

Learning Personalized Risk Preferences for Recommendation

Yingqiang Ge[†], Shuyuan Xu[†], Shuchang Liu[†], Zuohui Fu[†], Fei Sun[‡], Yongfeng Zhang[†]

[†]Rutgers, The State University of New Jersey [‡]Alibaba Group

{yingqiang.ge,shuyuan.xu,sl1471,zuohui.fu}@rutgers.edu, ofey.sunfei@gmail.com, yongfeng.zhang@rutgers.edu

ABSTRACT

The rapid growth of e-commerce has made people accustomed to shopping online. Before making purchases on e-commerce websites, most consumers tend to rely on rating scores and review information to make purchase decisions. With this information, they can infer the quality of products to reduce the risk of purchase. Specifically, items with high rating scores and good reviews tend to be less risky, while items with low rating scores and bad reviews might be risky to purchase. On the other hand, the purchase behaviors will also be influenced by consumers' tolerance of risks, known as the risk attitudes. Economists have studied risk attitudes for decades. These studies reveal that people are not always rational enough when making decisions, and their risk attitudes may vary in different circumstances.

Most existing works over recommendation systems do not consider users' risk attitudes in modeling, which may lead to inappropriate recommendations to users. For example, suggesting a risky item to a risk-averse person or a conservative item to a risk-seeking person may result in the reduction of user experience. In this paper, we propose a novel risk-aware recommendation framework that integrates machine learning and behavioral economics to uncover the risk mechanism behind users' purchasing behaviors. Concretely, we first develop statistical methods to estimate the risk distribution of each item and then draw the Nobel-award winning Prospect Theory into our model to learn how users choose from probabilistic alternatives that involve risks, where the probabilities of the outcomes are uncertain. Experiments on several e-commerce datasets demonstrate that by taking user risk preferences into consideration, our approach can achieve better performance than many classical recommendation approaches, and further analyses also verify the advantages of risk-aware recommendation beyond accuracy.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; • **Applied computing** → **Economics**; • **Computing methodologies** → **Intelligent agents**.

KEYWORDS

Recommender Systems; Personalization; Risk Attitudes; Prospect Theory; Computational Economics

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '20, July 25–30, 2020, Virtual Event, China

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8016-4/20/07...\$15.00

<https://doi.org/10.1145/3397271.3401056>

ACM Reference Format:

Yingqiang Ge[†], Shuyuan Xu[†], Shuchang Liu[†], Zuohui Fu[†], Fei Sun[‡], Yongfeng Zhang[†]. 2020. Learning Personalized Risk Preferences for Recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)*, July 25–30, 2020, Virtual Event, China. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3397271.3401056>

1 INTRODUCTION

Designing personalized recommender systems is able to help users find relevant items efficiently in the context of web information overload. A well-informed recommender system is capable of not only saving consumers' exploration time but also benefiting the revenue of various online economic platforms. Traditional recommendation algorithms mostly focus on optimizing rating- or ranking-oriented metrics. However, previous recommendation algorithms seldom consider users' economic incentives when modeling user behaviors and generating recommendations.

In fact, consumers' economic incentives play an important role when making decisions in online economic systems such as e-commerce [14, 51]. A common observation in practical recommendation systems is that customers may use rating scores and review information to support their buying decisions [26, 28, 33]. According to this information, buyers can infer the quality of products to avoid wasting time and reduce the risk of purchase. However, the judgments of risky items are different from person to person. Actually, users with different risk attitudes could make different decisions, even when facing the same situation. Concretely, risk-averse consumers would be more likely to buy products that are "safe choices", i.e., that have many good reviews, while risk-appetite consumers may more likely try new products even though they yet to have sufficiently many good reviews. Meanwhile, based on real-user psychological experiments, behavioral economists [21] have revealed that most consumers will overweight small probability gains and underweight moderate or high probability losses. For instance, when facing 1% probability to win \$500 or 100% probability to win \$5, most people would take the risk and choose the first choice, although the expected gains are the same; however, when facing 1% probability to lose \$500 or 100% probability to lose \$5, most people would mitigate the risk and choose the second one, although the expected losses are the same. Thus, understanding consumer decisions under risk can assist researchers to better build personalized recommendation algorithms according to users' risk preferences, and an informed recommendation system with risk-awareness can provide appropriate recommendations that make consumers comfortable. Fortunately, integrating machine learning and the established economic principles can help to model the risk attitude of the user decision process based on large-scale user transaction logs.

In this paper, we propose a novel risk-aware recommendation framework, which takes users' risk attitudes into consideration according to several behavioral economic principles. In particular, we simulate the risk distribution of each item based on its rating distribution by conventional statistical methods. Combining the Nobel-prize winning theory in economics — the Prospect Theory [21] — with machine learning algorithms, our model learns to predict the consumer decisions with risk-awareness.

The key contributions of the paper can be summarized as follows:

- We take consumers' risk attitudes and weighted event probability into consideration for economic recommendation, which better simulates real-world online economic environment, where users have to make decisions under potential risks of dissatisfaction.
- Our model can be considered as a personalized version of the Prospect Theory, i.e., unlike the original Prospect Theory, which assumes a uniform risk function with the same parameters for all subjects, our method endeavors to figure out the personalized risk preference for each user, which adapts the economic theory under different risk attitudes for different users.
- Experimental results on real-world e-commerce verify that our approach can not only achieve better recommendation performance than both classical and economic recommendation baselines but also successfully adapt the Prospect Theory to learn each user's personalized risk attitudes.

The following contents of the paper will be organized as follows: we first review some related work in section 2, and then introduce the key economic preliminaries in section 3 to prepare readers with the underlying economic backgrounds used in this work. The proposed model and recommendation strategies are introduced in section 4, followed by experimental results in section 5. Finally, we conclude the work with possible future research directions in section 6.

2 RELATED WORK

In this section, we will briefly introduce some background knowledge to help the readers get a better understanding of the areas that are related to our work.

2.1 Collaborative Filtering

Collaborative Filtering (CF) has been one of the most dominant approaches to recommender systems. Early methods of CF consider the user-item rating matrix and conduct rating prediction task with user-based [23, 38] or item-based [29, 40] collaborative filtering methods. In these methods, the user and item rating vector are considered as the representation vector for each user and item.

With the development of dimension reduction methods, latent factor models, such as singular value decomposition (SVD) [25], non-negative matrix factorization [27], and probabilistic matrix factorization [32], are later widely adopted in recommender systems. In these aforementioned matrix factorization approaches, each user and item is learned as a latent factor representation to calculate the matching score of the user-item pair.

Deep models have recently been further extended to collaborative filtering methods for the recommendation tasks. The relevant

methods can be roughly divided into two subcategories: similarity learning methods and representational learning methods. The similarity learning approaches adopt simple user/item representations (such as one-hot) and learn a complex prediction network as the similarity function to calculate user-item matching scores [19]. Meanwhile, the representation learning approaches learn rich user/item representations and adopt a simple similarity function (e.g., inner product) for matching score calculation [49].

Another important research direction is learning to rank (LTR) for recommendation, which learns the relative ordering of items rather than the absolute preference scores. Probably, the most representative method in this direction is Bayesian Personalized Ranking (BPR) [37], which is a pair-wise learning-to-rank method for recommendation. It is also further generalized to take other information sources such as images [18] and text [13, 49] into consideration.

2.2 Economic Recommendation

For a long time, the focus of recommendation system research has been on the above-mentioned rating- or ranking-related tasks, such as rating prediction and top-N recommendation, however, the related approaches rarely consider the economic value that the recommendation lists could bring to the users or the system, although it is one of the most important goals for real-world recommendation systems. Fortunately, we can find some recent research on economic recommendation has begun to take care of the economic value of personalized recommendation. For example, [47] studied the user's sense of value for items in terms of utility in recommender systems, and [52] conducted large-scale experiments with real-world users to verify consumers' sense of utility for personalized promotions. Further, [51] bridged economic principles and machine learning methods to maximize the social surplus for recommendation, and [53] proposed to learn the substitute and complementary relations between products so as to maximize consumer utility. Researchers have also considered using economic models for user behavior analysis in search systems [3]. Although current economic recommendation approaches may improve the economic value, their underlying motivation is to maximize a total utility function for each user to generate recommendations, which is quite different from ours. Established behavioral economic principles show that consumers usually make decisions under risk conditions, and the risk attitude/preference of different users could be different, which may influence their decisions in economic systems [21]. This motivates us to estimate the risk preference of users for economics-driven risk-aware recommendation.

There are only a few works studying the effect of user risk attitudes in recommendation systems. [36] is motivated by the intuition that "the pain of loss is more powerful than the pleasure from a similar gain", and designs an algorithm suggesting fewer negative items to users. [34] shows that most of the existing recommendation algorithms are either risk-neutral or risk-averse, and it designs a risk-seeking strategy to recommend items. However, the users' risk attitudes are personalized, which can hardly be modeled by a single global assumption. As a result, we decide user risk attitudes according to different scenarios (gains and losses) based on Prospect Theory, take advantage of machine learning to personalize their risk attitudes, and use prospect values to make top-K recommendations.

Another stream of risk-aware recommendation research belongs to Bouneffouf et al. [5, 6]. In this research, the risk is defined as the exploration-exploitation trade-off in contextual bandit algorithms, which is different from our perception of consumers' risk attitudes and the consumer decision-making process under uncertainty.

2.3 Decision-Making Under Uncertainty

The real-world is filled with plenty of uncertainties, and it involves various risk factors. People have to make a large number of decisions under uncertainty in their daily life, such as making purchases or investments. As a result, decision-making under risk is one of the most fundamental research subjective for economists. There have been two major types of theory on decision-making under risk, and the critical difference between them is how to understand the definition of uncertainty. The first type of theory holds the view that uncertainty is objective, which means that the uncertainty of an event can be represented as an objective probability distribution of all possible outcomes. One representative theory of this view is the von Neumann-Morgenstern Expected Utility Theory [46]. The other type of risk theory believes that the uncertainty of events is not objective, but it is based on people's subjective judgments of all possible outcomes. In particular, Savage [41] puts forward a theory of subjectivity based on personal probability and statistics, which constitutes the research line underlying Bayesian statistics.

Overall, there exists more than one economic theory on risk modeling, and among them, the Expected Utility Theory was generally adopted and practically applied due to its intuition and convenience. However, economists find that these theories still cannot correctly explain all the known consumer behaviors, and most importantly, researchers gradually realized the limitations of these theories, because of the emerging of several classical paradoxes out of the theories, which were raised by Bernoulli [4], Allais [2], and Ellsberg [11]. To solve these paradoxes, Daniel Kahneman and Amos Tversky renewed the study based on human behaviors and proposed the Prospect Theory [21] in the year of 1979 to model the user risk preferences, and later in 2002, Kahneman was awarded the Nobel Prize in Economics for the success of Prospect Theory in modeling consumer risks in economic behaviors.

3 PRELIMINARIES

In this section, we introduce some essential economic preliminaries about the above-mentioned risk attitudes, and prospect theory to help readers better understand our proposed model later.

3.1 Risk Attitudes

Users usually have to take the risk for their own choices when making decisions under uncertainty. In recommender systems, consumers hardly know for sure if they would be satisfied with a product or not, because a purchase that is believed to be a good match may still result in dissatisfaction due to various factors, such as the quality or delayed shipping, which can be considered as potential risk factors in e-commerce. As we have discussed before, different people may have different attitudes toward risks [1, 20, 45]. In general, consumer risk attitudes can be divided into three categories, which are introduced as follows:

Risk Aversion: Even though there exist many methods to describe an individual's risk aversion attitude, the essence of them is the same, i.e., comparing all of the possible returns under different choices that have the same expected value, a risk-aversion consumer would always prefer the certain returns.

Risk Appetite: The definition of risk appetite is counterpart against risk aversion, i.e., comparing all of the possible returns under different choices that have the same expected value, a risk-appetite consumer would always prefer the risky returns.

Risk Neutral: Risk neutral is a mindset where an individual is indifferent to risk when making a decision under uncertainty.

It would be easier to understand the difference between different risk preferences with an example. We still take gambling as an instance, but this time, we focus on people's risk attitudes instead of overweight small probability gains and underweight moderate or high probability losses. Suppose there are two choices in a gamble, one is to lose either \$100 or \$0 with 50% probability each, and the other choice is to lose \$50 with 100% probability. Though the expected utility of both choices is the same (\$-50), different users may make different choices based on their risk attitude. A risk-aversion user would prefer to choose a sure \$50 loss, and a risk-appetite user would prefer to risk for a potential \$0 loss, while a risk-neutral user is indifferent to the two choices.

3.2 Prospect Theory

Prospect theory is a commonly accepted and widely applied economic principle in behavioral economics, which describes the consumer's decision-making process between probabilistic alternatives involving risks. It believes that people have different evaluations for gains and losses, and describes the decision-making process in two steps. Firstly, people will divide all potential outcomes into gains or losses by setting a reference point that they consider as indifferent. Consequently, those outcomes less than the reference point are considered as losses, and those that are greater are considered as gains. Secondly, people have different prospect value functions for each outcomes and its probabilistic estimation. With this, people will calculate prospect values based on all the possible outcomes and their respective probabilities, and finally, choose the best action according to prospect values instead of raw outcomes.

According to Kahneman [21, 44], the prospect value function is a product of the value function and the weighting function. While the value function measures the subjective value of the gain/loss, the weighting function can be seen as a subjective distortion of the original probabilities, which will be introduced later.

On one hand, through a large number of psychological experiments, Kahneman depicted some general consumer psychological principles and summarized that people usually are risk-aversion for gains and risk-appetite for losses, because they have stronger feelings for losses than for gains, i.e., the happiness out of gaining \$100 usually cannot compensate for the hurt of losing \$100. The hypothesis of the value function is based on the above observations and is shown in Fig. 1a. For gains, the function is concave to show risk-aversion; for loss, it is convex to show risk-appetite. Also, it is steeper for losses than gains indicating that losses outweigh gains.

On the other hand, according to the observations of many behavioral economic experiments [8, 48], people usually overweight

the low probabilities and underweight the high probabilities for both gains and losses. This trend, which is not reflected in the value function, is reflected by the individual differences in the weighting function. Specifically, there exists a probability distortion that people generally do not look at the value of probability uniformly between 0 and 1. Lower probability will be overweighted (that means a person is over concerned with the low probability outcome), while higher probability is underweighted (that means a person is not concerned enough with the high probability outcome). Based on these observations, the hypothesis of the shape of the weighting function is shown in Fig.1b.




Figure 1: The left figure shows a hypothetical value function, while the right figure is a hypothetical weighting function.

4 THE FRAMEWORK

In this section, we further discuss how to bridge the fundamental economic concepts with machine learning algorithms to establish a risk-aware recommendation framework.

4.1 Personalized Prospect Theory

In order to implement prospect theory in real-world recommendation scenario, we first define a few concepts and then specify their function in e-commerce. Let S be a finite set of states of nature; subsets of S are called events. In [44], it is assumed that exactly one state obtains, which is unknown to the decision-maker (e.g., one item only can be rated with 1 or 2 or ... or 5 stars on Amazon by a certain user). Let X be a set of consequences/outcomes and assume that X includes a neural outcome, denoted as 0, then we interpret all other elements as gains or losses through comparisons. An uncertain prospect f , which is a function from S to X , is represented as a sequence of pairs (x_i, s_i) , which yields x_i if s_i occurs, and (s_i) is a partition of S . When a prospect $f = (x_i, s_i)$ is given by a probability distribution $p(s_i) = p_i$, it can be viewed as a risky prospect (x_i, p_i) .

In an e-commerce scenario, the consumer satisfaction of purchasing an item is usually reflected by the corresponding rating scores [12], which can be seen as a potential state of the purchase. Therefore, we define the ratings, which are usually in the range of $\{1, 2, 3, 4, 5\}$, as states and get $S = \{r_1, r_2, r_3, r_4, r_5\}$, where $r_i = i$. For the potential outcome of one purchase, we assume it is related to both the price of a given item and the expectation of the consumer to that item. We use the reference point, denoted by \hat{r}_u , to represent the user u 's expectation to a given item. The difference between the final rating and the reference point will cause a gain

or a loss to the consumer. Then, for every possible state r_i , the outcomes (gains or losses) can be designed as:

$$x_i = \text{price} \cdot \tanh(r_i - \hat{r}_u),$$

where $r_i = i, i \in 1, \dots, 5$

The outcome x_i is positive for gains (i.e., the rating is greater than the reference point) and negative for losses (i.e., the rating is less than the reference point). Since the difference between r_i and \hat{r}_u won't be too large, we choose the hyperbolic tangent function $\tanh(\cdot)$ to normalize $(r_i - \hat{r}_u)$ into $[-1, 1]$ so that we get $x_i \in [-\text{price}, \text{price}]$. The probability distribution $p(r_i)$ under the e-commerce setting will be discussed in detail in Sec. 4.2; for now, we assume there exists certain $p(r_i) = p_i$, where $r_i \in S$.

4.1.1 The Value Function. To accord with the basic ideas of the Prospect Theory — a value function is concave for gains, convex for losses, and steeper for losses than for gains, we use the following function $v(x)$ as the value function, which comes from Kahneman's work in 1992 [44].

$$v(x_i) = \begin{cases} \lambda x_i^\alpha & \text{if } x_i \geq 0 \\ -(-x_i)^\beta & \text{if } x_i < 0 \end{cases} \quad (2)$$

where $0 < \alpha, \beta, \lambda < 1$. $\lambda \in (0, 1)$ is to make sure that the value function is steeper for losses than for gains. Since our model needs to learn the personalized risk attitudes of a user u under the circumstance that he/she is given a certain item v , we will use α_{uv} , β_{uv} , λ_{uv} in the following sections to represent their use in certain user-item pair. The above function reflects the required value function shape in prospect theory, as shown in Fig.1a.

4.1.2 The Weighting Function. The weighting function is used to represent the decision weight for all possible outcomes based on their corresponding probabilities, and we use the following function $\pi(p)$ as the weighting function, which is still consistent in the way of Kahneman's study [44].

$$\pi(p_i) = \begin{cases} \frac{p_i^\gamma}{(p_i^\gamma + (1 - p_i)^\gamma)^{1/\gamma}} & \text{for gain} \\ \frac{p_i^\delta}{(p_i^\delta + (1 - p_i)^\delta)^{1/\delta}} & \text{for loss} \end{cases} \quad (3)$$

where p_i is the probability of the gain (or loss), depending on the corresponding x_i ; $\pi(0) = 0$ and $\pi(1) = 1$; γ and δ are positive parameters ($0 < \gamma, \delta < 1$) used to control the shape of the weighting function, which will be learned in the model. It should be noted that the weight $\pi(p)$ is no longer a probability, and it only represents the importance of a gain/loss for the consumer. For the same purpose of personalization, we will use γ_{uv} and δ_{uv} to replace them in the following sections.

4.1.3 Prospect Value. According to Kahneman and Tversky's work in 1979 [21] and 1992 [44], we use a simple mathematical formula to calculate the prospect value defined as:

$$\mathcal{V} = \sum_{i=1}^l v(x_i) \pi(p_i) \quad (4)$$

where \mathcal{V} represents the prospect value for a certain choice, which consists of $v(x)$ as the value function and $\pi(p)$ as the weighting

function; $v(x_1), \dots, v(x_l)$ are the values of gains or losses for all potential states, and p_1, \dots, p_l are the corresponding probabilities.

4.1.4 Calculation Paradigm. In order to clarify the calculation of the proposed prospect value and help readers better understand our framework, we give an example to show the process. Suppose we have a user-item pair, and assume we know the user's expected score of that item is 2.5 (which indicates $\hat{r}_u = 2.5$). Suppose we also know the rating distribution of that item, e.g. {1 star : p_1 , 2 stars : p_2 , ..., 5 stars : p_5 }. Then, we have Table 1 showing all possible states with their corresponding outcomes and weights. We can see that since 1 star and 2 stars are below 2.5, which make x_1 and x_2 below to zero, they are losses for the user and need to use $-(-x_i)^\beta$ as the value function and $\frac{p_i^\delta}{(p_i^\delta + (1-p_i)^\delta)^{1/\delta}}$ as the weighting function; on the contrary, 3 to 5 stars are gains for the user, thus, we use the remaining parts. Based on Eq. 4, the prospect value of this user-item pair is calculated as follows:

$$\mathcal{V} = v_-(x_1)\pi_-(p_1) + v_-(x_2)\pi_-(p_2) + v_+(x_3)\pi_+(p_3) + v_+(x_4)\pi_+(p_4) + v_+(x_5)\pi_+(p_5) \quad (5)$$

where $v_-(x_i) = -(-x_i)^\beta$ and $\pi_-(p_i) = \frac{p_i^\delta}{(p_i^\delta + (1-p_i)^\delta)^{1/\delta}}$ represent the negative part of the value function and the weighting function, respectively (when $x_i < 0$), while $v_+(x_i) = \lambda x_i^\alpha$ and $\pi_+(p_i) = \frac{p_i^\gamma}{(p_i^\gamma + (1-p_i)^\gamma)^{1/\gamma}}$ represent the positive part when $x_i \geq 0$.

Table 1: Examples showing how to calculate the prospect value for a certain user-item pair.

States	x_i	Value	Weight
1 star	$price \cdot \tanh(1 - 2.5) < 0$	$v_-(x_1)$	$\pi_-(p_1)$
2 stars	$price \cdot \tanh(2 - 2.5) < 0$	$v_-(x_2)$	$\pi_-(p_2)$
3 stars	$price \cdot \tanh(3 - 2.5) > 0$	$v_+(x_3)$	$\pi_+(p_3)$
4 stars	$price \cdot \tanh(4 - 2.5) > 0$	$v_+(x_4)$	$\pi_+(p_4)$
5 stars	$price \cdot \tanh(5 - 2.5) > 0$	$v_+(x_5)$	$\pi_+(p_5)$

4.2 Estimation of Rating Distribution

Our objective is to solve the classical problem of decision-making under uncertainty, and therefore, in this section, we define the uncertainty in e-commerce. Consumers often use other information about the products, such as rating scores and reviews, to estimate their probability of satisfaction and support their purchase decision. In this work, we assume buyers use rating scores to estimate the product quality, as the logged rating score is an important indicator that reflects the satisfaction of the consumers who had already purchased the item; in our future work, we plan to combine rating and reviews for better risk estimation.

Intuitively, we can directly use an item's rating distribution as the risk distribution in our model. However, there exist some problems if we do so. For example, some items' rating distribution may contain 0 probabilities, which means they get no feedback on certain rating scores. This could happen because the dataset can still be seen as a sampling from the ground truth, and it does not contain all possible outcomes. Thus, except for filtering out items without enough actions, we also estimate the rating distributions

for those items contain 0 probabilities to get more plausible risk distributions for them; and for those who do not, we use their raw rating distributions directly. Inspired by [50], we adopt Weibull distribution to reconstruct the continuous rating distribution from limited discrete observations, since Weibull distribution is able to fit very flexible curve shapes. The probability density function (PDF) of the numerical rating z is modeled as follows:

$$f(z) = \mu \eta z^{\mu-1} \exp(-\eta z^\mu) \quad (6)$$

where η is the scale parameter and μ is the shape parameter. Since the rating is a discrete value representing an approximation of user satisfaction, we use the interval integration of the estimated Weibull distribution as the probabilities of the potential outcomes. Specifically, $p_i = \Pr(z = i) = \int_{i-0.5}^{i+0.5} f(z)dz$ is the probability that the rating score of the item is i ($i = 1, 2, 3, 4$ or 5 , for $i = 1$, integral lower bound is $-\infty$ and for $i = 5$, upper bound is $+\infty$).

4.3 Risk-aware Recommendation (RARE)

In this work, we propose an optimization framework based on discrete choice model for our risk-aware recommendation learning framework – **RARE**. Online shopping requires consumers to choose their desired items among alternatives, which can be considered as a discrete choice problem. This discrete choice problem describes a situation when a consumer chooses an option between two or even more discrete alternatives. More formally, consumer u chooses item v over a set of some other alternative products $\Omega_u(v)$. We define the total choice set as $\Pi_u = \{v, \Omega_u(v)\}$ and its k -th element is Π_u^k (i.e., $\Pi_u^1 = v$). The probability that consumer u chooses alternative v is denoted as P_{uv} . We can simply view v and $\Omega_u(v)$ as positive and negative training records, respectively.

Researchers in economics have utilized Random Utility Models (RUM) to deal with discrete choice problem [9]. Different from traditional RUMs, we adopt the idea of choosing the alternative item that provides the highest utility in choosing the alternative item with the highest prospect value. In this way, we have:

$$\hat{\mathcal{V}}_u(\Pi_u^k) = \mathcal{V}_u(\Pi_u^k) + \epsilon_k \quad (7)$$

where $\mathcal{V}_u(\Pi_u^k)$ represents the true prospect value of the k -th product in item set Π_u , and $\hat{\mathcal{V}}_u(\Pi_u^k)$ represents the observed prospect value of that product and ϵ_k is a random variable capturing the impact of all unknown factors. Thus, the probability of a customer choosing Π_u^1 (which is positive item v) over other alternatives is:

$$P_{uv}(\hat{\mathcal{V}}_u(\Pi_u^1) > \hat{\mathcal{V}}_u(\Pi_u^k)) = P_{uv}(\epsilon_k - \epsilon_1 < \mathcal{V}_u(\Pi_u^1) - \mathcal{V}_u(\Pi_u^k)) \quad (8)$$

where $k = 2, \dots, |\Pi_u|$. If ϵ_1 and ϵ_k follow an i.i.d. extreme value distribution, it can be shown that the probability of choosing Π_u^1 is the following multinomial logistic (MNL) model [16],

$$P(y_{uv} = 1) = P_{uv}(\hat{\mathcal{V}}_u(\Pi_u^1) > \hat{\mathcal{V}}_u(\Pi_u^k)) = \frac{\exp(\mathcal{V}_u(\Pi_u^1))}{\sum_{k=1}^{|\Pi_u|} \exp(\mathcal{V}_u(\Pi_u^k))} \quad (9)$$

where y_{uv} is an indication function that is

$$y_{uv} = \begin{cases} 1 & \mathcal{V}_u(\Pi_u^1) > \mathcal{V}_u(\Pi_u^k) \quad \forall k \neq 1 \\ 0 & \text{Otherwise} \end{cases} \quad (10)$$

$y_{uv} = 1$ indicates user u chooses item v among all the other alternatives $\Omega_u(v)$. Therefore, $P(y_{uv} = 1)$ tends to maximize the probability of choosing v among Π_u as $\Pi_u^1 = v$, further more, it tends

to maximize the probability that the prospect value of v is greater than those of its alternatives, namely, $\mathcal{V}_u(v) > \mathcal{V}_u(\Pi_u^k), \forall k \neq 1$. In the training process, we set the number of alternatives as two and pick them randomly from items that the user has no interaction before; thus, we get $\Pi_u = \{v, \text{negative_sample}_1, \text{negative_sample}_2\}$ for each user-item pair in our train set.

Given the observed transactions and MNL, the model parameters can be learned by maximizing the probability of choosing the purchased items, which is

$$\max_{(u,v) \in \mathcal{R}} \sum \log(P(y_{uv} = 1)) - \lambda \|\Phi\|^2, \quad (11)$$

where \mathcal{R} is the training dataset; Φ is the parameter set to be learned in the corresponding loss function, which will be crystallized in the following section.

4.4 Model Specification

Economic principles mainly focus on finding out the general rules for the majority of people, however, with the help of machine learning and massive data, we now gain the ability to design a model to fit the personalized preference of each individual consumer.

To obtain personalized measurements, each user-item pair comes with a set of parameters, i.e. $\alpha_{uv}, \beta_{uv}, \lambda_{uv}, \gamma_{uv}, \delta_{uv}, \hat{r}_u$. In particular, each parameter (except for the reference point \hat{r}_u) can be decomposed into the global bias, user bias, item bias, and K -dimensional latent factors of users and items, which is the same idea implemented in SVD based CF [39]. For example, the parameter α in the value function Eq.(2) can be rewritten as:

$$\alpha_{uv} = g^\alpha + b_u^\alpha + l_v^\alpha + \mathbf{p}_u^{\alpha T} \mathbf{q}_v^\alpha, \quad (12)$$

where g^α is the global bias of α , b_u^α is the user bias of α , l_v^α is the item bias of α , \mathbf{p}_u^α and \mathbf{q}_v^α are k -dimensional latent factors of user u and item v . The above reparametrization is also used for $\beta_{uv}, \lambda_{uv}, \gamma_{uv}$ and δ_{uv} .

Besides, since each consumer has his/her own reference point, the reference point \hat{r}_u of each consumer u also needs to be learned. In this way, the optimization problem in Eq.(11) can be specialized as follow:

$$\max_{(u,v) \in \mathcal{R}} \sum \log \left(\frac{\exp(\mathcal{V}_u(v))}{\sum_{k=1}^{|\Pi_u|} \exp(\mathcal{V}_u(\Pi_u^k))} \right) - \lambda \|\Phi\|^2, \quad (13)$$

s. t. $\alpha_{uv}, \beta_{uv}, \lambda_{uv}, \gamma_{uv}, \delta_{uv} \in (0, 1)$

where $\mathcal{V}_u(\Pi_u^k)$ represents the prospect value of the k -th product in itemset Π_u and can be calculated by Eq.(2), Eq.(3) and Eq.(4). Meanwhile, we set the reference point \hat{r}_u in Eq.(2) as a learnable parameter for each consumer u . As a result, the parameter set is

$$\Phi = \{g^\alpha, g^\beta, g^\lambda, g^\gamma, g^\delta, \\ \forall u: \{\hat{r}_u, b_u^\alpha, \mathbf{p}_u^\alpha, b_u^\beta, \mathbf{p}_u^\beta, b_u^\lambda, \mathbf{p}_u^\lambda, b_u^\gamma, \mathbf{p}_u^\gamma, b_u^\delta, \mathbf{p}_u^\delta\}, \\ \forall v: \{l_v^\alpha, \mathbf{q}_v^\alpha, l_v^\beta, \mathbf{q}_v^\beta, l_v^\lambda, \mathbf{q}_v^\lambda, l_v^\gamma, \mathbf{q}_v^\gamma, l_v^\delta, \mathbf{q}_v^\delta\}\}. \quad (14)$$

Suppose we have n users and m items in our dataset and we set the dimension of the latent embedding as k , for each reparameterized parameter, which can be seen as a matrix X ($X = \alpha, \beta, \lambda, \gamma, \delta$) with elements represented as x_{uv} , we will have a one-dimension global bias vector, a n -dimension user bias vector, a m -dimension item bias vector, a $n \times k$ user embedding matrix and a $m \times k$ item

embedding matrix. Thus, the total number of the learning parameters is $5 + 5 \times (n + m)(k + 1)$ plus n , represented for a n -dimension reference point vector \hat{r} with elements represented as \hat{r}_u . We use log softmax loss as the loss function and batch stochastic gradient descent (SGD) as an optimizer to learn the model parameters. Meanwhile, people's risk attitudes may not be static and may vary over time. To overcome this, we can retrain our model based on their recent interaction with the system and update consumers' personalized risk preferences.

4.5 Top-K Recommendation

Once we learned the model parameters $\alpha_{uv}, \beta_{uv}, \lambda_{uv}, \gamma_{uv}, \delta_{uv}$ and \hat{r}_u according to our model, we can calculate the prospect value \mathcal{V}_{uv} for each user-item pair to rank all the products for a user. Then we select the top K items to generate the top-K recommendation list.

5 EXPERIMENTS

5.1 Dataset Description

We use the consumer transaction data from different sources — *Amazon*¹ [17, 31], *MovieLens*² [15], *Ciao* and *Epinions*³ [42, 43] — in our experiments to verify the recommendation performance of **RARE**⁴. The *Amazon* dataset includes user transaction (user id, item id, rating, etc.) and item metadata (item id, price, related item, etc.) on twenty-four product categories lasting from May 1996 to July 2014. We pick the largest three categories (namely *Books*, *Movies&TV*, and *Electronics*), which have different sizes and data sparsity, for experiments. The original data is huge and highly sparse, which makes it challenging to evaluate. Therefore, we first filter out items without prices, as the value function needs item prices as input; then, similar to previous work [19, 22], we filter out users and items with fewer than ten interactions. Meanwhile, we choose *MovieLens1M* dataset, which includes one million user transactions (user id, item id, rating, timestamp, etc.). This dataset is much denser compared with the *Amazon* datasets, *Ciao* and *Epinions*, where each user has at least 20 ratings. Considering that three *Amazon* datasets and *MovieLens1M* only contain one certain category of items, we do the experiments on another two widely used datasets — *Ciao* and *Epinions*, which contain items from different categories. Especially, *Ciao* contains items purchased by users in 28 categories, and *Epinions* has 27 categories in total. Since we need prices to calculate the value function, which are not recorded by *MovieLens1M*, *Ciao*, and *Epinions*, we set the price of each item in these three datasets equal to the same number, for example, \$1. For each dataset, we sort the transactions of each consumer according to the purchase timestamp and then split the records into training, validation, and testing sets chronologically. We then adopt the widely used leave-one-out evaluation [7, 19, 22], concretely, the test set contains user's most recent transaction, the validation set contains user's second most recent transaction, and all remaining interactions are for training. Some basic statistics of the experimental datasets are shown in Table 2. We notice that *MovieLens1M* is the densest dataset and has adequate actions per

¹<http://jmcauley.ucsd.edu/data/amazon/>

²<https://grouplens.org/datasets/movielens/>

³<https://www.cse.msu.edu/~tangjili/datasetcode/truststudy.htm>

⁴<https://github.com/TobyGE/Risk-Aware-Recommendation-Model>

Table 2: Basic statistics of the experimental datasets.

Dataset	#users	#items	#act./user	#act./item	#act.	density
Movies& TV	25,431	10,470	28.6	69.4	726,857	0.273%
Electronics	40,983	16,286	13.5	34.2	556,227	0.083%
Books	157,510	126,578	29.5	36.8	4653780	0.023%
MovieLens1M	6040	3706	166	270	1,000,209	4.468%
Ciao	2248	16861	16.0	2.14	36065	0.095%
Epinions	22164	296277	41.6	3.11	922267	0.014%

Table 3: Hyperparameter settings for each dataset.

Dataset	Latent Size	Learning Rate	L2 Penalty
Movies&TV	64	1e-3	1e-1
Electronics	16	1e-3	1e-1
Books	256	5e-3	1e-1
MovieLens1M	16	1e-4	1e-1
Ciao	16	1e-2	1e-1
Epinions	64	1e-2	1e-1

user and per item, while the three *Amazon* datasets have more users, items, and interactions compared with *MovieLens1M*.

5.2 Experimental Setup

We compare our model with the following baselines, including both economic and non-economic methods. For economic methods, we involve baselines that do not consider risk preferences to illustrate the importance of risk consideration.

SVD: Collaborative Filtering based on matrix factorization is a representative method for rating prediction. Basically, the user and item rating vectors are considered as the representation vector for each user and item. In this experiment, we use CF based on Singular Value Decomposition techniques [25, 39].

BPR-MF: Bayesian Personalized Ranking [37] is one of the most widely used ranking-based methods for the top-N recommendation. It is considered a classification problem with two classes that bought and not bought. Instead of classifying with the prediction of a single good, BPR tried to use the difference between two predictions. In the implementation, we conduct balanced negative sampling on unpurchased items for model learning.

NCF: Neural Collaborative Filtering [19] is one of the state-of-the-art recommender algorithms, which is based on deep neural networks. In the evaluation part, we choose Neural Matrix Factorization to conduct the experiments, fusing both Generalized Matrix Factorization (GMF) and Multiple Layer Perceptron (MLP) under the NCF framework.

MPUM: Multi-Product Utility Maximization for recommendation [53], which is an economic recommendation approach that maximizes the utility of product combinations for recommendation.

We implement SVD, BPR-MF, NCF and MPUM using *Pytorch* [35] with Adam optimizer. For all the methods, we consider latent dimensions d from {16, 32, 64, 128, 256}, learning rate lr from {1e-1, 5e-2, 1e-2, ..., 5e-4, 1e-4}, and the L2 penalty is chosen from {0.01, 0.1, 1}. We tune the hyper-parameters using the validation set and terminate training when the performance on the validation set does not change too much within 20 epochs. Meanwhile, to avoid heavy computation on all testing user-item pairs, we followed the mechanism in [10, 19, 24]. For each user i , we randomly sample

Table 4: Summary of the performance on six datasets. We evaluate for ranking (F_1 and $NDCG$, in percentage (%) values), and K is the length of recommendation list. When RARE is the best, its improvements against the best baseline are significant at $p < 0.01$.

Dataset	Metrics	F_1 Score (%)			$NDCG$ (%)		
	K	5	10	20	5	10	20
Movies&TV	SVD	5.288	4.393	3.302	12.29	15.13	17.86
	BPR-MF	6.010	5.496	4.488	13.56	17.72	22.07
	NCF	6.357	5.346	4.192	15.51	19.17	23.05
	MPUM	<u>6.732</u>	<u>5.865</u>	<u>4.573</u>	15.31	<u>19.43</u>	<u>23.77</u>
	RARE	7.144	6.024	4.710	16.67	20.68	24.89
Relative Improvements		+6.12	+2.71	+3.00	+7.48	+6.34	+4.71
Electronics	SVD	5.277	3.932	2.778	12.96	14.96	16.90
	BPR-MF	6.213	5.144	4.005	14.64	17.94	21.50
	NCF	<u>6.815</u>	5.397	3.968	<u>16.29</u>	<u>19.46</u>	22.55
	MPUM	6.783	5.566	4.129	15.94	19.45	22.74
	RARE	7.395	5.881	4.279	17.66	21.14	24.38
Relative Improvements		+8.51	+5.66	+3.63	+8.41	+8.63	+7.21
Books	SVD	5.330	3.992	2.815	13.04	15.10	17.05
	BPR-MF	9.639	7.489	5.265	23.12	27.33	30.98
	NCF	<u>10.32</u>	<u>7.643</u>	<u>5.282</u>	<u>25.72</u>	<u>29.52</u>	<u>32.90</u>
	MPUM	7.724	6.487	4.886	17.99	22.27	26.31
	RARE	10.79	7.902	5.395	27.00	30.81	34.22
Relative Improvements		+4.55	+3.39	+2.14	+4.98	+4.37	+4.01
MovieLens	SVD	13.02	9.037	5.763	32.77	36.44	39.24
	BPR-MF	14.51	10.44	6.944	36.19	40.97	44.97
	NCF	15.17	<u>10.60</u>	<u>6.861</u>	39.69	<u>44.10</u>	<u>47.66</u>
	MPUM	<u>15.53</u>	10.29	6.468	<u>40.67</u>	44.09	47.03
	RARE	16.88	11.21	7.064	44.25	47.92	51.25
Relative Improvements		+8.69	+5.75	+2.96	+8.80	+8.66	+7.53
Ciao	SVD	7.103	4.812	3.258	18.54	20.33	22.32
	BPR-MF	<u>10.49</u>	7.190	<u>4.588</u>	<u>26.97</u>	<u>29.75</u>	<u>31.98</u>
	NCF	7.770	5.670	4.029	19.63	22.33	25.19
	MPUM	10.02	<u>7.401</u>	4.550	23.66	27.33	29.19
	RARE	11.54	8.250	5.275	28.23	31.96	34.56
Relative Improvements		+10.0	+11.5	+15.0	+4.67	+7.43	+8.07
Epinions	SVD	11.98	7.604	4.345	31.69	33.71	34.76
	BPR-MF	<u>19.01</u>	11.73	6.766	<u>50.79</u>	<u>53.40</u>	<u>55.09</u>
	NCF	18.06	11.46	6.791	47.86	50.92	53.07
	MPUM	18.31	<u>12.03</u>	<u>7.136</u>	46.74	50.65	52.90
	RARE	19.91	12.78	7.416	52.31	55.95	58.11
Relative Improvements		+4.73	+6.23	+3.92	+2.99	+4.78	+5.48

100 negative items (items that are not interacted by the user) and rank these items with the positive sample in the test set. Based on the ranking results of the 101 items, we adopt two common Top-K metrics — F_1 Score and $NDCG$ — to evaluate each recommender model’s performance. F_1 Score considers both the precision and the recall of the test to compute the score, while $NDCG$ is a position-aware metric, involving a discount function over the rank. We implement **RARE** with *Pytorch* and fine-tune hyperparameters on our validation set, and the detailed settings on each dataset for it are shown in Table 3.

5.3 Experimental Results

The major experimental results are shown in Table 4, besides, we also plot the $NDCG$ in Figure 2 under different length of recommendation list K (from 1 to 50). We analyze and discuss the results in terms of the following perspectives.

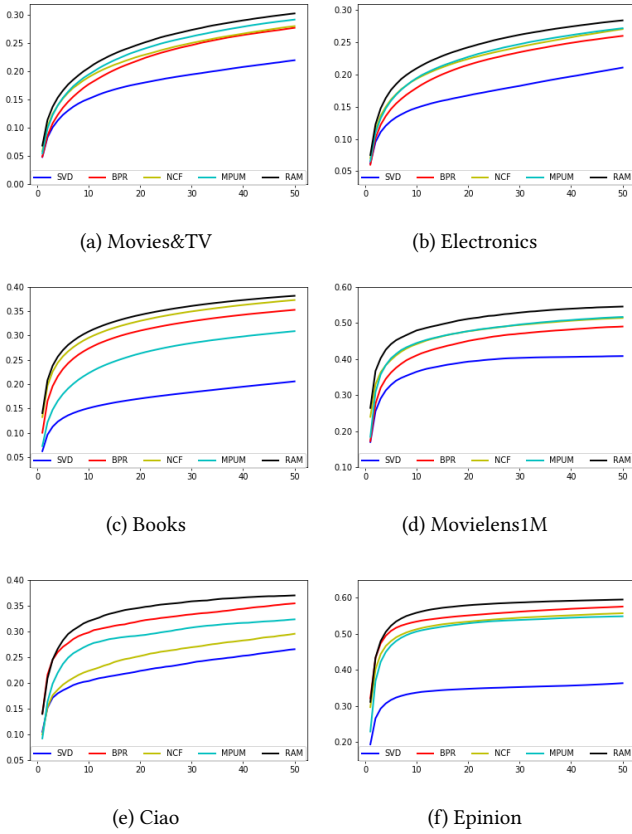


Figure 2: NDCG on six datasets. x -axis is the length of the recommendation list and y -axis is the NDCG.

i) Recommendation Performance: The results of the recommendation performance are shown in Table 4. The largest value on each dataset and for each evaluation measure is significant at 0.01 level. Among all the baseline models, we can see that all pair-wise learning methods (BPR, NCF, MPUM, and RARE) are much better than the simple point-wise (SVD) method, which demonstrates the superiority of pair-wise learning to rank methods on top-K ranking tasks. Further, among all the baselines, BPR is the strongest: when averaging across all datasets and recommendation lengths, BPR gets 41.0% improvement than SVD, 2.27% than NCF, and 1.60 % than MPUM.

Our RARE approach achieves the best top-K recommendation performance against all baselines on six datasets from four different sources. On the one hand, for F1 Score, we get 10.11% improvement than the BPR baseline when averaged across K on all of the six datasets, and the largest improvement (18.87%) is achieved when $K = 5$ on *Movies&TV* dataset. On the other hand, for NDCG, RARE relatively improves 12.83% than BPR on average. Compared with deep recommendation baseline NCF, our model gets 13.82% improvement on F1 on average, and especially a 48.52% improvement for F1@5 on *Ciao* dataset; furthermore, our model also gets 13.46% improvement than NCF for NDCG on average, and a 43.81% improvement for NDCG@5 on *Ciao* dataset. Meanwhile, in order to prove the validity of our improvement, we also plot the NDCG for different lengths of recommendation lists on all datasets, shown

Table 5: Summary of the performance against ablations on six datasets. When RARE is the best, its improvements against the best ablation are significant at $p < 0.01$.

Dataset	Metrics	F_1 Score (%)			NDCG (%)		
	K	5	10	20	5	10	20
Movies&TV	RARE-WF	4.749	3.772	2.702	11.11	13.34	15.31
	RARE-VF	2.015	1.961	2.013	4.666	6.274	8.919
	RARE-RP	3.872	3.100	2.191	9.294	11.15	12.69
	RARE	7.144	6.024	4.710	16.67	20.68	24.89
Electronics	RARE-WF	6.585	5.418	3.182	16.28	18.96	22.05
	RARE-VF	1.805	2.082	2.126	3.784	5.830	8.619
	RARE-RP	4.387	3.118	2.101	10.96	12.37	13.64
	RARE	7.395	5.881	4.279	17.66	21.14	24.38
Books	RARE-WF	4.979	3.620	2.450	12.48	14.19	15.70
	RARE-VF	2.116	2.322	2.405	4.508	6.687	9.880
	RARE-RP	4.670	3.392	2.292	11.68	13.28	14.68
	RARE	10.79	7.902	5.395	27.00	30.81	34.22
Movielens	RARE-WF	14.42	8.823	5.065	39.28	41.11	42.33
	RARE-VF	2.130	1.860	1.788	5.320	6.610	8.786
	RARE-RP	8.742	5.205	4.092	24.75	25.55	29.16
	RARE	16.88	11.21	7.064	44.25	47.92	51.25
Ciao	RARE-WF	9.356	6.584	4.135	23.51	26.34	28.22
	RARE-VF	4.864	3.680	2.453	11.53	13.49	14.92
	RARE-RP	9.060	6.446	4.007	22.55	25.41	27.11
	RARE	11.54	8.250	5.275	28.23	31.96	34.56
Epinions	RARE-WF	9.475	5.866	3.370	25.46	26.79	27.61
	RARE-VF	8.727	7.501	4.589	20.14	25.32	27.19
	RARE-RP	8.484	5.284	3.153	14.79	24.12	25.16
	RARE	19.91	12.78	7.416	52.31	55.95	58.11

in Fig.2. We can find that RARE outperforms all baseline for larger K , for example, $K = 50$. As shown, the black curve represented as RARE, is always above the best baselines in all six datasets. These observations imply that by modeling user behaviors under uncertainty based on established risk-aware principles, our model has the ability to capture better user preferences resulting in better recommendation results.

ii) Behavioral Economics vs. Classical Economics: We also have a strong economic recommender baseline — MPUM. When comparing with MPUM, our model gets 12.97% improvement on average. Especially, for the F1 Score, RARE gets 10.45% improvement against MPUM; for NDCG, the improvement is larger, increasing to 15.48 %. We believe there are several reasons to explain this improvement. First, the economic intuition behind MPUM is to make use of the substitutability and complementarity between products. However, most of our datasets only contain items from one category, which may decrease their performance. Second, the utility function implemented in their model is the well-known Constant Elasticity Substitution (CES) utility function, which is always risk-aversion [30], while Prospect Theory Value Function is a more delicate and comprehensive method that involves the risk attitudes and considers the difference between gains and losses, which allows for finer-grained modeling of consumer behaviors in online commercial environments.

5.4 Ablation Studies

In order to have a better understanding of the role of each part of our model, i.e., the value function and weighting function, we

further conduct ablation studies to analyze our model and answer the following research questions:

- Whether users have their own judgments of the extent of gains and losses?
- Whether users have their own sensitivity of the happening of possible events?
- Whether users have their own expectations of the outcome of probabilistic events?

Based on the above questions, we design three variants of RAM and compare RAM with each of them. The detailed forms of these variants are shown as following.

RARE without Value Function (RARE-VF): Prospect theory for recommendation without using the personalized value function, is a variant of our model that does not consider the personalized value judgments introduced by value function. We keep the global bias of each parameter in the original value function (e.g. g_α , g_β , g_λ) and remove the remaining personalized parts to create the non-personalized one, shown as Eq. (15); meanwhile, the weighting function in this variant is still be the same as Eq. (3).

$$v(x_i) = \begin{cases} g_\lambda x_i^{g_\alpha} & \text{if } x_i \geq 0 \\ -(-x_i)^{g_\beta} & \text{if } x_i < 0 \end{cases} \quad (15)$$

RARE without Weighting Function (RARE-WF): Prospect theory for recommendation without using the personalized weighting function is another variant of our model, which does not consider the personalized nonlinear transformation introduced by weighting function. Instead, it uses the original probabilities directly, which is shown in Eq. (16); meanwhile, the value function of this variant is still the same as Eq. (2).

$$\pi(p_i) = p_i \quad (16)$$

RARE without Reference Point (RARE-RP): Prospect theory for recommendation without using the reference point, which is the third variant of our model. A very fundamental idea in prospect theory is that people have different risk attitudes towards gains and loss, by removing reference point, we eliminate this cornerstone from the theory. Since we do not use reference point \hat{r} in this case, there will only exist gains for consumers, because $x_i = \text{price} \cdot \tanh(r_i) \geq 0$. The (positive) value function and (positive) weighting function used in this variant are shown as follows,

$$v(x_i) = x_i^\alpha \quad (17)$$

$$\pi(p_i) = \frac{p_i^\gamma}{p_i^\gamma + (1 - p_i)^\gamma} \quad (18)$$

Key experimental results of the ablation study are shown in Table 5. We compare, analyze, and discuss the three variants of RARE in terms of the following perspectives.

5.4.1 The role of Personalized Value Function. In RARE-VF (RARE w/o Personalized Value Function), we set $\alpha = g_\alpha$, $\lambda = g_\lambda$ and $\beta = g_\beta$ in the value function and maintain the weighting function same as RARE. From Table 5, comparing with our RARE, RARE-VF decreases 66.55 % when averaged across K on all datasets. We can also notice that among all six datasets, RARE-VF is always the worst, which, in turn, indicates the importance of the personalized value function. Thus, we can conclude that people do have their

own subjective judgments over the objective values of the products, particularly, gains and losses.

5.4.2 The role of Personalized Weighting Function. When we use the original probability without the weighting function, the whole method is noted as RARE-WF (RARE w/o Personalized Weighting Function). According to Table 5, our RARE model has a better performance against RARE-WF. When averaged across all datasets and recommendation lengths, RARE-WF has a decline of 33.70 % on F1 Score and 30.46 % on NDCG. The decline indicates that users with different risk attitudes should have different estimations, even with the same event probabilities. This personalized weighting function, as a nonlinear transformation of the probability scale, will predict a subjective probability for each user, so it can help the model better estimate users' risk attitudes in order to get better recommendation performances.

5.4.3 The role of Personalized Reference Point. The reference point is an essential part of Prospect Theory. It implies the consumer's subjective threshold of gain and loss. In the third variance of RARE, i.e., RAW-RP, we set the reference point equal to 0, which means that every outcome will be gain for consumers. In Table 5, we can see that RARE achieves higher scores than RARE-RP. When averaged across all datasets and recommendation lengths, RARE-RP relatively decreases more than 46% on average. In fact, setting all reference points to zero is equal to eliminate the loss of wealth; in other words, all the items are "seemingly" safe to the users. This might let the model recommend risky items to risk-aversion people, which causes bad performances. Thus, consumers possess a personalized threshold of gains and losses when facing different outcomes.

6 CONCLUSIONS AND FUTURE WORK

In this paper, we introduce risk attitudes into personalized recommendation systems and propose risk-aware recommendation. In particular, we bridge prospect theory and machine learning algorithms together to predict individual risk attitudes. We believe that online shopping could be risky with uncertainty, and the risk attitudes of different consumers may affect their decision making processes under uncertainty/risk. Understanding the risk attitude of consumers can help the system to predict human behaviors accurately for better services. Meanwhile, we advance prospect theory into a personalized version based on machine learning over large-scale consumer transaction logs. Experimental results verified the effectiveness of our model in terms of top-K recommendation. In the future, we will consider user risk attitudes in other online systems beyond e-commerce recommendation, and consider other economic principles and/or learning methods to benefit recommendation systems both effectively and economically.

ACKNOWLEDGEMENT

The authors thank the anonymous reviewers for the careful reviews and constructive suggestions.

REFERENCES

- [1] Binay Kumar Adhikari and Anup Agrawal. 2016. Does local religiosity matter for bank risk-taking? *Journal of Corporate Finance* 38 (2016), 272–293.
- [2] Maurice Allais. 1953. L'extension des théories de l'équilibre économique général et du rendement social au cas du risque. *Econometrica, Journal of the Econometric Society* (1953), 269–290.
- [3] Leif Azzopardi. 2014. Modelling interaction with economic models of search. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*. 3–12.
- [4] Daniel Bernoulli. 1668. *Specimen theoriae novae de mensura sortis*. Gregg.
- [5] Djallel Bouneffouf. 2014. R-UCB: a Contextual Bandit Algorithm for Risk-Aware Recommender Systems. *CoRR abs/1408.2195* (2014). arXiv:1408.2195 <http://arxiv.org/abs/1408.2195>
- [6] Djallel Bouneffouf, Amel Bouzeghoub, and Alda Lopes Ganarski. 2013. Risk-Aware Recommender Systems. In *Neural Information Processing*, Minho Lee, Akira Hirose, Zeng-Guang Hou, and Rhee Man Kil (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 57–65.
- [7] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishii Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ipsir, et al. 2016. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*. ACM, 7–10.
- [8] Michele Cohen, Jean-Yves Jaffray, and Taniols Said. 1987. Experimental comparison of individual behavior under risk and under uncertainty for gains and for losses. *Organizational behavior and human decision processes* 39, 1 (1987), 1–22.
- [9] T.A. Domencich and D. McFadden. 1975. *Urban travel demand: a behavioral analysis*. NORTH-HOLLAND. <https://books.google.com/books?id=Nx4pAQAAIAAJ>
- [10] Ali Mamdouh Elkahky, Yang Song, and Xiaodong He. 2015. A multi-view deep learning approach for cross domain user modeling in recommendation systems. In *Proceedings of the 24th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 278–288.
- [11] Daniel Ellsberg. 1961. Risk, ambiguity, and the Savage axioms. *The quarterly journal of economics* (1961), 643–669.
- [12] Andrew J. Flanagan, Miriam J. Metzger, Rebekah Pure, Alex Markov, and Ethan Hartsell. 2014. Mitigating risk in e-commerce transactions: perceptions of information credibility and the role of user-generated ratings in product quality and purchase intention. *Electronic Commerce Research* 14, 1 (01 Mar 2014), 1–23. <https://doi.org/10.1007/s10660-014-9139-2>
- [13] Zuohui Fu, Yikun Xian, Ruoyuan Gao, Jieyu Zhao, Qiaoying Huang, Yingqiang Ge, Shuyuan Xu, Shijie Geng, Chirag Shah, Yongfeng Zhang, and Gerard de Melo. 2020. Fairness-Aware Explainable Recommendation over Knowledge Graphs. In *Proceedings of SIGIR 2020*. ACM, New York, NY, USA.
- [14] Yingqiang Ge, Shuyuan Xu, Shuchang Liu, Shijie Geng, Zuohui Fu, and Yongfeng Zhang. 2019. Maximizing marginal utility per dollar for economic recommendation. In *The World Wide Web Conference*. 2757–2763.
- [15] F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. *ACM Trans. Interact. Syst.* 5, 4, Article 19 (Dec. 2015), 19 pages. <https://doi.org/10.1145/2827872>
- [16] Jerry Hausman and Daniel McFadden. 1984. Specification tests for the multinomial logit model. *Econometrica: Journal of the Econometric Society* (1984), 1219–1240.
- [17] Ruining He and Julian McAuley. 2016. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *proceedings of the 25th international conference on World Wide Web*. 507–517.
- [18] Ruining He and Julian McAuley. 2016. VBPR: Visual Bayesian Personalized Ranking from Implicit Feedback. In *AAAI*.
- [19] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. In *WWW*. 173–182.
- [20] David Hillson and Ruth Murray-Webster. 2017. *Understanding and managing risk attitude*. Routledge.
- [21] Daniel Kahneman. 1979. Prospect theory: An analysis of decisions under risk. *Econometrica* 47 (1979), 278.
- [22] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 197–206.
- [23] Joseph A Konstan, Bradley N Miller, David Maltz, Jonathan L Herlocker, Lee R Gordon, and John Riedl. 1997. GroupLens: applying collaborative filtering to Usenet news. *Commun. ACM* 40, 3 (1997), 77–87.
- [24] Yehuda Koren. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 426–434.
- [25] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 8 (2009), 30–37.
- [26] Georg Lackermair, Daniel Kailer, and Kenan Kanmaz. 2013. Importance of online product reviews from a consumer's perspective. *Advances in economics and business* 1, 1 (2013), 1–5.
- [27] Daniel D Lee and H Sebastian Seung. 2001. Algorithms for non-negative matrix factorization. In *Advances in neural information processing systems*. 556–562.
- [28] Juha Leino and Kari-Jouko Räihä. 2007. Case Amazon: Ratings and Reviews As Part of Recommendations. In *Proceedings of the 2007 ACM Conference on Recommender Systems (RecSys '07)*. ACM, New York, NY, USA, 137–140. <https://doi.org/10.1145/1297231.1297255>
- [29] Greg Linden, Brent Smith, and Jeremy York. 2003. Amazon.Com Recommendations: Item-to-Item Collaborative Filtering. *IEEE Internet Computing* 7, 1 (Jan. 2003), 76–80. <https://doi.org/10.1109/MIC.2003.1167344>
- [30] Robert E. Marks. 2015. Searching for Agents' Best Risk Profiles. In *Proceedings of the 18th Asia Pacific Symposium on Intelligent and Evolutionary Systems, Volume 1*, Hisashi Handa, Hisao Ishibuchi, Yew-Soon Ong, and Kay Chen Tan (Eds.). Springer International Publishing, Cham, 297–309.
- [31] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton van den Hengel. 2015. Image-based recommendations on styles and substitutes. In *SIGIR*. ACM.
- [32] Andriy Mnih and Ruslan R Salakhutdinov. 2008. Probabilistic matrix factorization. In *Advances in neural information processing systems*. 1257–1264.
- [33] Zan Mo, Yan-Fei Li, and Peng Fan. 2015. Effect of online reviews on consumer purchase behavior. *Journal of Service Science and Management* 8, 03 (2015), 419.
- [34] Shameem Ahamed Puthiya Parambath, Nishant Vijayakumar, and Sanjay Chawla. 2019. Risk Aware Ranking for Top-k Recommendations. *CoRR abs/1904.05325* (2019). arXiv:1904.05325 <http://arxiv.org/abs/1904.05325>
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*. 8024–8035.
- [36] Bibek Paudel, Sandro Luck, and Abraham Bernstein. 2018. Loss Aversion in Recommender Systems: Utilizing Negative User Preference to Improve Recommendation Quality. *CoRR abs/1812.11422* (2018). arXiv:1812.11422 <http://arxiv.org/abs/1812.11422>
- [37] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence*. AUAI Press, 452–461.
- [38] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, and John Riedl. 1994. GroupLens: an open architecture for collaborative filtering of netnews. In *Proceedings of the 1994 ACM conference on Computer supported cooperative work*. ACM, 175–186.
- [39] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2011. Introduction to recommender systems handbook. In *Recommender systems handbook*. Springer, 1–35.
- [40] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*. ACM, 285–295.
- [41] Leonard J Savage. 1972. *The foundations of statistics*. Courier Corporation.
- [42] Jiliang Tang, Huiji Gao, Xia Hu, and Huan Liu. 2013. Exploiting Homophily Effect for Trust Prediction. In *WSDM*.
- [43] Jiliang Tang, Huiji Gao, Huan Liu, and Atish Das Sarma. 2012. eTrust: Understanding Trust Evolution in an Online World. In *KDD*.
- [44] Amos Tversky and Daniel Kahneman. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty* 5, 4 (1992), 297–323.
- [45] Lev Virine and Michael Trumper. 2016. *ProjectThink: why good managers make poor project choices*. Routledge.
- [46] John Von Neumann and Oskar Morgenstern. 2007. *Theory of games and economic behavior (commemorative edition)*. Princeton university press.
- [47] Jian Wang and Yi Zhang. 2011. Utilizing marginal net utility for recommendation in e-commerce. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*. ACM, 1003–1012.
- [48] Menahem E. Yaari. 1965. Convexity in the Theory of Choice Under Risk. *The Quarterly Journal of Economics* 79, 2 (1965), 278–290. <http://www.jstor.org/stable/1880632>
- [49] Yongfeng Zhang, Qingyao Ai, Xu Chen, and W Bruce Croft. 2017. Joint representation learning for top-n recommendation with heterogeneous information sources. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. ACM, 1449–1458.
- [50] Yongfeng Zhang, Min Zhang, Yi Zhang, Guokun Lai, Yiqun Liu, Honghui Zhang, and Shaoping Ma. 2015. Daily-aware personalized recommendation based on feature-level time series analysis. In *Proceedings of the 24th international conference on world wide web*. 1373–1383.
- [51] Yongfeng Zhang, Qi Zhao, Yi Zhang, Daniel Friedman, Min Zhang, Yiqun Liu, and Shaoping Ma. 2016. Economic recommendation with surplus maximization. In *Proceedings of the 25th International Conference on World Wide Web*. 73–83.
- [52] Qi Zhao, Yi Zhang, Daniel Friedman, and Fangfang Tan. 2015. E-commerce recommendation with personalized promotion. In *Proceedings of the 9th ACM Conference on Recommender Systems*. ACM, 219–226.
- [53] Qi Zhao, Yongfeng Zhang, Yi Zhang, and Daniel Friedman. 2017. Multi-product utility maximization for economic recommendation. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. ACM, 435–443.