additional value from data that we ignored.

The second channel arises because the platform provides a customized recommendation in our framework. Although the target item is seemingly unattractive to the target user, it is possible that the user strictly benefits from trying the item if the item has a strong correlation with another item (or set of items) that the user has not tried yet and would like to learn about. One might think of a user trying an item that is representative of a new music or movie genre that they have not tried before and they expect not to like, just so they can be sure they are not missing out on potentially a lot of other items they could actually enjoy.

Since our main focus was only on whether a target user should try one target item or not, the findings of the current paper is not affected by this type of exploration. However, because there is informational externality generated from the target user's experience with a target item which might affect subsequent recommendations, it can be of interest to study the optimal order of recommendation when there are multiple items to try and when users have multi-unit demands. For this reason, we focus on this second channel and study the exploration issue that arises from customization, which we believe is new to the literature.

We first consider a two-item platform that has accumulated data from $N-1$ previous users and makes recommendations to user N. We assume that the recommendations are user-centric in the sense that the platform's threshold aligns with the users. The platform's decision is two-fold: It decides which item to recommend first and the threshold level above which this item is recommended to the user. We call the recommendation policy that recommends the item with the highest myopic expected payoff first the *naïve order*. On the other hand, we refer the recommendation threshold that induces the highest myopic expected payoff as the *naïve threshold*. When the recommendation rule takes into account dynamic considerations in recommending the first item, the optimal rule is referred to as a *dynamically optimal policy*.

In the Online Appendix F we formally show the following findings: (1) for any two-item platforms, the naïve order is dynamically optimal, and (2) the naïve threshold is not dynamically optimal unless the second item to be recommended is expected to produce negative utility regardless of the experience with the first item. The results hold for any discount factor $\delta \in (0, 1]$.

The *naïve threshold* is not dynamically optimal because the threshold does not fully take the dynamic benefit of within-user exploration into account. Even if the first item is expected to harm the user to some extent, if the user experience with the first item releases relevant information about the other item, the first item can be worth trying to get access to the relevant information. This implies the user's dynamically optimal threshold can start lower than the *naïve threshold* when the user has more items to try on the platform, and then eventually becomes the *naïve threshold* when there are no more items to try. Although it is shown that the *naïve order* is dynamically optimal, this no longer has to be true when the platform takes more items into consideration in making sequential recommendations. In the Online Appendix F we show, by means of an example, that when there are three items, the *naïve order* may no longer be dynamically optimal. When a particular item reveals more relevant information about all other items than does any of the other items, users are better off by trying the said item even if the expected utility from the item is strictly less than what they expect to get from some other item.

6.2 Incentive compatibility of user feedback

In our baseline model, we have assumed that the data the platform collects from users is reliable (i.e., that users give truthful feedback), and that the platform's recommendations are effective (i.e., users follow the recommendations from the platform). However, it is conceivable that sophisticated forward-looking users may be able to strategically distort their feedback to manipulate the recommender system and, as a result, receive more useful information from it.

Lee (2021) explores this possibility. He shows that the manipulation motive of a user can sometimes prevail even though the platform is committed to the user-optimal recommendations. The motivation is stronger when the platform has highly unbalanced data, meaning that, using our terminology, the number of data points from history r users is significantly larger than those from history r' users. As we have shown in Proposition 5, this results in quality differences in the recommendations that users with different history receive. As a result, history r' users may have an incentive to masquerade as history r users to receive higher quality information from the platform. Although this strategic deception can strictly benefit some users *ex-post*, we also show that leaving truthful feedback and subsequently following the recommendation is optimal for users in expectation, under a wide variety of user beliefs about the item correlation structure, including uniform priors. Furthermore, it is also shown that the strategic motivation to manipulate the platform by leaving false feedback and (or) not following recommendations vanishes as the platform accumulates more data points from users of both histories, and it completely disappears in the limit. This supports our assumption of truthful feedback in this paper, at least for the settings where we focus on limit results, which was the case for much of our theoretical analysis.

6.3 The shape of learning curves

Positive but diminishing marginal return to data has been found in the field experiment of Claussen et al. (2021) and the empirical studies of Bajari et al. (2019) and Schaefer and Sapi (2021). Although our empirical result using the Jester dataset also exhibits the same property, this is not a property that will always hold for all data or all payoffs. To illustrate, we generate random ratings over three items, item 3 being the target item. We consider a situation that the target item is positively correlated with item 1 and negatively correlated with item 2.[23] We carry out the simulations based on our benchmark setting that estimate the average utility of users in terms of the number of data points similar to the last exercise we conducted in Section 5.2, but under different user payoff structures (v_1, v_0) and the corresponding user-optimal thresholds.[24]

Fixing $v_1 = 1$, we take $v_0 = -1, -3, -5, -7, -10$ and estimate the average utility of users. As summarized in Figure 8, when the disutility from a negative experience outweighs the utility from a positive experience, the return from an initial period of learning exhibits increasing returns, and then it eventually diminishes. The S-shaped learning curves in Figure 8, for example in case $v_0 = -10$, reflect that when recommending the incorrect item induces a very negative outcome for consumers, the user-optimal threshold ensures that recommendations are not made unless the platform is very sure of a positive outcome. This leads to no recommendations being made and zero expected utility until the recommender system has access to enough training data, after which additional data becomes increasingly valuable for a while before the usual diminishing marginal return to data property sets in. Such an S-shaped learning curve may be relevant for applications where the penalty for trying something that users have a bad experience with is very negative, such as for doctors relying on AI-powered medical imaging for recommending conditions or treatments (e.g. Behold.ai), or airport security relying on AI-based readings of X-rays of bags for threat detection (e.g. seeTrue.ai).

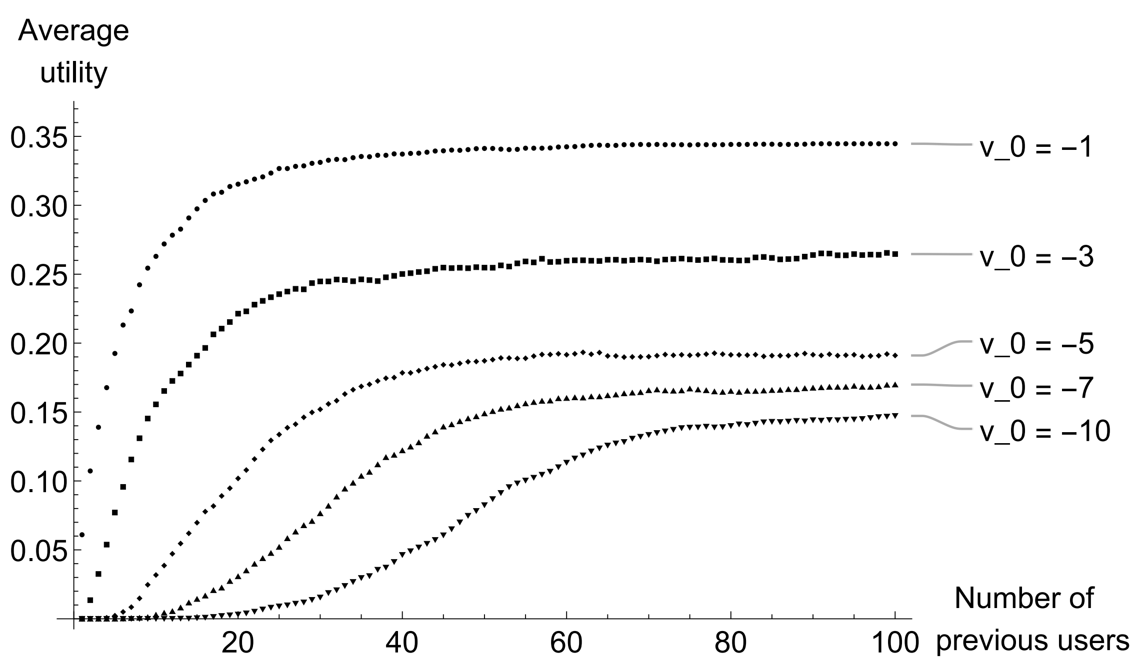

Figure 8: Different shapes of learning curves

[23]The underlying correlation structure is $p = \{0.175, 0.075, 0.2475, 0.0025, 0.0025, 0.2475, 0.075, 0.175\}$.

[24]To be more precise, we generate 73,412 simulated user data points over three items. In each simulation, which we run 1,000 times, N training group users and 10,000 test group users are drawn at random. We take $C = 2$ and $I = 1$, meaning that we make predictions about the target item 3 based on the ratings on item 1 and item 2. We vary N from 1 to 100 to study how learning affects user utility.

7 Conclusion

In this paper we provide a framework to help understand the value created by a recommender system, and use it to identify the underlying sources of value creation, both theoretically and empirically. The value creation process consists of three key elements: history-specific customized predictions, the selection of the best item, and screening items not suitable for users of each history. We theoretically explore how each of these elements adds value to users, and empirically confirm that each of them contributes significantly to the value created. We provide theoretical conditions under which the data network effect generated from a recommender system is positive but diminishes with more users, a result we find also holds empirically based on a publicly available dataset containing over four million anonymous joke ratings from 73,421 users. We show theoretically that the value of customization, while normally positive, can turn negative if the platform's interests diverge enough from the interests of users. While such a divergence does not lower consumer utility when our model is applied to the jokes dataset, the increase in value from customization becomes negligible when the divergence of interests is maximal. And from applying our model to the jokes dataset, we find strong complementarity between the value created by across-user learning and customization when both the number of users and the degree of customization is low, but that this complementarity disappears quickly as the system accumulates data from more users.

We view our framework as a first attempt to model how a recommender system based on collaborative filtering works. There remains much that can be done to build on it. In our extensions we provided some initial analysis of a few directions for future research. These include allowing the platform to optimize its threshold for making each of its recommendations taking into account the future benefits of exploration when it is still learning the correlation structure, and allowing users to decide whether to give truthful feedback or not. In our main theoretical analysis, we sidestepped these issues by largely focusing on limit results. Another direction is to use our prediction framework to say something about how different items should be ranked. Thus, our framework could provide a useful starting point for analyzing optimal ranking mechanisms that take into account collaborative filtering. Finally, it would be useful to apply our Bayesian recommender system to other datasets to quantify the value of data across different environments.

References

Acemoglu, D., Dahleh, M. A., Lobel, I., and Ozdaglar, A. (2011). Bayesian learning in social networks. *Review of Economic Studies*, 78:1201–1236.

Acemoglu, D., Makhdoumi, A., Malekian, A., and Ozdaglar, A. (2021). Too much data: Price and inefficiencies in data markets. *American Economic Journal: Microeconomics*, (Forthcoming).

Aggarwal, C. C. (2016). Knowledge-based recommender systems. In *Recommender systems*, pages 167–197. Springer.

Bajari, P., Chernozhukov, V., Hortaçsu, A., and Suzuki, J. (2019). The impact of big data on firm performance: An empirical investigation. In *AEA Papers and Proceedings*, volume 109, pages 33–37.

Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817.

Bergemann, D., Bonatti, A., and Gan, T. (2021). The economics of social data. Working paper.

Bergemann, D. and Ozmen, D. (2006). Optimal pricing with recommender systems. Proceedings of the 7th ACM Conference on Electronic Commerce, Ann Arbor MI.

Biglaiser, G., Calvano, E., and Crémer, J. (2019). Incumbency advantage and its value. *Journal of Economics & Management Strategy*, 28(1):41–48.

Böge, W. and Möcks, J. (1986). Learn-merge invariance of priors: a characterization of the dirichlet distributions and processes. *Journal of multivariate analysis*, 18(1):83–92.

Burke, R. (2000). Knowledge-based recommender systems. *Encyclopedia of library and information systems*, 69(Supplement 32):175–186.

Calvano, E., Calzolari, G., Denicolò, V., and Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10):3267–97.

Che, Y.-K. and Hörner, J. (2018). Recommender systems as mechanisms for social learning. *Quarterly Journal of Economics*, 133:871–925.

Chien, Y.-H. and George, E. I. (1999). A bayesian model for collaborative filtering. *The 7th International Workshop on Artificial Intelligence and Statistics, Fort Lauderdale, Florida, 4–6 January 1999.*

Chiou, L. and Tucker, C. (2017). Search engines and data retention: Implications for privacy and antitrust. Technical report.

Choi, J. P., Jeon, D.-S., and Kim, B.-C. (2019). Privacy and personal data collection with information externalities. *Journal of Public Economics*, (May):113–124.

Claussen, J., Peukert, C., and Sen, A. (2021). The editor vs. the algorithm: Economic returns to data and externalities in online news. *Available at SSRN 3479854*.

de Cornière, A. and Taylor, G. (2020). Data and competition: A general framework with applications to mergers, market structure, and privacy policy. Working paper.

Dellarocas, C., Dini, F., and Spagnolo, G. (2006). *Designing reputation mechanisms*, page 446–482. Cambridge University Press.

Diaconis, P., Ylvisaker, D., et al. (1979). Conjugate priors for exponential families. *The Annals of statistics*, 7(2):269–281.

Diniz, M. A., De Bock, J., and Van Camp, A. (2016). Characterizing dirichlet priors. *The American Statistician*, 70(1):9–17.

Elahi, M., Ricci, F., and Rubens, N. (2016). A survey of active learning in collaborative filtering recommender systems. *Computer Science Review*, 20:29–50.

Farboodi, M., Mihet, R., Philippon, T., and Veldkamp, L. (2019). Big data and firm dynamics. *AEA Papers and Proceedings*, 109:38–42.

Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, 41(2):148–177.

Goldberg, K., Roeder, T., Gupta, D., and Perkins, C. (2001). Eigentaste: A constant time collaborative filtering algorithm. *Information Retrieval*, 4(2):133–151.

Gu, Y., Madio, L., and Reggiani, C. (2021). Data brokers co-opetition. Working paper.

Gupta, S., Chaudhari, S., Joshi, G., and Yağan, O. (2019). Multi-armed bandits with correlated arms. *arXiv preprint arXiv:1911.03959*.

Hagiu, A. and Wright, J. (2021). Data-enabled learning, network effect and competitive advantage. Working paper.

Ichihashi, S. (2020a). The economics of data externalities. Working paper.

Ichihashi, S. (2020b). Competing data intermediaries. *RAND Journal of Economics*, (Forthcoming).

Kremer, I., Mansour, Y., and Perry, M. (2014). Implementing the "wisdom of the crowd". *Journal of Political Economy*, 122(5):988–1012.

Lee, G. (2021). Data reliability and collaborative filtering. Working paper.

Lops, P., de Gemmis, M., and Semeraro, G. (2011). *Content-based Recommender Systems: State of the Art and Trends*, pages 73–105. Springer US, Boston, MA.

Lu, J., Wu, D., Mao, M., Wang, W., and Zhang, G. (2015). Recommender system application developments: A survey. *Decision Support Systems*, 74:12–32.

Mosimann, J. E. (1962). On the compound multinomial distribution, the multivariate β-distribution, and correlations among proportions. *Biometrika*, 49(1/2):65–82.

Olszewski, W. and Weber, R. (2015). A more general pandora's rule? *Journal of Economic Theory*, 160:429–437.

Pandey, S., Chakrabarti, D., and Agarwal, D. (2007). Multi-armed bandit problems with dependent arms. In *Proceedings of the 24th international conference on Machine learning*, pages 721–728.

Papanastasiu, Y., Bimpikis, K., and Savva, N. (2018). Crowdsourcing exploration. *Management Science*, 64(4):1727–1746.

Pazzani, M. J. and Billsus, D. (2007). Content-based recommendation systems. In *The adaptive web*, pages 325–341. Springer.

Prüfer, J. and Schottmüller, C. (2021). Competing with big data. *Journal of Industrial Economics*, (Forthcoming).

Rothschild, M. (1974). Searching for the slowest price when the distribution of prices is unknown. *Journal of Political Economics*, 82(4):689–711.

Sarwar, B., Karypis, G., Konstan, J., and Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms. *The 10th International Conference on the World Wide Web*, pages 285–295.

Schaefer, M. and Sapi, G. (2021). Learning from data and network effects: The example of internet search. Working paper.

Schafer, J. B., Frankowski, D., Herlocker, J., and Sen, S. (2007). *Collaborative Filtering Recommender Systems*, pages 291–324. Springer Berlin Heidelberg, Berlin, Heidelberg.

Smith, L. and Sørensen, P. (2000). Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398.

Srivastava, V., Reverdy, P., and Leonard, N. E. (2015). Correlated multiarmed bandit problem: Bayesian algorithms and regret analysis. *arXiv preprint arXiv:1507.01160.*

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica*, 47(3):641–654.

Wolitzky, A. (2018). Learning from others' outcomes. *American Economic Review*, 108(10):2763–2801.

A Appendix

A.1 Proof of Proposition 1 and Corollary 3

Proof. We focus on the user's expected utility evaluated at the limit as $N \to \infty$ in the proof. The same logic can be applied to measure the platform surplus or the social surplus.

Let the user-optimal threshold be $\tau^u = \frac{-v_0}{v_1 - v_0}$. By the definition of a consistent recommender system, i is recommended to history r' user *if and only if* the following inequality holds:

$$\frac{p^i_{(r',1)}}{p^i_{r'}} \geq \max_{j \in \mathscr{I}} \left\{ \frac{p^j_{(r',1)}}{p^j_{r'}},\ \frac{-v_0}{v_1 - v_0} \right\}.$$

Since $p^i_{(r',1)} + p^i_{(r',0)} = p^i_{r'} = p^j_{r'},\ \forall i,\ j \in I$, the condition is equivalent to the following expressions

$$\begin{aligned}
& \frac{p^i_{(r',1)}}{p^i_{r'}} \geq \max_{j \in \mathscr{I}} \left\{ \frac{p^j_{(r',1)}}{p^j_{r'}},\ \frac{-v_0}{v_1 - v_0} \right\} \\
\Leftrightarrow\ & (v_1 - v_0) p^i_{(r',1)} \geq \max_{j \in \mathscr{I}} \left\{ (v_1 - v_0) p^j_{(r',1)},\ \ - v_0 (p^i_{(r',1)} + p^i_{(r',0)}) \right\} \\
\Leftrightarrow\ & v_1 p^i_{(r',1)} + v_0 p^i_{(r',0)} \geq \max_{j \in \mathscr{I}} \left\{ v_1 p^j_{(r',1)} + v_0 p^j_{(r',0)},\ 0 \right\}.
\end{aligned}$$

Since this holds for any $r' \in R'$ and r' happens with probability $p^i_{r'}$, $\forall i \in I$, the result in Proposition 1 follows immediately.

Similarly, for a general threshold level τ^p, i is recommended to a history r' user *if and only if*

$$\frac{p^i_{(r',1)}}{p^i_{r'}} \geq \max_{j \in \mathscr{I}} \left\{ \frac{p^j_{(r',1)}}{p^j_{r'}},\ \tau^p \right\}.$$

An equivalent representation of the above condition is

$$v_1 p^i_{(r',1)} + v_0 p^i_{(r',0)} \geq \max_{j \in \mathscr{I}} \left\{ v_1 p^j_{(r',1)} + v_0 p^j_{(r',0)},\ p_{r'}(\tau v_1 + (1 - \tau) v_0) \right\}.$$

Again, since history r' happens with probability $p_{r'}$, we have the expression in the corollary. ■

A.2 Proof of Corollary 2

Proof. Consider the following expression from (5):

$$p_{(r',0)} \mathbf{1}\left\{ (1 - \tau) p_{(r',1)} \geq \tau p_{(r',0)} \right\} + p_{(r',1)} \mathbf{1}\left\{ (1 - \tau) p_{(r',1)} < \tau p_{(r',0)} \right\}.$$

Regardless of r', the expression is minimized at $\tau = \frac{1}{2}$. Thus, the error bound in (5) is minimized at $\tau^u(r') = \frac{1}{2}$. Taking $\tau^u(r') = \frac{1}{2}$, the expression (5) is equivalent to

$$\begin{aligned}
&\sum_{r' \in R'} p_{r'} \left(\frac{p_{(r',0)}}{p_{r'}} \mathbf{1}\big\{p_{(r',1)} \geq p_{(r',0)}\big\} + \frac{p_{(r',1)}}{p_{r'}} \mathbf{1}\big\{p_{(r',1)} < p_{(r',0)}\big\} \right) \\
&= \sum_{r' \in R'} p_{(r',0)} \mathbf{1}\big\{p_{(r',1)} \geq p_{(r',0)}\big\} + p_{(r',1)} \mathbf{1}\big\{p_{(r',1)} < p_{(r',0)}\big\} \\
&= \sum_{r' \in R'} \min\{p_{(r',1)}, p_{(r',0)}\},
\end{aligned}$$

which completes the proof. ■

A.3 Proof of Proposition 2

Proof. For each history of a target user r', we construct a set of I correlation structures $q = \{q^i\}_{i \in \mathscr{I}}$ under which an extra degree in customization strictly hurts the target user. Since the proof is done by construction, it is without loss of generality to assume $I = 1$. For any cases with $I > 1$, we can simply let q^j, $j \neq C + I = M$, satisfies the following inequality and focus on q^M only:

$$z_j(r') = \frac{q^j_{(r',1)}}{q_{r'}} < \tau(r'), \ \forall j \neq M \ and \ \forall r' \in R'.$$

Let $I = 1$ and $\tau \neq \tau^u$, where τ is the threshold level the system adopts. Similar to the construction of R', we denote the set of all outcomes that can be generated by the first $C - 1$ conditioning items by R''. Similarly, $(r'', k, 1)$ and $(r'', k, 0)$ represents the positive and negative user experience with the target item of a history (r'', k) user, $k \in \{1, \cdots, n_C\}$.

We will show that for any $\tau \neq \tau^u$, there exists q^M such that the target user with history $r'' \in R''$ receive strictly lower utility when the recommendation is customized based on C conditioning items and the target item than when it is customized based on first $C - 1$ conditioning items and the target item. By an induction argument, this will prove the existence of a correlation structure that users are strictly worse off from an extra degree in customization.

Firstly, if $\tau < \tau^u$, then consider the following correlation structure q^M. For each $r'' \in R''$ and $k \in \{1, \cdots, n_C\}$, the following equalities and inequality hold:

$$\begin{cases} \frac{q^M_{(r'',k,1)}}{q^M_{(r'',k)}} = \tau & \text{if } k \neq n_C \\ \frac{q^M_{(r'',k,1)}}{q^M_{(r'',k)}} < \tau & \text{if } k = n_C. \end{cases}$$

Thus, we have $\frac{\sum_{k \in \{1, \cdots, n_C\}} q^M_{(r'',k,1)}}{\sum_{k \in \{1, \cdots, n_C\}} q^M_{(r'',k)}} < \tau$. Under this correlation structure, when the recommendation is fully customized, item M is recommended to all users except the history (r'', n_C) user, whereas when the system omits conditioning item C in making recommendations, no user is recommended to try item M.

The user's utility under fully customized recommendations can be represented as follows using Lemma 1:

$$\begin{aligned}&\sum_{r''\in R''}\sum_{k\in\{1,\cdots,n_C\}} q^M_{(r'',k)}\mathbf{1}\left\{\frac{q^M_{(r'',k,1)}}{q^M_{(r'',k)}}\geq\tau\right\}\left(v_1\frac{q^M_{(r'',k,1)}}{q^M_{(r'',k)}}+v_0\frac{q^M_{(r'',k,0)}}{q^M_{(r'',k)}}\right)\\ =&\sum_{r''\in R''}\sum_{k\in\{1,\cdots,n_C\}} q^M_{(r'',k)}\mathbf{1}\left\{\frac{q^M_{(r'',k,1)}}{q^M_{(r'',k)}}\geq\tau\right\}(v_1-v_0)\left(\frac{q^M_{(r'',k,1)}}{q^M_{(r'',k)}}-\tau^u\right).\end{aligned}$$

Since $\frac{q^M_{(r'',k,1)}}{q^M_{(r'',k)}}=\tau<\tau^u$ for $k\neq n_C$, the user utility is strictly negative for users whose history is not $(r'',n_C,1)$. On the other hand, the history $(r'',n_C,1)$ user does not try the item. By construction, the item is recommended to no users when the system omits C, and the resulting user utility is zero.

Secondly, suppose we have $\tau>\tau^u$. Again, for each $r''\in R''$ and $k\in\{1,\cdots,n_C\}$, consider a correlation structure q^M that satisfies the following equality and inequality:

$$\begin{cases}\frac{q^M_{(r'',k,1)}}{q^M_{(r'',k)}}>\tau & \text{if } k=1\\ \frac{q^M_{(r'',k,1)}}{q^M_{(r'',k)}}=\tau & \text{if } k\neq 1 \textit{ or } n_C\\ \frac{q^M_{(r'',1,1)}+q^M_{(r'',n_C,1)}}{q^M_{(r'',1)}+q^M_{(r'',n_C)}}=\tau. & \end{cases}$$

Under this correlation structure, item M is recommended to all users when the system omits conditioning item C in making predictions. However, if it takes all conditioning items into account, the history (r'',n_C) user does not receive a recommendation. Since trying M is actually beneficial to all users, there is a missing opportunity if recommendations are fully customized. ■

A.4 Proof of Proposition 3

Proof. Without loss of generality, we only look at the case of $\frac{p_{(r',1,1)}}{p_{(r',1)}}\geq\frac{p_{(r',0,1)}}{p_{(r',0)}}$. The opposite case can be shown using the exact same logic.

To begin, note that the expected utility of r' user before the extra degree in customization is

$$\mathbf{1}\left\{\frac{p_{(r',1,1)}+p_{(r',0,1)}}{p_{(r',1)}+p_{(r',0)}}\geq\tau\right\}\left(v_1\frac{p_{(r',1,1)}+p_{(r',0,1}}{p_{r'}})+v_0\frac{p_{(r',1,0)}+p_{(r',0,0)}}{p_{r'}}\right).$$

On the other hand, the expected utility of the user after the extra degree in customization is

$$\frac{p_{(r',1)}}{p_{r'}}\mathbf{1}\left\{\frac{p_{(r',1,1)}}{p_{(r',1)}}\geq\tau\right\}\left(v_1\frac{p_{(r',1,1)}}{p_{(r',1)}}+v_0\frac{p_{(r',1,0)}}{p_{(r',1)}}\right)+\frac{p_{(r',0)}}{p_{r'}}\mathbf{1}\left\{\frac{p_{(r',0,1)}}{p_{(r',0)}}\geq\tau\right\}\left(v_1\frac{p_{(r',0,1)}}{p_{(r',0)}}+v_0\frac{p_{(r',0,0)}}{p_{(r',0)}}\right).$$

Therefore, for each $r' \in R'$, the benefit of the marginal customization is

$$\begin{cases} 0 & \text{if } \tau \le \frac{p_{(r',0,1)}}{p_{(r',0)}} \text{ or } \tau \ge \frac{p_{(r',1,1)}}{p_{(r',1)}} \\ -(v_1 p_{(r',0,1)} + v_0 p_{(r',0,0)})/p_{r'} & \text{if } \frac{p_{(r',0,1)}}{p_{(r',0)}} \le \tau < \frac{p_{(r',1,1)}+p_{(r',0,1)}}{p_{(r',1)}+p_{(r',0)}} \\ (v_1 p_{(r',1,1)} + v_0 p_{(r',1,0)})/p_{r'} & \text{if } \frac{p_{(r',1,1)}+p_{(r',0,1)}}{p_{(r',1)}+p_{(r',0)}} \le \tau < \frac{p_{(r',1,1)}}{p_{(r',1)}}. \end{cases} \tag{Difference}$$

First, let item $C+1$ and the target item are strictly correlated conditional on r'. Because we focus on $\frac{p_{(r',1,1)}}{p_{(r',1)}} \ge \frac{p_{(r',0,1)}}{p_{(r',0)}}$, the two items are strictly positively correlated. That is, we have

$$v_1 p_{(r',0,1)} + v_0 p_{(r',0,0)} < 0 \text{ and } v_1 p_{(r',1,1)} + v_0 p_{(r',1,0)} > 0.$$

Thus, for any r' that does not satisfies (6), the expected utility related to r' is strictly positive.

Conversely, suppose that the two items are not correlated conditional on r'. By definition of correlation, it is either both $v_1 p_{(r',1,1)} + v_0 p_{(r',1,0)}$ and $v_1 p_{(r',0,1)} + v_0 p_{(r',0,0)}$ are positive or both are negative. Thus, there always exists τ such that the utility represented in (Difference) is negative. ■

A.5 Proof of Corollary 4

Proof. To capture the marginal customization effect of the item $C+1$ delivered through the target item $C+2$, consider first a situation that the mechanism uses item 1 to item $C+1$ in making a prediction about the target item. We can derive the value to the target user using Proposition 1. Let R' be the set of outcomes that the first C items can possibly generate. The value to users is given by

$$\sum_{r' \in R'} \max\{v_1 p_{(r',1,1)} + v_0 p_{(r',1,0)}, 0\} + \max\{v_1 p_{(r',0,1)} + v_0 p_{(r',0,0)}, 0\}.$$

On the other hand, if the mechanism omits item $C+1$ in making a prediction about item $C+2$, the value delivered to the target user is

$$\sum_{r' \in R'} \max\{v_1 p_{(r',1,1)} + v_0 p_{(r',1,0)} + v_1 p_{(r',0,1)} + v_0 p_{r'(0,0)}, 0\}.$$

Thus, the marginal customization effect is the difference between the two values above.

To verify the second statement in the proposition, suppose $C+1$ and $C+2$ are correlated. That is, for some $r' \in R'$, we have $(v_1 p_{(r',1,1)} + v_0 p_{(r',1,0)} \ge 0$ and $v_1 p_{(r',0,1)} + v_0 p_{(r',0,0)} \le 0)$ or $(v_1 p_{r'(1,1)} + v_0 p_{(r',1,0) \le 0}$ and $v_1 p_{(r',0,1)} + v_0 p_{(r',0,0)} \ge 0)$, with at least one strict inequality in at least one case. Suppose we have $v_1 p_{(r',1,1)} + v_0 p_{(r',1,0)} > 0$ and $v_1 p_{(r',0,1)} + v_0 p_{(r',0,0)} \le 0$ for some r'. Then, by $v_1 p_{r'(0,1)} + v_0 p_{(r',0,0)} \le 0$, we have

$$\max\{v_1 p_{(r',1,1)} + v_0 p_{(r',1,0)}, 0\} > \max\{v_1 p_{(r',1,1)} + v_0 p_{(r',1,0)} + v_1 p_{(r',0,1)} + v_0 p_{(r',0,0)}, 0\}.$$

The same logic applies to other cases. Now, suppose we have

$$\begin{aligned}&\sum_{r'\in R'}\max\{v_1p_{(r',1,1)}+v_0p_{(r',1,0)},0\}+\max\{v_1p_{(r',0,1)}+v_0p_{(r',0,0)},0\}\\&>\sum_{r'\in R'}\max\{v_1p_{(r',1,1)}+v_0p_{(r',1,0)}+v_1p_{(r',0,1)}+v_0p_{(r',0,0)},0\}.\end{aligned}$$

Since $\max\{v_1p_{(r',1,1)}+v_0p_{(r',1,0)},0\}+\max\{v_1p_{(r',0,1)}+v_0p_{(r',0,0)},0\}\geq\max\{v_1p_{(r',1,1)}+v_1p_{(r',0,1)}+v_0p_{(r',1,0)}+v_0p_{(r',0,0)},0\}$, for every $r'\in R'$, we should have at least one $r''\in R'$ such that we have

$$\begin{aligned}&\max\{v_1p_{(r'',1,1)}+v_0p_{(r'',1,0)},0\}+\max\{v_1p_{(r'',0,1)}+v_0p_{(r'',0,0)},0\}\\&>\max\{v_1p_{(r'',1,1)}+v_0p_{(r'',1,0)}+v_1p_{(r'',0,1)}+v_0p_{(r'',0,0)},0\}.\end{aligned}$$

For such r'', if either $v_1p_{(r'',1,1)}+v_0p_{(r'',1,0)}\geq 0$ and $v_1p_{(r'',0,1)}+v_0p_{(r'',0,0)}\geq 0$ or $v_1p_{(r'',1,1)}+v_0p_{(r'',1,0)}\leq 0$ and $v_1p_{(r'',0,1)}+v_0p_{(r'',0,0)}\leq 0$, we should have

$$\begin{aligned}&\max\{v_1p_{(r'',1,1)}+v_0p_{(r'',1,0)},0\}+\max\{v_1p_{(r'',0,1)}+v_0p_{(r'',0,0)},0\}\\&=\max\{v_1p_{(r'',1,1)}+v_1p_{(r'',0,1)}+v_0p_{(r'',1,0)}+v_0p_{(r'',0,0)},0\}.\end{aligned}$$

Hence, it should be the case that we have either $(v_1p_{(r'',1,1)}+v_0p_{(r'',1,0)}\geq 0$ and $v_1p_{(r'',0,1)}+v_0p_{(r'',0,0)}\leq 0)$ or $(v_1p_{(r'',1,1)}+v_0p_{(r'',1,0)}\leq 0$ and $v_1p_{(r'',0,1)}+v_0p_{(r'',0,0)}\geq 0)$ with at least one inequality holding with strict inequality. That is, item $C+1$ and $C+2$ are correlated. ■

A-4 Proof of Proposition 4

Proof. We compare user utilities before and after item $C+1$ is added to the system. As before, let r' denote the ratings over C conditioning items. Using Lemma 1, the history r' user's utility before item $C+1$ is added is

$$\mathbf{1}\left\{\frac{p_{(r',1,1)}+p_{(r',0,1)}}{p_{(r',1)}+p_{(r',0)}}\geq\frac{-w_0}{w_1-w_0}\right\}\left(v_1\frac{p_{(r',1,1)}}{p_{r'}}+v_0\frac{p_{(r',1,0)}}{p_{r'}}+v_1\frac{p_{(r',0,1)}}{p_{r'}}+v_0\frac{p_{(r',0,0)}}{p_{r'}}\right).$$

On the other hand, the user utility after the addition of the item is

$$\begin{aligned}&\frac{p_{(r',1)}}{p_{r'}}\mathbf{1}\left\{\frac{p_{(r',1,1)}}{p_{(r',1)}}\geq\frac{-w_0}{w_1-w_0}\right\}\left(v_1\frac{p_{(r',1,1)}}{p_{(r',1)}}+v_0\frac{p_{(r',1,0)}}{p_{(r',1)}}\right)\\&+\frac{p_{(r',0)}}{p_{r'}}\mathbf{1}\left\{\frac{p_{(r',0,1)}}{p_{(r',0)}}\geq\frac{-w_0}{w_1-w_0}\right\}\left(v_1\frac{p_{(r',0,1)}}{p_{(r',0)}}+v_0\frac{p_{(r',0,0)}}{p_{(r',0)}}\right).\end{aligned}$$

We will only deal with the case when $p_{(r',1,1)}\geq p_{(r',0,1)}$ as the exact same logic can be applied to the other case. In this case, we have

$$\frac{p_{(r',0,1)}}{p_{(r',0)}}\leq\frac{p_{(r',1,1)}+p_{(r',0,1)}}{p_{(r',1)}+p_{(r',0)}}\leq\frac{p_{(r',1,1)}}{p_{(r',1)}}.$$

Thus, if the item is recommended to the user in the system without item $C+1$, it will also be recommended to the history $(r',1)$ user in the system with item $C+1$. Conversely, if the item is not recommended to the user in the system without item $C+1$, it will not be recommended to the user with history $(r',0)$ in the system with item $C+1$. Therefore, there are two cases when the extra degree in customization strictly hurts the user. The first is

$$\tau^u < \frac{p_{(r',0,1)}}{p_{(r',0)}} < \tau^p \leq \frac{p_{(r',1,1)} + p_{(r',0,1)}}{p_{(r',1)} + p_{(r',0)}}.$$

That is, the item is recommended to the target user in the system without item $C+1$, but in the system with item $C+1$ it is only recommended to the users whose history is $(r',1)$ even though it is expected to generate positive utility to history $(r',0)$ users.

On the other hand, there is also a case that the item is recommended to history $(r',1)$ users under the system with item $C+1$ even though it is not expected to generate positive utility to the users and the system without item $C+1$ does not recommend the item to users. This case arises when

$$\tau^u > \frac{p_{(r',1,1)}}{p_{(r',1)}} \geq \tau^p > \frac{p_{(r',1,1)} + p_{(r',0,1)}}{p_{(r',1)} + p_{(r',0)}}.$$

These two conditions coincide with the conditions presented in the proposition. ■

A.6 Proof of Proposition 5 and Corollary 5

Proof. For notational simplicity, let $P[y_{(1,1)} \geq N_1 - y_{(1,1)}] = T(m, 2m-1, s)$. That is,

$$T(m, 2m-1, s) = \sum_{k=m}^{2m-1} \binom{2m-1}{k} s^k (1-s)^{2m-1-k}.$$

We first need to show that $v(m|s)$ increases in m. Let $X_n \sim Binomial(n, s)$. Note that the associated cumulative distribution function of the binomial random variable is

$$P[X_n \leq k] = \sum_{i=0}^{k} \binom{n}{i} s^i (1-s)^{n-i} = 1 - \sum_{i=k+1}^{n} \binom{n}{i} s^i (1-s)^{n-i} = 1 - T(k+1, n, s).$$

Using this, we have

$$\begin{aligned}
T(m+1,2m+1,s) =& 1-P[X_{2m+1}\le m]\\
=&1-P[X_{2m+1}\le m|X_{2m-1}\le m-2]P[X_{2m-1}\le m-2]\\
&-P[X_{2m+1}\le m|X_{2m-1}=m-1]P[X_{2m-1}=m-1]\\
&-P[X_{2m+1}\le m|X_{2m-1}=m]P[X_{2m-1}=m]\\
=&1-P[X_{2m-1}\le m-2]-(1-s^2)P[X_{2m-1}=m-1]-(1-s)^2P[X_{2m-1}=m]\\
=&T(m-1,2m-1,s)-(1-s^2)\binom{2m-1}{m-1}s^{m-1}(1-s)^m\\
&-(1-s)^2\binom{2m-1}{m}s^m(1-s)^{m-1}.
\end{aligned}$$

Here, note that $\binom{2m-1}{m-1}=\binom{2m-1}{m}$ and $T(m-1,2m-1,s)=T(m,2m-1,s)+\binom{2m-1}{m-1}s^{m-1}(1-s)^m$. The last expression can be simplified to

$$T(m+1,2m+1,s)=T(m,2m-1,s)-(1-2s)\binom{2m-1}{m}s^m(1-s)^m.$$

Thus, we have

$$\begin{aligned}
v(m+1|s)-v(m|s)=&-(1-2s)(T(m+1,2m+1,s)-T(m,2m-1,s))\\
=&(1-2s)^2\binom{2m-1}{m}s^m(1-s)^m\ge 0.
\end{aligned}$$

That is, for any realization of s, the user always expects weakly higher utility when there are more accumulated data points.

To show the diminishing marginal return property, let $\Delta v(m|s)$ be the marginal externality that the $(m+1)^{th}$ user contributes to the subsequent user. That is,

$$\Delta v(m|s)=v(m+1|s)-v(m|s).$$

Using the derivation in the above proposition, it has a closed form representation of

$$\Delta v(m|s)=(1-2s)^2\binom{2m-1}{m}(s(1-s))^m.$$

Consider the ratio between the following two increments:

$$\begin{aligned}\frac{\Delta v(m+1|s)}{\Delta v(m|s)} =& \frac{v(m+2|s)-v(m+1|s)}{v(m+1|s)-v(m|s)}\\ =& \frac{\binom{2m+1}{m+1}s^{m+1}(1-s)^{m+1}}{\binom{2m-1}{m}s^{m}(1-s)^{m}}\\ =& \frac{2(2m+1)}{m+1}s(1-s) < 4s(1-s) \leq 1.\end{aligned}$$

That is, the increment diminishes.

The proof of Corollary 5 immediately follows since $v(m+1|s)$ is strictly higher than $v(m|s)$ for all s except $s=\frac{1}{2}$. ■

A.7 Proof of Proposition 6

Proof. To begin, let x be the history of the target user. Note that the Dirichlet distribution is stable with respect to any aggregation.[25] Given this property, we have

$$\left(\sum_{r\in R_i^+(x)} q_r^1, \sum_{r\in R_i^-(x)} q_r^1, \sum_{r\notin R_i^+(x)\cup R_i^-(x)} q_r^1\right) \sim Dir\left(\sum_{r\in R_i^+(x)} \alpha_r^1, \sum_{r\in R_i^-(x)} \alpha_r^1, \sum_{r\notin R_i^+(x)\cup R_i^-(x)} \alpha_r^1\right).$$

For the sake of notation, we denote $\sum_{r\in R_i^+(x)} q_r^1$ and $\sum_{r\in R_i^-(x)} q_r^1$ by b_1 and b_2 respectively. Similarly, let β_1, β_2 and β_3 denote $\sum_{r\in R_i^+(x)} \alpha_r^1$, $\sum_{r\in R_i^-(x)} \alpha_r^1$ and $\sum_{r\notin R_i^+(x)\cup R_i^-(x)} \alpha_r^1$ respectively. The conditional probability of positive experience with i can be represented as $\frac{b_1}{b_1+b_2}$ when $(b_1, b_2, 1-b_1-b_2) \sim Dir(\beta_1, \beta_2, \beta_3)$.

For $k \in \{1,2,3\}$, define an independent set of random variables γ_k each of which follows a gamma distribution with a shape parameter β_k and a rate parameter θ for some $\theta > 0$; i.e., $\gamma_k \sim Gamma(\beta_k, \theta)$.

It is well known that $\left(\frac{\gamma_1}{\gamma_1+\gamma_2+\gamma_3}, \frac{\gamma_2}{\gamma_1+\gamma_2+\gamma_3}, \frac{\gamma_3}{\gamma_1+\gamma_2+\gamma_3}\right) \sim Dir(\beta_1, \beta_2, \beta_3)$. Thus, we have the following equality in distribution:

$$(b_1, b_2) \overset{d}{=} \left(\frac{\gamma_1}{\gamma_1+\gamma_2+\gamma_3}, \frac{\gamma_2}{\gamma_1+\gamma_2+\gamma_3}\right).$$

Here, note that if $X \overset{d}{=} Y$, then $h(X) \overset{d}{=} h(Y)$ for any deterministic function h. Letting $h_1(x_1, x_2) = \frac{x_1}{x_1+x_2}$ and $h_2(x_1, x_2) = \frac{x_2}{x_1+x_2}$, we have

$$\frac{b_1}{b_1+b_2} \overset{d}{=} \frac{\gamma_1}{\gamma_1+\gamma_2} \text{ and } \frac{b_2}{b_1+b_2} \overset{d}{=} \frac{\gamma_2}{\gamma_1+\gamma_2}.$$

Furthermore, using the relationship between the gamma distribution and the Dirichlet distribution

[25] If $(p_1, \cdots, p_n)$ is Dirichlet with parameter $\alpha_1, \cdots, \alpha_n$, then a collection of sums of elements also follows Dirichlet. For example, when $n=4$, $(p_1+p_4, p_2+p_3) \sim Dir(\alpha_1+\alpha_4, \alpha_2+\alpha_3)$.

once again, we conclude that

$$\left(\frac{b_1}{b_1+b_2}, \frac{b_2}{b_1+b_2}\right) \sim Beta(\beta_1, \beta_2).$$

That is, the predictive experience of the target user with i takes the form of the following random variable: $\hat{\mathbf{z}}_{\mathbf{i}}^{\mathbf{N}} \sim Beta(\beta_1, \beta_2)$. The first statement in the proposition can be obtained by simply taking an expectation of $\hat{\mathbf{z}}_{\mathbf{i}}^{\mathbf{N}}$:

$$\hat{z}_i^N = \frac{\beta_1}{\beta_1+\beta_2} = \frac{\sum_{r\in R_i^+(x)} \alpha_r^1}{\sum_{r\in R_i^+(x)\cup R_i^-(x)} \alpha_r^1}.$$

Lastly, recall that we already have shown that

$$\frac{\sum_{r\in R_i^+(x)} q_r^1}{\sum_{r\in R_i^+(x)\cup R_i^-(x)} q_r^1} \sim Beta\left(\sum_{r\in R_i^+(x)} \alpha_r^1, \sum_{r\in R_i^-(x)} \alpha_r^1\right),$$

for given data X and user rating x. Consider now that the data is collected from K users. Let K_1 be the number of ratings whose associated outcomes are in $R_i^+(x)$, K_2 be the number of ratings whose associated outcomes are in $R_i^-(x)$, and M be the total number of trials whose outcome is consistent with x, i.e., $K = K_1 + K_2$. We have

$$\lim_{K\to\infty} \mathbb{E}[\hat{\mathbf{z}}_{\mathbf{i}}^{\mathbf{K}}|X_K, x] \quad = \quad \lim_{K\to\infty} \frac{K_1 + \sum_{r\in R_i^+(x)} \alpha_r^0}{K_1 + K_2 + \sum_{r\in R_i^+(x)\cup R_i^-(x)} \alpha_r^0} \overset{a.s.}{=} \lim_{K\to\infty} \frac{K_1}{K}.$$

Here, X_K denotes the data from K users. By the law of large numbers, the last term is the same as z_i. On the other hand,

$$\begin{aligned}\lim_{K\to\infty} \mathbb{V}ar[\hat{\mathbf{z}}_{\mathbf{i}}^{\mathbf{K}}|X_K, x] \quad &= \quad \lim_{K\to\infty} \frac{(K_1 + \sum_{r\in R_i(x)} \alpha_r^0)(K_2 + \sum_{r\in R_i^+(x)} \alpha_r^0)}{(K + \sum_{r\in R_i^+(x)\cup R_i^-(x)} \alpha_r^0)^2 (K + 1 + \sum_{r\in R_i^+(x)\cup R_i^-(x)} \alpha_r^0)} \\ &\leq \quad \lim_{K\to\infty} \frac{1}{K + 1 + \sum_{r\in R_i^+(x)\cup R_i^-(x)} \alpha_r^0} = 0,\end{aligned}$$

which gives the convergence in mean-square to z_i. ■

A-7 Proof of Proposition 7

Proof. The result immediately follows from the two expressions of the predictions:

$$\hat{x}_i^G = \frac{\sum_{r\in R_i^+(\emptyset)} y_r + \sum_{r\in R_i^+(\emptyset)} \alpha_r^0}{n - 1 + \sum_{r\in R} \alpha_r^0} \text{ and } \hat{x}_i^{AR} = \frac{\sum_{r\in R_i^+(\emptyset)} y_r}{n-1}.$$

As $n \to \infty$, it is obvious that we have $\hat{x}_i^G = \hat{x}_i^{AR}$, $\forall i$. ■