
A Cost-Effective Framework for Preference Elicitation and Aggregation

**Zhibing Zhao, Haoming Li,
Junming Wang**
RPI
Troy, NY, USA
{zhaoz6,lih14,wangj33}@rpi.edu

**Jeffrey O. Kephart,
Nicholas Mattei, Hui Su**
IBM Research
Yorktown, NY, USA
{kephart,n.mattei,huisuibmres}@us.ibm.com

Lirong Xia
RPI
Troy, NY, USA
xial@cs.rpi.edu

Abstract

We propose a cost-effective framework for preference elicitation and aggregation under the Plackett-Luce model with features. Given a budget, our framework iteratively computes the most cost-effective elicitation questions in order to help the agents make a better group decision.

We illustrate the viability of the framework with experiments on Amazon Mechanical Turk, which we use to estimate the cost of answering different types of elicitation questions. We compare the prediction accuracy of our framework when adopting various information criteria that evaluate the expected information gain from a question. Our experiments show carefully designed information criteria are much more efficient, i.e., they arrive at the correct answer using fewer queries, than randomly asking questions given the budget constraint.

1 INTRODUCTION

Consider the hiring decision problem [Bhattacharjya and Kephart, 2014]. With the aid of an intelligent system, a group of people (the *key group*) faces a hiring decision about many candidates who are characterized by attributes, such as experiences, technical skills, communication skills, etc. The goal is to help the key group make a group decision without directly eliciting their full preferences over all candidates, which is often infeasible given the vast number of candidates. Instead, the intelligent system may ask fellow employees (the *regular group*) about their preferences in order to learn about the key group’s preferences. How can the intelligent system decide which member in the regular group to ask and

which questions to ask? Note that we discuss the presence of two groups but our framework is applicable when there is only one group of decision makers as well.

This example illustrates the *preference elicitation* problem, which has been widely studied in the field of recommender systems [Loepp et al., 2014], healthcare [Chajewska et al., 2000, Weermink et al., 2014, Erdem and Campbell, 2017], marketing [Huang and Luo, 2016], stable matching [Drummond and Boutilier, 2014, Rastegari et al., 2016], combinatorial auctions [Sandholm and Boutilier, 2006], etc. Most previous works studied a special case of the aforementioned scenario, in which the regular group is the key group. The objective of preference elicitation is to achieve some goal using as few samples (data) as possible. A common approach is to adaptively ask questions that maximize expected information gain, measured by some information criteria.

Moreover, most previous work focused on specific types of elicitation questions, e.g. pairwise comparisons. In this paper, we consider a more general framework that asks a variety of elicitation questions and can accommodate one or more groups. The diversity of elicitation questions enables us to query cost-effectively. Intuitively, an agent’s preference order over 10 alternatives tells us more about her preference in general than just her top choice among the 10; however, it may take her longer to do so. The key question we want to answer in this paper is:

How can we compute the most cost-effective questions for preference elicitation under resource constraints?

1.1 OUR CONTRIBUTIONS

We propose a flexible cost-effective preference elicitation and aggregation framework to predict a single agent’s preference or help make a group decision. The main inputs include a budget W , a set of designs (i.e. questions to ask) \mathcal{H} , a cost function w , a randomized

voting rule and an information criterion. We model non-deterministic preferences using the Plackett-Luce model with features.

Cost-effectiveness. We propose a flexible, cost-effective preference elicitation framework that accommodates all randomized voting rules, ranking models, and information criteria. This iterative framework leverages the *optimal design* technique. In each iteration, we choose the question that provides the most information per unit cost. The response is then recorded as a data point, leading to an update of the posterior distribution of the parameter, which is treated as the prior for the next iteration. In any iteration, the posterior estimate of the parameter can be used to compute a winner distribution using a randomized voting rule. This procedure is illustrated in Figure 1.




Figure 1: Illustration of the proposed framework.

Randomized Voting Rules. We use randomized voting rules to compute the winning alternatives of a group decision, which outputs the distribution of winners (see Section 4 for details). The probability for each alternative to be the winner is proportional to its score based on the voting rule. These probability estimates are more informative than only recommending a winner as it provides a distribution over the candidates as well.

We prove that when people have non-deterministic preferences, the probability of an alternative to be the winner is proportional to the total expected score of this alternative for all agents (Theorem 1). This means the randomized counterpart of any scoring rule can be used in our framework as long as the expected score of each alternative for a single agent is easy to compute. Then we prove that under the Plackett-Luce model, the winner distributions of probabilistic plurality and probabilistic Borda are easy to compute (Corollary 1 and Theorem 2).

Information Criteria. An information criterion plays a key role in determining the next elicitation question by measuring the information in the distribution of a parameter. We propose the minimum pairwise certainty

(MPC) criterion, extended from the information criterion by Azari Soufiani et al. [2013], which maximizes the improvement of the least certain pairwise comparison. Other commonly-used information criteria include D-optimality [Wald, 1943, Mood et al., 1946] and E-optimality [Ehrenfeld, 1955], as well as asking a question uniformly at random. All these information criteria are based on the information of the posterior distribution of the model parameter, which is approximated by its asymptotic distribution, a multivariate Gaussian computed based on the composite marginal likelihood method [Pauli et al., 2011].

Empirical Studies & Experiments. We carry out Amazon Mechanical Turk experiments to estimate the cost of answering various types of questions for a target domain of ranking hotels. We compare the performances of MPC, D-optimality and E-optimality with simulations and observe that these criteria have similar performance in terms of prediction accuracy, and we observe that all of them significantly outperform random elicitation questions.

1.2 RELATED WORK AND DISCUSSIONS

Our work is related to cost-effective experimental designs, which were investigated by Wright et al. [2010], Volkov [2014] in the context of aquatic toxicology and drug development, respectively. Volkov [2014] modeled cost-effectiveness as different types of optimization problems, e.g., minimize cost under information constraints. We take a greedy approach, similar to an algorithm proposed by Wright et al. [2010], and choose the design (elicitation question) that maximizes the expected information gain per unit cost. Our cost varies depending on the type of questions and is estimated empirically, similar to the idea in Volkov [2014]. To our best knowledge, this paper is the first work to apply cost-effective experimental design to preference elicitation.

The greedy approach is also called one-step-lookahead policy, which can be arbitrarily worse than optimal t -step-lookahead (t -step myopic active search) policies for $t \geq 2$ [Garnett et al., 2012]. Arbitrary t -step myopic active search is hard to compute, as was shown by Jiang et al. [2017], which also proved that nonmyopic active search is computationally hard even to approximate and proposed an efficient searching algorithm. This algorithm is potentially useful in the preference elicitation context and is an interesting future direction.

Most previous works in preference elicitation assumed that people’s preferences are deterministic. For example, Bhattacharjya and Kephart [2014] proposed an even swap algorithm to reveal a single decision maker’s most

preferred alternative; Lu and Boutilier [2011a], Kalech et al. [2011] elicited preferences from a group of people in order to make a group decision under a (deterministic) voting rule. In contrast, we consider non-deterministic preferences of people, which is often the case in real-world. Moreover, we use randomized voting rules, which output the probability of each alternative to be the winner. These probabilities, which can be viewed as normalized scores over all alternatives, provide a quantitative measure of the quality of each alternative. For example, an alternative that wins with probability 0.8 can be seen as being much better than other alternatives.

Non-deterministic preferences were modeled by general random utility models by Azari Soufiani et al. [2013]. They proposed a preference elicitation framework for personalized choice and social choice (aggregated preference). We use the Plackett-Luce model with features, which is a special case of general random utility models but has easy-to-compute probabilities. More importantly, we use randomized voting rules for aggregation, which is very different from parametric modeling of social choices employed by [Azari Soufiani et al., 2013].

Pairwise elicitation questions may be the most widely explored in the literature due to their simplicity [Branke et al., 2017, Eric et al., 2008, Houlby et al., 2012, Lu and Boutilier, 2011a, Pfeiffer et al., 2012]. In contrast, Azari Soufiani et al. [2013] focused on elicitation of full rankings, though their proposed framework also allows for partial orders. Drummond and Boutilier [2014], Lu and Boutilier [2011b] studied a larger set of queries, which includes asking a person to rank her top k choices over all alternatives. In this paper, we consider an even broader set of queries, asking an agent to rank her top k choices over a subset of l alternatives ($k < l$). This enables us to elicit preferences in a more cost-effective manner.

As a key role in preference elicitation, information criteria have been widely investigated for different applications. Standard information criteria include D-optimality (used in [Houlby et al., 2011, 2012, Pfeiffer et al., 2012]) and E-optimality. Drummond and Boutilier [2014] and Lu and Boutilier [2011a] use minimax-regret-based criterion for stable matching and aggregation respectively. Azari Soufiani et al. [2013] proposed yet another criterion, defined on the certainty of the least certain pairwise comparison over the intrinsic utilities (part of the parameter of their general random utility models) of all alternatives. Our MPC criterion extends the criterion by Azari Soufiani et al. [2013]. To predict a single agent’s top k preference, we search over a subset of all pairwise comparisons (see Section 3.3). To help make a group decision, we search over all pairwise comparisons

of all agents in the key group to find the least certain pairwise comparison (Equation (3)).

2 PRELIMINARIES

Let $\mathcal{A} = \{a_1, a_2, \dots, a_m\}$ denote a set of m alternatives and $\{1, \dots, n_1, n_1 + 1, \dots, n_1 + n_2\}$ denote $n_1 + n_2$ agents, where the first n_1 agents belong to the key group, who will be making a group decision. The remaining n_2 agents belong to the regular group. For all $i = 1, \dots, m$, a_i is characterized by a real-valued column vector of K attributes \vec{z}_i . For all $j = 1, \dots, n_1 + n_2$, agent j is characterized by a real-valued column vector of L attributes \vec{x}_j . A full ranking R is often denoted by $a_{i_1} \succ a_{i_2} \succ \dots \succ a_{i_m}$, where “ \succ ” means “is preferred over”. We denote the budget by W , where the money is used to pay the agents for answering elicitation questions.

For $n_1 = 1$, we want to predict the single key agent’s full or top k ranking with as much certainty as possible given a budget W . For $n_1 \geq 2$, the goal is to predict the winning alternative of the key group by eliciting preferences from the regular group in the most cost-effective way. More concretely, given W , we want to output a distribution of winning alternatives, w.r.t. a randomized voting rule, which will be defined in Section 4.

2.1 THE PLACKETT-LUCE MODEL WITH FEATURES

Let the parameter $B = [b_{\kappa\ell}]_{K \times L}$ be a matrix of real-valued coefficients, transforming features to utilities. Each value $b_{\kappa\ell}$ corresponds to the κ -th attribute from an alternative and ℓ -th attribute from an agent. The parameter space Θ is a set of all real-valued $K \times L$ matrices. Then the utility of an alternative a_i to an agent j is

$$u_{ji} = \vec{x}_j^\top B \vec{z}_i. \quad (1)$$

For any agent j and any full ranking $R_j = a_{i_1} \succ a_{i_2} \succ \dots \succ a_{i_m}$, the probability of R_j is

$$\Pr(R_j) = \frac{\exp(u_{ji_1})}{\sum_{q=1}^m \exp(u_{ji_q})} \times \frac{\exp(u_{ji_2})}{\sum_{q=2}^m \exp(u_{ji_q})} \times \dots \times \frac{\exp(u_{ji_{m-1}})}{\exp(u_{ji_{m-1}}) + \exp(u_{ji_m})}.$$

Given the Plackett-Luce model with features, the probability of alternative a_{i_1} to be ranked at the top among $\{a_{i_1}, \dots, a_{i_l}\}$ by agent j is $\frac{\exp(u_{ji_1})}{\sum_{q=1}^l \exp(u_{ji_q})}$. Specifically, for any two alternatives a_1 and a_2 , the probability of $a_1 \succ a_2$ by agent j is $\frac{\exp(u_{j1})}{\exp(u_{j1}) + \exp(u_{j2})}$.

2.2 ONE-STEP BAYESIAN EXPERIMENTAL DESIGN

Given any probabilistic model parameterized by $B \in \Theta$ and any prior distribution $\pi(B)$, a one-step Bayesian experimental design consists of two parts: (i) a set of designs \mathcal{H} , where each $h \in \mathcal{H}$ is composed of an agent and a question; (ii) an information measure $G(\cdot)$, which maps any distribution of B over Θ to a real-valued scalar: a measure of information in this distribution.

For any design $h \in \mathcal{H}$, the distribution of responses can be computed using the ground truth parameter B^* . We use D to denote the set of all possible responses. Given a ground truth parameter B^* , the probability of any data $d \in D$ can be computed as $\Pr(d|h)$. Further, we can compute the posterior distribution of parameter $\pi(B|d, h)$ over the parameter space Θ and the corresponding information criterion $G(\pi(B|d, h))$. The expected information is

$$E[G(\pi(B|h))] = \sum_{d \in D} G(\pi(B|d, h)) \Pr(d|h),$$

where the expectation is taken over all possible responses. The goal is to find the design h that maximize the expected information gain, which is $E[G(\pi(B|h))] - G(\pi(B))$, per unit cost. Let $w(h)$ denote the cost function, which maps the 2-tuple (agent, question) to a positive cost. Given the cost function $w(h)$, we can compute the optimal design h^* that maximizes the expected information gain per unit cost by

$$h^* = \arg \max_h \frac{E[G(\pi(B|h))] - G(\pi(B))}{w(h)}. \quad (2)$$

3 COST-EFFECTIVE PREFERENCE ELICITATION

In our proposed framework, we iteratively adapt the one-step experimental design by querying the most cost-effective question in each iteration. At any iteration t , the prior distribution of B is the posterior distribution given data D^t , i.e. $\pi(B^t|D^t)$. Given this posterior, we find the most cost-effective design h^t , which consists of one agent and one question, and query h^t . The response is combined with D^t to form D^{t+1} . Then the budget W and the set of designs \mathcal{H} are updated before going to the next iteration. Finally, when $n_1 = 1$, we compute the predicted preference of this agent; when $n_1 \geq 2$, we compute the distribution of winners based on a randomized voting rule. This framework is formally illustrated in Algorithm 1.

For the rest of this section, we will explain how to approximate the posterior distribution $\pi(B^t|D^t)$ and how $G(\pi(B))$ is computed.

Algorithm 1 Cost-Effective Preference Elicitation

Input: Budget W , randomize voting rule r , cost function $w(h)$, information criterion $G(\pi(B))$, the set of designs \mathcal{H} where for any $h \in \mathcal{H}$, $w(h) \leq W$.

Output: A predicted preference when $n_1 = 1$ or a distribution of winning alternatives for group decision when $n_1 \geq 2$.

Initialization: Randomly initialize data D^1 .

while \mathcal{H} is not empty **do**

 Compute/approximate $\pi(B^t|D^t)$;

 Compute $h^t \in \mathcal{H}$ using (2);

 Implement h^t (query an agent a question). Let R^t denote her answer. Then $D^{t+1} \leftarrow D^t \cup \{R^t\}$, $\mathcal{H} \leftarrow \mathcal{H} - h^t$, $W \leftarrow W - w^t$;

 Remove all h^t 's from \mathcal{H} where $w(h^t) > W$.

end while

Compute the predicted preference when $n_1 = 1$ or a distribution of winning alternatives according to the voting rule r when $n_1 \geq 2$.

3.1 APPROXIMATION OF POSTERIOR DISTRIBUTION

For any prior $\pi(B)$ and data D , the posterior distribution is given by $\pi(B|D) = \frac{\Pr(D|B)\pi(B)}{\int_{\Theta} \Pr(D|B)\pi(B)dB}$ according to Bayes' rule. This posterior is often hard to compute. A commonly-used approach is to approximate it by its asymptotic distribution, which is a multivariate Gaussian distribution characterized by the composite marginal likelihood (CML) method [Pauli et al., 2011].

For convenience we vectorize B as a column vector, denoted by $\vec{\beta} = \text{vec}(B)$. The composite marginal likelihood method [Lindsay, 1988, Zhao and Xia, 2018] computes the estimate of the ground truth parameter from marginal events, e.g., pairwise comparisons. Let $\{\mathcal{E}_1, \dots, \mathcal{E}_q\}$ denote q selected marginal events. Then the composite marginal likelihood method computes the estimate $\vec{\beta}_{\text{CML}}$ by

$$\vec{\beta}_{\text{CML}} = \arg \max_{\vec{\beta} \in \Theta} \text{CLL}(\vec{\beta}) = \arg \max_{\vec{\beta} \in \Theta} \sum_{\lambda=1}^q \ln \Pr(\mathcal{E}_\lambda | \vec{\beta}),$$

where $\text{CLL}(\vec{\beta})$ denotes the composite log-likelihood function. Under our Plackett-Luce model with features, $\text{CLL}(\vec{\beta})$ is twice differentiable for all $\vec{\beta} \in \Theta$, i.e. $J(\vec{\beta}) = -\nabla_{\vec{\beta}}^2 \text{CLL}(\vec{\beta})$ exists. From Pauli et al. [2011], asymptotically, $\pi(\vec{\beta}|D)$ is a multivariate Gaussian distribution, whose mean is $\vec{\beta}_{\text{CML}}$ and covariance matrix is $J^{-1}(\vec{\beta})$.

Computing $J(\vec{\beta})$ requires computation of second order partial derivatives of $\ln \Pr(\mathcal{E}_\lambda | \vec{\beta})$ for all λ . We will show the close-form second order partial derivative formula for

any response from an agent.

3.2 THE SET OF DESIGNS

Each design $h \in \mathcal{H}$ is a combination of an agent and a question about her preferences. The agent can be anyone from $\{1, \dots, n_1 + n_2\}$. In this paper, for simplicity, we consider the case where only the agents from the regular group $\{n_1 + 1, \dots, n_1 + n_2\}$ are queried. For any integers $k < l \leq m$, we may ask an agent to rank her top k alternatives over a subset of l alternatives. When $k = 1, l = 2$, the question is a pairwise comparison; when $k = 1, l > 2$, the question is to query an agent's top alternative among a subset of alternatives; when $k = l - 1$, we are asking a full ranking over a subset of alternatives. The advantage of this type of questions is that the probabilities of responses of these questions are easy to compute, as well as their partial derivatives. W.l.o.g. let $R_j = a_1 \succ a_2 \succ \dots \succ a_k \succ \text{others}$ be the answer from agent j . Then we have $\Pr(R_j|B) = \prod_{p=1}^k \frac{\exp(u_{jp})}{\sum_{i=p}^l \exp(u_{ji})}$, and $\ln \Pr(R_j|B) = \sum_{p=1}^k (u_{jp} - \ln \sum_{i=p}^l \exp(u_{ji}))$.

For any $1 \leq \kappa \leq K$ and $1 \leq \iota \leq L$, let $b_{\kappa\iota}$ be the (κ, ι) entry of B . We have

$$\frac{\partial \ln \Pr(R_j|B)}{\partial b_{\kappa\iota}} = \sum_{p=1}^k \left(\frac{\partial u_{jp}}{\partial b_{\kappa\iota}} - \frac{\sum_{i=p}^l \exp(u_{ji}) \frac{\partial u_{ji}}{\partial b_{\kappa\iota}}}{\sum_{i=p}^l \exp(u_{ji})} \right),$$

where $\frac{\partial u_{jp}}{\partial b_{\kappa\iota}}$ and $\frac{\partial u_{ji}}{\partial b_{\kappa\iota}}$ are constants (products of an agent's attribute and an alternative's attribute) by definition. Therefore, for diagonal entries, the second order partial derivatives are given by

$$\frac{\partial^2 \ln \Pr_j(R|B)}{\partial b_{\kappa\iota}^2} = \sum_{p=1}^k \left(\left(\frac{\sum_{i=p}^l \exp(u_{ji}) \frac{\partial u_{ji}}{\partial b_{\kappa\iota}}}{\sum_{i=p}^l \exp(u_{ji})} \right)^2 - \frac{\sum_{i=p}^l \exp(u_{ji}) \left(\frac{\partial u_{jp}}{\partial b_{\kappa\iota}} \right)^2}{\sum_{p=1}^l \exp(u_{jp})} \right),$$

and for non-diagonal entries, we have

$$\begin{aligned} & \frac{\partial^2 \ln \Pr_j(R|B)}{\partial b_{\kappa_1 \iota_1} \partial b_{\kappa_2 \iota_2}} \\ &= \sum_{p=1}^k \left(\frac{(\sum_{i=p}^l \exp(u_{ji}) \frac{\partial u_{ji}}{\partial b_{\kappa_1 \iota_1}}) (\sum_{i=p}^l \exp(u_{ji}) \frac{\partial u_{ji}}{\partial b_{\kappa_2 \iota_2}})}{(\sum_{i=p}^{m'} e^{u_{ji}})^2} \right. \\ & \left. - \frac{\sum_{i=p}^l \exp(u_{ji}) \left(\frac{\partial u_{ji}}{\partial b_{\kappa_1 \iota_1}} \right) \left(\frac{\partial u_{ji}}{\partial b_{\kappa_2 \iota_2}} \right)}{\sum_{i=p}^l \exp(u_{ji})} \right). \end{aligned}$$

3.3 INFORMATION CRITERIA

An information criterion maps the distribution of a parameter to a real-valued quality. Standard information

criteria are mostly directly computed from the covariance matrix $J^{-1}(\vec{\beta})$ or its inverse $J(\vec{\beta})$. For example, D-optimality [Wald, 1943, Mood et al., 1946] computes the determinant of $J(\vec{\beta})$; E-optimality [Ehrenfeld, 1955] computes the minimum eigenvalue of $J(\vec{\beta})$. We propose the following minimum pairwise certainty (MPC) criterion by extending the criterion from [Azari Soufiani et al., 2013] to our domain.

MPC for Case $n_1 = 1$. We consider two types of purposes: predicting the agent's (unordered) top k alternatives and predicting the agent's ranked top k alternatives. We note that the criterion by Azari Soufiani et al. [2013] only applies to full rankings, which is a special case of our ranked top k . The intuition of this criterion is to maximize the certainty of the least certain pairwise comparison among a subset of pairwise comparisons. Formally, let \mathcal{A}_k denote the set of predicted top k alternatives for this key agent.

- Unordered top- k where $1 \leq k < m$:

$$G(\pi(\vec{\beta})) = \min_{i_1 \in \mathcal{A}_k, i_2 \notin \mathcal{A}_k} \frac{|\text{mean}(u_{1i_1} - u_{1i_2})|}{\text{std}(u_{1i_1} - u_{1i_2})}.$$

- Ranked top- k where $1 < k < m$:

$$G(\pi(\vec{\beta})) = \min_{i_1 \in \mathcal{A}_k, i_2 \neq i_1} \frac{|\text{mean}(u_{1i_1} - u_{1i_2})|}{\text{std}(u_{1i_1} - u_{1i_2})}.$$

In the above equations, $\text{mean}(u_{ji_1} - u_{ji_2})$ is computed using $\vec{\beta}_{\text{CML}}$ and $\text{std}(u_{ji_1} - u_{ji_2})$ is computed using the approximated covariance matrix $J^{-1}(\vec{\beta})$ as follows.

Because $u_{ji_1} - u_{ji_2}$ is linear with $\vec{\beta}$ (see Equation (1) and recall that $\vec{\beta}$ is the vectorization of B), we write it as $u_{ji_1} - u_{ji_2} = \sum_{\kappa, \iota} c_{\kappa\iota} b_{\kappa\iota}$, where $c_{\kappa\iota}$'s are constants computed from attributes of a_{i_1}, a_{i_2} and agent j . Then we have $\text{std}(u_{ji_1} - u_{ji_2}) = \sqrt{\sum_{(\kappa_1, \iota_1), (\kappa_2, \iota_2)} c_{\kappa_1 \iota_1} c_{\kappa_2 \iota_2} \text{Cov}(b_{\kappa_1, \iota_1}, b_{\kappa_2, \iota_2})}$. When $\kappa_1 = \kappa_2 = \kappa$ and $\iota_1 = \iota_2 = \iota$, $\text{Cov}(b_{\kappa_1, \iota_1}, b_{\kappa_2, \iota_2})$ reduces to $\text{Var}(b_{\kappa\iota})$. Both $\text{Cov}(b_{\kappa_1, \iota_1}, b_{\kappa_2, \iota_2})$ and $\text{Var}(b_{\kappa\iota})$ are entries of $J^{-1}(\vec{\beta})$.

MPC for Case $n_1 \geq 2$. Our MPC for this case is different from the criterion by Azari Soufiani et al. [2013] in that we find the least certain pairwise comparison across all agents in the key group. Formally, our MPC for $n_1 \geq 2$ is

$$G(\pi(\vec{\beta})) = \min_{j \in \{1, \dots, n_1\}, i_2 \neq i_1} \frac{|\text{mean}(u_{ji_1} - u_{ji_2})|}{\text{std}(u_{ji_1} - u_{ji_2})}, \quad (3)$$

where the computation of $\text{mean}(u_{ji_1} - u_{ji_2})$ and $\text{std}(u_{ji_1} - u_{ji_2})$ are similar to the $n_1 = 1$ case.

4 RANDOMIZED VOTING RULES

We use randomized voting rules to aggregate the key group's preferences. A randomized voting rule computes the distribution of winners given the preferences of the agents from the key group. Under non-deterministic preferences, this distribution can be computed from the parameter of the model. This section shows that under the Plackett-Luce model with features, probabilistic plurality and probabilistic Borda are easy to compute.

A randomized voting rule assigns a probability for each alternative to be the winner according to the data, usually based on a scoring function. For example, the probabilistic plurality rule, which is equivalent to random dictatorship [Gibbard, 1977], samples a winner from a distribution where the probability of each alternative being the winner is proportional to the plurality score of this alternative. Other randomized voting rules can be defined similarly, including probabilistic Borda [Heckelman, 2003]. The voting rule must have scores associated with it, but this is a very mild restriction because many commonly-studied voting rules including all positional scoring rules, Copeland, range voting, and approval voting, have randomized counterparts.

Recall that n_1 agents are making a group decision among m alternatives. Let P denote the preference profile that consists of n_1 full rankings over m alternatives from the key group. Let $s_r(a_i, P)$ denote the score of alternative a_i under voting rule r and $\Pr_r(a_i|P)$ be the probability for a_i to win under the randomized analogy of r given P . Then $\Pr_r(a_i|P)$ is computed by $\Pr_r(a_i|P) = \frac{s_r(a_i, P)}{\sum_{i=1}^m s_r(a_i, P)}$.

Example 1 Suppose the set of alternatives is $\{a_1, a_2, a_3\}$ and the votes are $\{a_1 \succ a_2 \succ a_3, a_1 \succ a_3 \succ a_2, a_2 \succ a_1 \succ a_3\}$. The plurality and Borda scores are shown in Table 1. Under probabilistic plurality rule, a_1 wins with probability $2/3$ and a_2 wins with probability $1/3$. Under probabilistic Borda, a_1, a_2, a_3 win with probabilities $5/9, 3/9, 1/9$ respectively.

	a_1	a_2	a_3
plurality	2	1	0
Borda	5	3	1

Table 1: Scores under plurality and Borda

We consider non-deterministic preferences from agents, where the preferences from the key agents are independent of each other. Because each agent has $m!$ possible rankings, there are $(m!)^{n_1}$ possible preference profiles. Then we have $\Pr_r(a_i) = \sum_{q=1}^{(m!)^{n_1}} \Pr(P_q) \Pr_r(a_i|P_q)$,

where P_q denotes the q -th possible preference profile.

Given a voting rule r , let X_{ji} be the score of a_i for agent j . X_{ji} is a random variable due to the uncertainty of agent j 's preference. The following theorem shows that the probability of a_i being the winner is proportional to the sum of expected score of a_i for each agent.

Theorem 1 For any $1 \leq i \leq m$, $\Pr_r(a_i) \propto \sum_j^{n_1} EX_{ji}$.

Proof: It suffices to prove $\Pr_r(a_1) \propto \sum_j^{n_1} EX_{j1}$.

By definition, $\Pr_r(a_1) = \sum_{q=1}^{(m!)^{n_1}} \Pr(P_q) \Pr_r(a_1|P_q)$. Let S denote the score of a_1 under rule r . Then S is a random variable defined over the $(m!)^{n_1}$ cases. Let s_q denote the value that S takes for case q . In any case q , we have $\Pr_r(a_1|P_q) \propto s_q$ by the definition of randomized voting rules. We re-write it as $\Pr_r(a_1|P_q) = \frac{s_q}{M}$, where M is the normalization factor. Observe that across all the $(m!)^{n_1}$ cases, M does not change because the voting rule r and the set of agents does not change. So we have

$$\begin{aligned} \Pr_r(a_1) &= \frac{\sum_{q=1}^{(m!)^{n_1}} \Pr(P_q) s_q}{M} \\ &\propto \sum_{q=1}^{(m!)^{n_1}} \Pr(P_q) s_q = ES. \end{aligned} \quad (4)$$

Since $S = \sum_{j=1}^{n_1} X_{j1}$, due to linearity of expectation, we have $ES = E[\sum_{j=1}^{n_1} X_{j1}] = \sum_j^{n_1} EX_{j1}$. By (4), we have $\Pr_r(a_1) \propto \sum_j^{n_1} EX_{j1}$. ■

For probabilistic plurality, the expected score of a_i for agent j is exactly the probability of a_i being ranked at the top by agent j . So we have the following corollary:

Corollary 1 Let $p_j^{a_i}$ be the probability of a_i being ranked at the top by agent j . For any $1 \leq i \leq m$, $\Pr_{\text{plurality}}(a_i) = \frac{1}{n_1} \sum_j^{n_1} p_j^{a_i}$.

Theorem 2 Let $p_j^{a_i \succ a_{i'}}$ denote the probability for agent j to prefer alternative a_i over $a_{i'}$. Then for any $1 \leq i \leq m$, $\Pr_{\text{Borda}}(a_i) \propto \sum_j^{n_1} \sum_{i' \neq i} p_j^{a_i \succ a_{i'}}$.

Proof: For any $i \in \{1, \dots, m\}$, we have $\Pr_{\text{Borda}}(a_i) \propto \sum_j^{n_1} EX_{ji}$ by Theorem 1, where X_{ji} here denotes the score of a_i for agent j under Borda. We only need to prove $EX_{ji} = \sum_{i' \neq i} p_j^{a_i \succ a_{i'}}$. This is a known result, but we were not able to find a formal proof in literature, except a proof for three alternatives by Chen and Heckelman [2005], which is easy to be extended for arbitrary number of alternatives. For completeness we provide a short proof.

By definition of Borda, we have $EX_{ji} = \sum_{k=1}^{m-1} (m-k) \sum_{R: a_i \text{ at } k\text{th position of } R} \Pr_j(R)$, where R is any full

ranking over the m alternatives and $\Pr_j(R)$ is the probability of R by agent j . Imagine $m-1$ bins, each of which is labeled with $a_i \succ a_{i'}$ for all the remaining $m-1$ $a_{i'}$'s. Observe that there are $m-k$ copies of $\Pr_j(R)$ for all R where a_i beats exactly $m-k$ other alternatives. We can distribute the $m-k$ copies to the $m-k$ bins (one in each) for all $a_{i'}$'s that are ranked after a_i . We do this for all possible rankings and in the end, each bin labeled by $a_i \succ a_{i'}$ gets the probabilities of all rankings compatible with $a_i \succ a_{i'}$. This finishes the proof. ■

We note that for the Plackett-Luce model, $p_j^{a_i}$'s and $p_j^{a_i \succ a_{i'}}$'s are easy to compute (see Section 2.1).

5 EXPERIMENTS

We first introduce an example of empirically estimating the cost of asking different types of questions on MTurk. Then, we show the result of a simulation of cost-effective preference elicitation using synthetic data.

Rank the options, then

1	3.8 of 5.0 \$116 25 min 54 min
2	3.9 of 5.0 \$125 7 min 32 min
3	4.4 of 5.0 \$179 5 min 35 min
4	3.9 of 5.0 \$101 50 min 4 min
<hr style="border: 1px solid green;"/>	
5	3.9 of 5.0 \$112 42 min 16 min
6	4.4 of 5.0 \$249 8 min 41 min
7	4.3 of 5.0 \$102 9 min 28 min
8	4.1 of 5.0 \$143 6 min 40 min
9	3.2 of 5.0 \$100 46 min 10 min
10	4.7 of 5.0 \$164 4 min 32 min

Figure 3: The user interface for a Turker to submit her ranked top 4 over 10 alternatives. The attributes are average ratings, prices per night, time to Times Square, and time to the nearest airport.

5.1 ESTIMATING $w(h)$

We recall that a question is defined by a pair of parameters (k, l) , where l is the number of alternatives that are presented to an agent and k is the number of alternatives that the agent is asked to rank at the top k positions.

In order to map the question types to the time to answer them, we run 2 experiments with multiple tasks on MTurk. Each task required MTurk workers to report their preferences over a set of hotels. We recorded the time they spent on each task, in order to learn such mapping in the following two cases:

- $k, l \in [2, 10], k = l - 1$: full rankings;
- $k \in [1, 10], l = 10$: ranked top k alternatives over 10.

Experiment Setting. For the first case, we looked for information on the first 54 Hotels in New York City in alphabetical order. We split the 54 randomly into 9 sets, each containing 2, 3, ..., 10 hotels. We then showed the 9 sets to MTurk workers, randomizing the order of the 9 sets as well as the initial display order of alternatives within each set, and asked them to rearrange by drag-and-dropping the alternatives according to their preferences. The alternatives were anonymous and represented by 4 attributes: average guest rating on a popular travel website, price per night, time to Times Square and time to the nearest airport.

For the second case, a separate experiment is run with another 10 sets of NYC hotels, drawn randomly again from the first 100 hotels in NYC in alphabetical order. In each task, we placed a horizontal green bar underneath the alternative above which are the k alternatives of interest. We instructed the MTurk workers before the experiment started that only the alternatives above the green bar would count, i.e. only needed to rank-order top- k . All of these were done with goal of minimizing the overhead time for workers to understand the instruction so that the recorded time accurately reflect the time of decision-making. An example of the UI is show in Figure 3, where we asked Turkers to rank-order her top-4 favorite hotels over a set of 10.

The following analysis is made possible by responses from 408 MTurk workers (202 for the first case and 206 for the second).

Experiment Results. Although Volkov [2014] considered both linear and quadratic cost function and argued for the superiority of the latter, for simplicity, we perform a linear regression on the dataset to obtain a linear cost function. We regress the time to rearrange the alternatives and submit a full ranking on the number of alternatives in the set. We find that, on average, the time a Turker spent on rank-ordering a full ranking over l alternatives is $t_{\text{full-}l} = 5.33l$ (Figure 2 left), and that on rank-ordering her top- k alternatives over 10 alternatives is $t_{\text{top-}k} = 1.38k + 32.25$ (Figure 2 right). In addition, the 408 workers spent an average of 341.5 seconds on the tasks and were each paid \$0.3. Therefore, the monetary cost of elicitation on average is $w = 0.00088t$, which correspond to an hourly wage of \$3.16.




Figure 2: The left subfigure shows the average time a user spent to submit a full ranking over 2, . . . , 10 alternatives; the right subfigure shows the average time a user spent to give her ranked top 1, . . . , 10 alternatives when 10 alternatives were proposed.

Combining these two functions with the hourly wage, we propose the following cost functions, which estimates the cost (in USD) of elicitation about hotel preferences given 4 alternative attributes: $w_{\text{full}-l} = 0.0047l$ and $w_{\text{top}-k} = 0.0012k + 0.028$. We observe that the time a user spent is not very sensitive to k . This is sensible, as when a MTurk worker ranks her top k choices, she may follow the following procedure: 1. read the descriptions of all hotels, 2. form their preferences, and 3 choose top k . Step 1 and 2 do not depend on k and dominates the time for step 3, as illustrated in the right figure of Figure 2. This suggests that when a fixed number of alternatives is proposed to an agent, it’s likely that the most cost-effective question to ask is a full ranking, as we will see in the next subsection.

5.2 COST-EFFECTIVE PREFERENCE ELICITATION

We demonstrate the viability of our cost-effective framework and compare performances of different information criteria on synthetic data.

Synthetic Data. We randomly generated 10 alternatives, each of which has 3 attributes, independently normally distributed $N(0, 1)$. We then randomly generated 5 agents that forms the key group and 20 the regular group. Each agent also has 3 attributes, independently normally distributed $N(0, 1)$. B was generated from Dirichlet distribution $\text{Dir}(\vec{1})$. The result is averaged over 400 trials.

To echo the motivating example from the beginning of this paper, we simulated the process of eliciting key group’s preference by asking agents in the regular group questions. For simplicity, we consider 3 types of ques-

tions, represented in (k, l) : $(1, 2)$, $(1, 10)$ and $(9, 10)$. We run Algorithm 1 using 3 different information criteria: D-Optimality, E-Optimality, and the proposed MPC. The three elicitation processes utilized the cost function estimated in Section 5.1. They were initialized with the same set of 50 randomly generated pairwise comparisons and was given a \$0.9 budget. Agents’ answers to the elicitation questions are generated from the Plackett-Luce model.

Metrics. We use total variation distance to measure the difference between the winner distributions computed from the ground truth parameter and the estimates, denoted by ψ^* and ψ respectively. The total variance distance is defined as $\delta(\psi^*, \psi) = \frac{1}{2} \sum_{i=1}^m |\psi^*(a_i) - \psi(a_i)|$. To plot the results, at each cost w , we used the data point that is below w but closest to w in each trial. These points were averaged over all trials.

Observations. We observe that for both probabilistic plurality and probabilistic Borda, the performances of MPC, D-optimality, and E-optimality are similar, all of which significantly outperforms random elicitation questions (see Figures 4). For example, for probabilistic plurality and probabilistic Borda, at the budget of 0.85 dollars, MPC achieves 15% less total variation distance than that of random elicitation questions. As another example, to achieve the total variation distance of 0.064 under randomized plurality (respectively, randomized Borda), MPC uses 20% (23.5%) less money than that of random elicitation questions.

We also observe that D-optimality almost always choose a full ranking as the most cost-effective question, while MPC tends to choose more full rankings than pairwise comparisons at early stages (see Figure 5). Due to the budget limit, many trials finish after 19 iterations because




Figure 4: Total variation distance for probabilistic plurality (left) and probabilistic Borda (right).




Figure 5: Types of questions chosen by the MPC (left) and D-optimality (right). The legend “Full Ranking”, “Top Choice”, and “Pairwise” correspond to $(k = 9, l = 10)$, $(k = 1, l = 10)$ and $(k = 1, l = 2)$ respectively.

they only query full rankings. Others finish at different iterations. The distribution of types of questions for E-optimality is similar to MPC. Under all criteria except random, $l = 10, k = 1$ questions were rarely asked.

Discussions. The meaning of cost-effectiveness in this paper is twofold: (1) in the preference elicitation procedure, we ask elicitation questions that is expected to provide more information per unit cost; and (2) the presence of regular group gives us a belief on the key group’s preferences inexpensively. As we have seen in our experiments, a budget of \$0.9 gives us a reasonably good estimate of the key group’s preferences by querying the regular group.

6 CONCLUSIONS AND FUTURE WORK

We proposed a flexible and cost-effective framework for preference elicitation that can be adapted for any ranking

model, any information criterion, and any set of questions. We used randomized voting rules to help make group decisions and proposed MPC for both prediction of one agent’s preference and aggregation of a group of agents’ preferences. Experiments show that MPC and other commonly-used information criteria work better than asking random elicitation questions. For future work we will explore better information criteria for group decisions. We also plan to extend this framework to multiple types of resource constraints.

Acknowledgements

We thank all anonymous reviewers for helpful comments and suggestions. This work is supported by NSF #1453542 and ONR #N00014-17-1-2621.

References

- Hossein Azari Soufiani, David C. Parkes, and Lirong Xia. Preference elicitation for general random utility models. In *Proceedings of Uncertainty in Artificial Intelligence (UAI)*, Bellevue, Washington, USA, 2013.
- Debarun Bhattacharjya and Jeffrey Kephart. Bayesian interactive decision support for multi-attribute problems with even swaps. In *Proceedings of the Thirtieth Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-14)*, Corvallis, Oregon, 2014.
- Juergen Branke, Salvatore Corrente, Salvatore Greco, and Walter Gutjahr. Efficient pairwise preference elicitation allowing for indifference. *Computers & Operations Research*, 88:175–186, 2017.
- Urszula Chajewska, Daphne Koller, and Ron Parr. Making rational decisions using adaptive utility elicitation. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 363–369, Austin, TX, USA, 2000.
- Frederick H Chen and Jac C Heckelman. Winning probabilities in a pairwise lottery system with three alternatives. *Economic Theory*, 26(3):607–617, 2005.
- Joanna Drummond and Craig Boutilier. Preference elicitation and interview minimization in stable matchings. In *AAAI*, pages 645–653, 2014.
- Sylvain Ehrenfeld. On the efficiency of experimental designs. *The annals of mathematical statistics*, 26(2): 247–255, 1955.
- Seda Erdem and Danny Campbell. Preferences for public involvement in health service decisions: a comparison between best-worst scaling and trio-wise stated preference elicitation techniques. *The European Journal of Health Economics*, 18(9):1107–1123, 2017.
- Brochu Eric, Nando D Freitas, and Abhijeet Ghosh. Active preference learning with discrete choice data. In *Advances in neural information processing systems*, pages 409–416, 2008.
- Roman Garnett, Yamuna Krishnamurthy, Xuehan Xiong, Jeff Schneider, and Richard Mann. Bayesian optimal active search and surveying. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.
- Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45:665–681, 1977.
- Jac C Heckelman. Probabilistic Borda rule voting. *Social Choice and Welfare*, 21(3):455–468, 2003.
- Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. Bayesian active learning for classification and preference learning. *stat*, 1050:24, 2011.
- Neil Houlsby, Jose Miguel Hernandez-Lobato, Ferenc Huszar, and Zoubin Ghahramani. Collaborative Gaussian processes for preference learning. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, pages 2105–2113, Lake Tahoe, NV, USA, 2012.
- Dongling Huang and Lan Luo. Consumer preference elicitation of complex products using fuzzy support vector machine active learning. *Marketing Science*, 35(3):445–464, 2016.
- Shali Jiang, Gustavo Malkomes, Geoff Converse, Alyssa Shofner, Benjamin Moseley, and Roman Garnett. Efficient nonmyopic active search. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1714–1723, Sydney, Australia, 06–11 Aug 2017. PMLR.
- Meir Kalech, Sarit Kraus, Gal A Kaminka, and Claudia V Goldman. Practical voting rules with partial information. *Autonomous Agents and Multi-Agent Systems*, 22(1):151–182, 2011.
- Bruce G Lindsay. Composite likelihood methods. *Contemporary Mathematics*, 80, 1988.
- Benedikt Loepp, Tim Hussein, and Jüergen Ziegler. Choice-based preference elicitation for collaborative filtering recommender systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3085–3094. ACM, 2014.
- Tyler Lu and Craig Boutilier. Robust approximation and incremental elicitation in voting protocols. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, pages 287–293, Barcelona, Catalonia, Spain, 2011a.
- Tyler Lu and Craig Boutilier. Vote elicitation with probabilistic preference models: Empirical estimation and cost tradeoffs. In *International Conference on Algorithmic Decision Theory*, pages 135–149. Springer, 2011b.
- Alexander M Mood et al. On hotelling’s weighing problem. *The Annals of Mathematical Statistics*, 17(4): 432–446, 1946.
- Francesco Pauli, Walter Racugno, and Laura Ventura. Bayesian composite marginal likelihoods. *Statistica Sinica*, pages 149–164, 2011.
- Thomas Pfeiffer, Xi Alice Gao, Andrew Mao, Yiling Chen, and David G. Rand. Adaptive polling and information aggregation. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 122–128, Toronto, Canada, 2012.
- Baharak Rastegari, Paul Goldberg, and David Manlove. Preference elicitation in matching markets via inter-

- views: A study of offline benchmarks. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 1393–1394, 2016.
- Tuomas Sandholm and Craig Boutilier. Preference elicitation in combinatorial auctions. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 10, pages 233–263. MIT Press, 2006.
- Oleg Volkov. *Optimal Relaxed Designs of Experiments, with Pharmaceutical Applications*. PhD thesis, Queen Mary University of London, 2014.
- Abraham Wald. On the efficient design of statistical investigations. *The annals of mathematical statistics*, 14(2):134–140, 1943.
- Marieke GM Weernink, Sarah IM Janus, Janine A van Til, Dennis W Raisch, Jeannette G van Manen, and Maarten J IJzerman. A systematic review to identify the use of preference elicitation methods in healthcare decision making. *Pharmaceutical medicine*, 28(4):175–185, 2014.
- Stephen E Wright, Belle M Sigal, and A John Bailer. Workweek optimization of experimental designs: exact designs for variable sampling costs. *Journal of Agricultural, Biological and Environmental Statistics*, 15(4):491–509, 2010.
- Zhibing Zhao and Lirong Xia. Composite marginal likelihood methods for random utility models. In *Proceedings of the 35th International Conference on Machine Learning*, Stockholm, Sweden, 2018.