

On Estimating Multi-Attribute Choice Preferences using Private Signals and Matrix Factorization

Venkata Sriram Siddhardh Nadendla and Cedric Langbort

Coordinated Science Laboratory,
University of Illinois at Urbana-Champaign.
Email: {nadendla, langbort}@illinois.edu

Abstract—Revealed preference theory studies the possibility of modeling an agent’s revealed preferences and the construction of a consistent utility function. However, modeling agent’s choices over preference orderings is not always practical and demands strong assumptions on human rationality and data-acquisition abilities. Therefore, we propose a simple generative choice model where agents are assumed to generate the choice probabilities based on latent factor matrices that capture their choice evaluation across multiple attributes. Since the multi-attribute evaluation is typically hidden within the agent’s psyche, we consider a signaling mechanism where agents are provided with choice information through private signals, so that the agent’s choices provide more insight about his/her latent evaluation across multiple attributes. We estimate the choice model via a novel multi-stage matrix factorization algorithm that minimizes the average deviation of the factor estimates from choice data. Simulation results are presented to validate the estimation performance of our proposed algorithm.

I. MOTIVATION

In this paper, we consider the problem of modeling and estimating choice preferences across multiple attributes in a non-parametric manner via active data acquisition, where private signals are used to learn hidden preferences across multiple attributes. Such designs will significantly reduce search friction in current recommendation systems which are designed to model the average behavior of a population of agents. Furthermore, such choice models can be used in modern cyber-social-physical systems (e.g. smart transportation systems), where persuasive signals can be designed to steer agents’ decisions towards a social optimum.

One of the first attempts to model human decisions was made by Samuelson in his seminal work in [1] in 1938 using *the theory of revealed preferences*. This model is based on the idea that human agents choose decisions to maximize their utility functions. Later, several models have been proposed with utility-maximization as

a cornerstone philosophy to characterize human decisions. Some well-accepted models in the literature are Von Neumann and Morgenstern’s expected utility theory (EUT) [2], Herbert Simon’s *bounded rationality* models [3] and random utility based discrete choice models [4]. However, the problem of finding utility functions that are consistent with the revealed choices remained open until Afriat proposed a constructive algorithm in [5] to compute utility functions from finite choice revelations.

Several parametric choice models have been proposed to capture a person’s choice evaluation based on multiple attributes in the past literature. The characterization of multi-attribute choice models can be broadly classified into two types, depending on how evaluations across multiple attributes are combined together by the agent before picking a choice. In the first type of models, utility functions across multiple attributes are linearly combined together using appropriate weights, which are estimated from data. Such models have been extensively studied in the context of parametric models, with one of the most accepted models being the conditional multinomial logit (MNL) model proposed by McFadden in [6]. Alternatively, in the second type of models, conditional preferences across multiple attributes are combined in the realm of choice probabilities. Some examples include generalized extreme-value (GEV) models [4] and the mixture-MNL (MMNL) model [7], [8].

Although utility maximization has provided a tractable framework to design systems for human agents, several experiments have been documented in the psychology literature where human agents demonstrate deviating behaviors from the framework of utility maximization. For example, principles of *transitivity* [9] and *substitutability* [10] have been violated by human agents under various choice settings. Therefore, we avoid the utilitarian regime and other structural assumptions considered traditionally to model human decisions, and propose a non-parametric choice probability model where the agent is assumed to evaluate choices based on controller’s

signals which provide information about choices in terms of multiple attributes. Such a choice model facilitates the design of signaling mechanisms that strategically influence agents’ decisions in a desired manner.

The design of signaling mechanisms have a wide range of applications. For example, consider the example of smart transportation system (Alice), where a commuter (Bob) chooses a route from the set of all routes that connect start and destination nodes. On the other hand, the transportation agencies (e.g. GPS industry) are interested in controlling network congestion via providing traffic signals. While public signals are constructed to serve a population of commuters, private signals are designed to leverage an individual’s cognitive preferences across the set of all routes. Since private signals outperform public signals in most practical cases [11], there is a greater potential to influence commuters and control traffic congestion via private signaling.

II. LITERATURE SURVEY & CONTRIBUTIONS

Non-parametric choice models have been proposed by several researchers in the past. The most well-known one is the rank-based choice model where the agent is assumed to pick a choice with the highest rank. The agent’s decisions are modeled using a demand function which is characterized by the probability mass function over the set of all possible orders of choice probabilities. This model was first proposed by Block and Marschak [12] in 1960, and was later studied and analyzed in various applications. The caveat in this model is that it is practically impossible to observe the entire choice ordering in any decision-theoretic context. A possible remedy for this issue was provided by Jagabathula and Shah in [13], where they have considered the problem of modeling choice orders in the presence of constrained data. Later, Farias *et al.* have used this choice model to solve a revenue prediction problem in [14], [15]. However, it is not always practical to obtain choice rankings from human decision makers. For example, such queries are not relevant to the decision framework and can unnecessarily irk human agents.

Therefore, we consider a choice model where we are concerned with revealed choices rather than preference orderings. This is primarily because preference orders cannot be practically observed in every decision-theoretic framework, and is also not needed in system which only rely on choice probabilities. Furthermore, people evaluate choices across multiple attributes using different utility functions which need not be compatible with each other. In an attempt to address all these issues, we assume that the choice preferences (conditioned on the choice information signaled strategically to the agent)

are characterized by a matrix factorization model based on stochastic matrices, where the matrix factors capture the agent’s choice evaluation across multiple attributes. In order to find the optimal matrix factors in our choice model, we also propose a novel multi-stage estimation procedure which relies on choice revelations that are collected after the agent receives the strategic signal regarding choice information.

Matrix factorization has been a very active topic for several decades, and has emerged as a powerful tool to analyze clusters and other features in datasets. In contrast to traditional approaches such as singular value decomposition (SVD) and LU decomposition, the framework of non-negative matrix factorization (NMF) has become a versatile tool for dimensionality reduction and has been used in various contexts and applications. For more details, the readers may refer to a comprehensive review on NMF in [16]. Stochastic matrix factorization (SMF) is a constrained NMF, where a stochastic matrix is approximated as a product of two stochastic matrices. Although several generative probabilistic models (e.g. Latent Dirichlet Allocation models [17]) have been studied using various inference methods, the application of SMF for estimating these generative models was first proposed by Arora *et al.* in the context of topic modeling in [18] and have studied uniqueness in SMF in the presence of separability conditions. More recently, Adams has studied the SMF problem in [19], and investigated necessary and sufficient conditions on the observed data for a unique factorization. In addition, Adams has also derived bounds on the parameters for any observed data and presented a consistent least squares estimator. For details about the various SMF algorithms proposed in literature, the reader may refer to [20] and references within.

The main drawback of SMF is that the factors are not necessarily unique in general. Although the authors in [18] and [19] present conditions for the existence of a unique solution, these conditions need not necessarily hold true in general, in the context of choice modeling. Therefore, uniqueness of our matrix factorization approach is not guaranteed, and can affect the estimation performance. This analysis is beyond the scope of this paper, and will be considered in our future work. Instead, we focus our attention on our novel estimation procedure that finds stochastic factors in our choice model with the aid of information signaling in a greedy fashion.

III. CHOICE MODEL & PROBLEM SETUP

Consider two interacting agents, Alice and Bob, where Bob picks choices from a choice set $\mathcal{C} = \{1, \dots, K\}$ independently across time, and Alice’s goal is to evaluate




Fig. 1: Signaling scheme employed by Alice

Bob's choice model based on conditional preferences across a given set of attributes $\mathbb{A} = \{1, \dots, L\}$. We assume that Alice constructs private signals regarding the choice set across all possible subsets of the attribute set \mathbb{A} . These signals can be represented as a message matrix \mathbb{M} as shown in Figure 1, where each column corresponds to choice information based on a specific attribute subset $A \subseteq \mathbb{A}$. Based on this assumption, we propose an active data acquisition procedure where Alice constructs a message $m(A)$ for Bob based on a chosen attribute subset $A \subseteq \mathbb{A}$, by choosing the corresponding column in the matrix \mathbb{M} . For the sake of clarity, we denote the message set as $\mathbb{M} = \{m_1, \dots, m_M\}$, although each message in \mathbb{M} can be uniquely mapped to a subset of \mathbb{A} . This notation is employed to differentiate our message labels from the labels of attribute subsets, even though $M = 2^L$. Furthermore, we also assume that Alice sends a null message to Bob (denoted as m_1) when $A_1 = \phi$.

Having received a message signal $m_t = m$, Bob picks a choice $c_t = c$ with probability $p_{m,c} = \mathbb{P}(c_t = c | m_t = m)$, where $p_{m,c}$ is computed as

$$p_{m,c} = \sum_{A \in 2^{\mathbb{A}}} w_{m,A} \cdot q_{A,c} \quad (1)$$

for all $c \in \mathbb{C}$, where $w_{m,A} = \mathbb{P}(A | m_t = m)$ is the weight that Bob assigns to the attribute set $A \subseteq \mathbb{A}$ based on the received message m , and $q_{A,c} = \mathbb{P}(c_t = c | A)$ is the probability that Bob would have preferred the choice c based on the attribute set $A \subseteq \mathbb{A}$. The choice model in Equation (1) can also be presented equivalently in matrix form as

$$P = WQ, \quad (2)$$

where P is a $M \times K$ stochastic matrix, W is a $M \times 2^L$ stochastic matrix and Q is a $2^L \times K$ stochastic matrix. In other words, if u_k is a vector of k ones for any positive

integer k , we have

$$Pu_K = u_M, \quad Wu_{2^L} = u_M, \quad \text{and} \quad Qu_K = u_{2^L}. \quad (3)$$

We assume that the data acquisition experiment is designed in a sequence of stages, where Alice picks an attribute set A_i in the i^{th} stage based on an increasing order of the size of attribute subsets in \mathbb{A} , until she exhausts all the subsets of \mathbb{A} . Although it does not matter which one is chosen between two subsets of same size, we assume

$$\begin{aligned} A_1 &= \phi \\ A_2 &= \{1\}, \quad \dots, \quad A_{L+1} = \{L\} \\ A_{L+2} &= \{1, 2\}, \quad \dots, \quad A_{L^2+1} = \{L-1, L\} \\ &\vdots \\ A_{2^L} &= \mathbb{A}. \end{aligned} \quad (4)$$

Generally, Bob evaluates the choices over all possible attribute sets in $2^{\mathbb{A}}$, since Alice does not reveal the attribute set A based on which recommendations are constructed. However, in this paper, we assume that Alice reveals the attribute set to Bob at every stage. Since Bob has the autonomy to either completely/partially accept Alice's recommendation, or even ignore it altogether, our choice model in Equation (1) reduces to

$$p_{m,c} = \sum_{B \subseteq A_i} w_{m,B} \cdot q_{B,c}, \quad (5)$$

where Alice constructs its message m based on the subset A_i in the i^{th} stage, and B is a subset of A_i over which Bob may evaluate its choice c . Furthermore, in the matrix representation of the choice model as given in Equation (2), the weight matrix in Equation (2) takes the following lower-triangular form at the end of the data-acquisition process:

$$W = \begin{bmatrix} 1 & 0 & \dots & 0 \\ w_{2,1} & 1 - w_{2,1} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ w_{M,1} & w_{M,2} & \dots & 1 - \sum_{i=1}^{2^L} w_{M,i} \end{bmatrix}, \quad (6)$$

with $w_{i,j} = 0$ if $A_j \not\subseteq A_i$.

Alice's goal is to find the factor-estimates \hat{W} and \hat{Q} for given data-set $\mathbb{P} = \{P(1), \dots, P(N)\}$ of N stochastic matrices, that best fits to the choice model in Equation (2). In practice, Alice can obtain the data-set \mathbb{P} by partitioning the revealed preference dataset into N partitions, and estimating $P(n)$ using the data from the n^{th} partition for all $n = 1, \dots, N$. Note that the signaling mechanism and the procedure of computing \mathbb{P} are abstracted out from the rest of this paper, in order to

focus on the estimation of Q and W .

Let us define the *average deviation* from the model parameters Q and W to the data set \mathbb{P} as

$$f(W, Q) = \frac{1}{N} \sum_{n=1}^N \|WQ - P(n)\|_F^2, \quad (7)$$

where $\|\Pi\|_F$ is the Frobenius norm of the matrix Π . Then, the problem can be formally stated as follows.

$$\begin{aligned} & \underset{W, Q}{\text{minimize}} && f(W, Q) \\ & \text{subject to} && 1. Wu_{2^L} = u_M \\ & && 2. Qu_K = u_{2^L} \\ & && 3. \mathbf{0}_{M \times 2^L} \leq W \leq \mathbf{1}_{M \times 2^L} \\ & && 4. \mathbf{0}_{2^L \times K} \leq Q \leq \mathbf{1}_{2^L \times K}. \\ & && 5. W \text{ satisfies Equation (6)}. \end{aligned} \quad (\text{P})$$

Let \hat{Q} and \hat{W} denote the solution to Problem (P). Note that the above optimization problem is nonconvex since the objective function has a bilinear structure. We decompose Problem (P) into multiple stages and present an approximate learning procedure that is inspired from the sequential signaling policy employed by Alice.

IV. SEQUENTIAL LEARNING ALGORITHM

For the sake of convenience, we employ the following notation in the rest of the paper. We represent the matrices P , Q and W as

$$P = \begin{bmatrix} \mathbf{p}_1^T \\ \vdots \\ \mathbf{p}_M^T \end{bmatrix}, \quad Q = \begin{bmatrix} \mathbf{q}_1^T \\ \vdots \\ \mathbf{q}_{2^L}^T \end{bmatrix} \quad \text{and} \quad W = \begin{bmatrix} \mathbf{w}_1^T \\ \vdots \\ \mathbf{w}_M^T \end{bmatrix}, \quad (8)$$

where \mathbf{p}_m^T , \mathbf{q}_m^T and \mathbf{w}_m^T are the m^{th} rows in P , Q and W respectively.

Given a message $\mathbf{m} \in \mathbb{M}$ based on some attribute subset $A \subseteq \mathbb{A}$, $P(\boldsymbol{\sigma}|\mathbf{m})$ depends on the terms $P(\boldsymbol{\sigma}|B)$ and $w(B|\mathbf{m})$, for all $B \subseteq A$. This motivates Alice to adopt the ordering of messages, which is identical to the ordering of alphabet subsets in Equation (4), so that the weight estimates from the previous stages can be used to reduce Problem (P) into a ordered sequence of 2^L convex programs, where each convex program (stage) corresponds to each message in the set of all possible messages at Alice. In the i^{th} stage, Alice collects revealed preference data from Bob corresponding to the i^{th} message m_i in \mathbb{M} .

Stage-1: In the first stage of the algorithm, Alice's message is based out of a null subset of attributes. This is equivalent to the case where the agent does not provide

any information regarding the attributes. In this special case, since $\mathbf{w}_1^T = [1 \ 0 \ \dots \ 0]$, the first rows of both P and Q shall be identical. This can be illustrated in the following sub-problem:

$$\begin{aligned} & \underset{\mathbf{q}_1}{\text{minimize}} && \frac{1}{N} \sum_{n=1}^N \|\mathbf{q}_1 - \mathbf{p}_1(n)\|^2 \\ & \text{subject to} && 1. \mathbf{q}_1^T u_K = 1 \\ & && 2. \mathbf{0} \leq \mathbf{q}_1 \leq \mathbf{1}. \end{aligned} \quad (\text{P1})$$

Problem P1 is a least-squares problem, and therefore, has a closed form solution which is given by

$$\hat{\mathbf{q}}_1 = \frac{1}{N} \sum_{n=1}^N \mathbf{p}_1(n). \quad (9)$$

Stage- i : Now, let us consider the i^{th} stage where Alice collects Bob's revealed preferences conditioned on the message m_i , for any $i = 2, \dots, 2^L$. In this stage, the choice model that is of interest is given by

$$\mathbf{p}_i^T = \mathbf{w}_i^T Q \quad (10)$$

for all $i = 2, \dots, 2^L$. Due to the lower-triangular structure of the weight matrix W , \mathbf{p}_i is reduced to

$$\mathbf{p}_i = \sum_{j=1}^{i-1} w_{i,j} \mathbf{q}_j + \left(1 - \sum_{j=1}^{i-1} w_{i,j}\right) \mathbf{q}_i. \quad (11)$$

Note that, in the above model, the vectors $\mathbf{q}_1, \dots, \mathbf{q}_{i-1}$ have been estimated in the previous stages. Therefore, we equivalently represent the above model as

$$\mathbf{p}_i = \Phi_i \mathbf{x}_i, \quad (12)$$

where

$$\Phi_i = \begin{bmatrix} \mathbf{q}_1 & \dots & \mathbf{q}_{i-1} & I_K \end{bmatrix}, \quad (13)$$

denotes the coefficient matrix (known at the current stage since we have estimates $\hat{\mathbf{q}}_1, \dots, \hat{\mathbf{q}}_{i-1}$ from the previous stages) in the i^{th} stage, and,

$$\mathbf{x}_i = \begin{bmatrix} w_{i,1} \\ \vdots \\ w_{i,i-1} \\ \left(1 - \sum_{j=1}^{i-1} w_{i,j}\right) \mathbf{q}_i \end{bmatrix} \quad (14)$$

denotes the vector to be estimated in this stage.

While the above representation in Equation (12) seems linear, the nonlinearity is hidden in the bilinear term

$\left(1 - \sum_{j=1}^{i-1} w_{i,j}\right) \mathbf{q}_i$ in \mathbf{x}_i . In our proposed approximate algorithm, we ignore this bilinear nature of \mathbf{x}_i and evaluate the solution to the following linear program:

$$\begin{aligned} & \underset{\mathbf{x}_i}{\text{minimize}} && \frac{1}{N} \sum_{n=1}^N \|\hat{\Phi}_i \mathbf{x}_i - \mathbf{p}_i(n)\|^2 \\ & \text{subject to} && 1. \mathbf{u}_{K+i-1}^T \mathbf{x}_i = 1 \\ & && 2. \mathbf{0} \leq \mathbf{x}_i \leq \mathbf{1}, \end{aligned} \quad (\text{P2})$$

where $\hat{\Phi}_i = [\hat{\mathbf{q}}_1 \ \cdots \ \hat{\mathbf{q}}_{i-1} \ I_K]$ is the coefficient matrix whose entries are estimated in the previous stages.

Let $\hat{\mathbf{x}}_i$ denote the solution to Problem P2. If $\hat{\Phi}_i^{-1}$ is the MoorePenrose pseudo-inverse of the matrix $\hat{\Phi}_i$, $\hat{\mathbf{x}}_i$ is given by

$$\hat{\mathbf{x}}_i = \frac{1}{N} \sum_{n=1}^N \hat{\Phi}_i^{-1} \mathbf{p}_i(n). \quad (15)$$

If $\hat{\Phi}_i^{-1}$ cannot be computed, the problem can also be solved using efficient convex optimization algorithms. Given the solution vector $\hat{\mathbf{x}}_i$, we find the estimates of $\{w_{i,1}, \dots, w_{i,i-1}\}$ and \mathbf{q}_i , as shown below.

$$\hat{w}_{i,j} = \hat{\mathbf{x}}_i(j), \text{ for all } j = 1, \dots, i-1, \quad (16)$$

$$\hat{\mathbf{q}}_i = \frac{1}{1 - \sum_{j=1}^{i-1} \hat{\mathbf{x}}_i(j)} \mathbf{x}_i(i : K+i-1) \quad (17)$$

Next, we analyze the existence of a unique solution to Problem (P), and the optimality of our proposed algorithm.

Lemma 1. *Problem (P) does not have a unique optimal solution. Furthermore, if \mathcal{S} denotes the set of all possible optimal solutions to Problem (P), then the solution (\hat{W}, \hat{Q}) delivered by our proposed algorithm lies in \mathcal{S} .*

Proof. Consider the average deviation $f(W, Q)$, defined in Equation (7). Due to the lower-triangular structure of W matrix, we find that $\|WQ - P(n)\|_F^2$ can be expanded as

$$\begin{aligned} & \|WQ - P(n)\|_F^2 \\ &= \sum_{i=1}^M \left\| \sum_{j=1}^{i-1} w_{i,j} \mathbf{q}_j + \left(1 - \sum_{j=1}^{i-1} w_{i,j}\right) \mathbf{q}_i - \mathbf{p}_i(n) \right\|^2. \end{aligned} \quad (18)$$

Rearranging the order of summation in $f(W, Q)$, we have

$$f(W, Q) = \sum_{i=1}^M \tilde{f}_i(\mathbf{w}_i, \mathbf{q}_1, \dots, \mathbf{q}_i), \quad (19)$$

where

$$\begin{aligned} & \tilde{f}_i(\mathbf{w}_i, \mathbf{q}_1, \dots, \mathbf{q}_i) \\ &= \frac{1}{N} \sum_{n=1}^N \left\| \sum_{j=1}^{i-1} w_{i,j} \mathbf{q}_j + \left(1 - \sum_{j=1}^{i-1} w_{i,j}\right) \mathbf{q}_i - \mathbf{p}_i(n) \right\|^2. \end{aligned} \quad (20)$$

Given that each term in the summation is non-negative in the above representation, we have

$$\begin{aligned} & \min_{W, Q} \sum_{i=1}^M \tilde{f}_i(\mathbf{w}_i, \mathbf{q}_1, \dots, \mathbf{q}_i) \\ &= \sum_{i=1}^M \min_{W, Q} \tilde{f}_i(\mathbf{w}_i, \mathbf{q}_1, \dots, \mathbf{q}_i), \end{aligned} \quad (21)$$

if all the problems within the summation on the right hand side of Equation (21) are carried out simultaneously. This is equivalent to the case where we solve the following system of equations simultaneously, where each equation corresponds to a minimized term in the summation:

$$\mathbf{q}_1^* = \frac{1}{N} \sum_{n=1}^N \mathbf{p}_1(n), \quad (22)$$

$$\sum_{j=1}^{i-1} w_{i,j}^* \mathbf{q}_j^* + \left(1 - \sum_{j=1}^{i-1} w_{i,j}^*\right) \mathbf{q}_i^* = \frac{1}{N} \sum_{n=1}^N \mathbf{p}_i(n), \quad (23)$$

for all $i = 2, \dots, M$.

Since there are more variables than equations in the above system of simultaneous equations, we do not have a unique solution to this problem. Note that our proposed algorithm presents one of the solution candidates to the above system of simultaneous equations, since it presents

$$(\hat{\mathbf{w}}_i, \hat{\mathbf{q}}_i) = \arg \min_{\mathbf{w}_i, \mathbf{q}_i} \tilde{f}_i(\mathbf{w}_i, \hat{\mathbf{q}}_1, \dots, \hat{\mathbf{q}}_{i-1}, \mathbf{q}_i)$$

in the i^{th} stage for all $i = 1, \dots, M$. \square

Given that the size of the message set \mathbb{M} increases exponentially with the number of attributes, the algorithm has exponential complexity in terms of the size of the attribute set. However, we have closed-form expressions to both \hat{Q} and \hat{W} , as given in Equations (9), (16) and (17). Furthermore, the hierarchical structure in the power set of \mathbb{A} provides us with multiple non-interfering stages which reduces to L (or, equivalently $\log_2 M$) effective stages with the aid of parallel processing.

V. RESULTS AND DISCUSSION

Consider an example setting where there are $L = 2$ attributes, and therefore, a total of $M = 4$ possible message signals. We choose the stochastic matrices Q




Fig. 2: Average Deviation for increasing number of data samples, when $K = 5$, $L = 2$, $N = 5$ and 100 Monte-Carlo runs.

and W at random so that they satisfy the structural constraints. Having chosen the matrices Q and W , we compute the P matrix using our choice model in Equation (2). We sample C choices from P matrix, which are distributed across N bins to compute the input data $P(1), \dots, P(N)$ to our algorithm.

In Figure 2, we run this experiment for different number of i.i.d. data samples per bin, when there are $N = 5$ bins, $K = 5$ choices and $L = 2$ attributes. We repeat this experiment over 100 Monte Carlo runs and plot the average deviation $f(\hat{W}, \hat{Q})$, along with $\|P - \hat{P}\|_F$, $\|Q - \hat{Q}\|_F$ and $\|W - \hat{W}\|_F$, where $\hat{P} = \hat{W}\hat{Q}$. We plot the estimation error for increasing number of data samples (number of the agent’s choice revelations). Note that the average deviation $f(W, Q)$ and the estimation errors of Q and W matrices decrease, but do not converge to zero. This phenomenon can be attributed to the fact that our matrix factorization framework in Problem (P) does not have a unique solution.

In the future, we will study special structures of Q and W matrices that guarantee the existence of a unique solution to Problem (P), since the average deviation $f(W, Q)$ does not converge to zero in general. Furthermore, we will also investigate persuasive mechanisms where the controller constructs signals based on the agent’s estimated preferences to steer their decisions in a desired manner. We will also study the impact of information-framing, where a given message is presented relative to some reference.

REFERENCES

- [1] P. A. Samuelson, “A Note on the Pure Theory of Consumer’s Behaviour,” *Economica*, vol. 5, no. 17, pp. 61–71, 1938.
- [2] J. V. Neumann and O. Morgenstern, *Theory of games and economic behavior*. Princeton University Press, 1944.
- [3] A. Rubinstein, *Modeling Bounded Rationality*. MIT press, 1998.
- [4] K. Train, *Discrete Choice Methods with Simulation*. Cambridge university press, 2003.
- [5] S. N. Afriat, “The Construction of Utility Functions from Expenditure Data,” *International Economic Review*, vol. 8, no. 1, pp. 67–77, 1967.
- [6] D. L. McFadden, “Conditional Logit Analysis of Qualitative Choice Behavior,” in *Frontiers of Econometrics*, P. Zarembka, Ed. New York: Academic Press, 1973, pp. 105–142.
- [7] P. Rusmevichientong, D. Shmoys, C. Tong, and H. Topaloglu, “Assortment Optimization under the Multinomial Logit Model with Random Choice Parameters,” *Production and Operations Management*, vol. 23, no. 11, pp. 2023–2039, 2014.
- [8] N. Kallus and M. Udell, “Revealed Preference at Scale: Learning Personalized Preferences from Assortment Choices,” in *Proceedings of the 2016 ACM Conference on Economics and Computation*, ser. EC ’16, 2016, pp. 821–837.
- [9] A. Tversky, “Intransitivity of Preferences,” *Psychological Review*, vol. 76, no. 1, pp. 31–48, January 1969.
- [10] A. Tversky and J. E. Russo, “Substitutability and Similarity in Binary Choices,” *Journal of Mathematical Psychology*, vol. 6, no. 1, pp. 1–12, 1969.
- [11] C. Hellwig, “Public Information, Private Information, and the Multiplicity of Equilibria in Coordination Games,” *Journal of Economic Theory*, vol. 107, no. 2, pp. 191–222, 2002.
- [12] H. Block and J. Marschak, “Random Orderings and Stochastic Theories of Response,” in *Contributions to Probability and Statistics*, H. M. Olkin, Ghurye and Mann, Eds. Stanford University Press, 1960, pp. 97–132.
- [13] S. Jagabathula and D. Shah, “Inferring Rankings Using Constrained Sensing,” *IEEE Transactions on Information Theory*, vol. 57, no. 11, pp. 7288–7306, Nov 2011.
- [14] V. F. Farias, S. Jagabathula, and D. Shah, “Sparse Choice Models,” in *46th Annual Conference on Information Sciences and Systems (CISS)*, Princeton University, Princeton, NJ, USA., 2012, pp. 1–28.
- [15] —, “A Nonparametric Approach to Modeling Choice with Limited Data,” *Management Science*, vol. 59, no. 2, pp. 305–322, 2013.
- [16] Y. X. Wang and Y. J. Zhang, “Nonnegative Matrix Factorization: A Comprehensive Review,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 25, no. 6, pp. 1336–1353, June 2013.
- [17] D. M. Blei, A. Y. Ng, and M. I. Jordan, “Latent Dirichlet Allocation,” *Journal of Machine Learning Research*, vol. 3, no. Jan, pp. 993–1022, 2003.
- [18] S. Arora, R. Ge, and A. Moitra, “Learning Topic Models – Going Beyond SVD,” in *Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, ser. FOCS ’12, Washington, DC, USA, 2012, pp. 1–10.
- [19] C. P. Adams, “Stochastic Matrix Factorization,” *SSRN Electronic Journal*, p. 124, 2016. [Online]. Available: <https://ssrn.com/abstract=2840852>
- [20] M. Luo, F. Nie, X. Chang, Y. Yang, A. Hauptmann, and Q. Zheng, “Probabilistic Non-Negative Matrix Factorization and Its Robust Extensions for Topic Modeling,” in *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 2017.