

Multi-Winner Social Choice with Incomplete Preferences

Tyler Lu

University of Toronto
 Department of Computer Science
 tl@cs.toronto.edu

Craig Boutilier

University of Toronto
 Department of Computer Science
 cebly@cs.toronto.edu

Abstract

Multi-winner social choice considers the problem of selecting a slate of K options to realize some social objective. It has found application in the construction of political legislatures and committees, product recommendation, and related problems, and has recently attracted attention from a computational perspective. We address the multi-winner problem when facing *incomplete voter preferences*, using the notion of minimax regret to determine a *robust* slate of options in the presence of preference uncertainty. We analyze the complexity of this problem and develop new exact and greedy robust optimization algorithms for its solution. Using these techniques, we also develop preference elicitation heuristics which, in practice, allow us to find near-optimal slates with considerable savings in the preference information required vis-à-vis complete votes.

1 Introduction

Social choice deals with the aggregation of individual preferences over a set of alternatives in order to select a suitable consensus option. While considerable research in computational social choice has focused on *single-winner* choice problems, only recently has much attention been paid to the algorithmic aspects of *multi-winner choice* [Potthoff and Brams, 1998; Procaccia *et al.*, 2008; Meir *et al.*, 2008; Lu and Boutilier, 2011a; Skowron *et al.*, 2013; Betzler *et al.*, 2011; Lucier and Oren, 2012]. Multi-winner problems are of critical importance in many settings, the classic being election of legislatures and committees using proportional representation [Chamberlin and Courant, 1983; Monroe, 1995]. However, they are also appropriate for resource allocation, product recommendation, and consumer segmentation problems where: a limited number of options can be offered; budget constraints preclude personalized recommendations; or budget can be traded off against group satisfaction [Kleinberg *et al.*, 2004; Lu and Boutilier, 2011a].

As with any social choice problem, practical deployment of multi-winner choice faces the challenge of *preference elicitation*: requiring individuals to specify a full preference ranking over the option space poses high, and often unnecessary, costs with respect to cognitive burden, communication and privacy. Furthermore, in settings such as consumer choice, preference information is often limited to revealed choice data (e.g., which products have been purchased) or surveys, in which

case complete preference information may simply be unavailable. Recent work in *vote elicitation* has begun to address this problem [Kalech *et al.*, 2011; Lu and Boutilier, 2011b; 2011c; Oren *et al.*, 2013] by designing schemes that effectively elicit only *relevant* preference information; and constructing methods for winner selection with only partial preferences. In this work, we address the same problem in the context of multi-winner choice, addressing two problems: (a) how to select a bounded slate of (at most) K options given partial preference information; and (b) how to elicit additional preference information incrementally to improve the quality of the slate.

Specifically, we consider the problem of *proportional representation* using positional scoring functions [Chamberlin and Courant, 1983] under partial preferences. We adapt the *minimax regret*-based approach of Lu and Boutilier [2011b] for single-winner problems to the problem of choosing a slate of K options, we show that *robust optimization* is NP-hard, but that *robust greedy optimization* (i.e., selecting the next best option for the partial slate) can be computed effectively—indeed, the model generalizes regret-based optimization for single-winner voting. We then address the problem of elicitation, adapting the *current solution heuristic* [Lu and Boutilier, 2011b] to multi-winner choice. Empirical results show that very good slates of options can be found with relatively little preference information.

2 Background

We begin with a brief review of relevant concepts from social choice, outline the Chamberlin-Courant proportional representation scheme, and briefly describe our model of partial preference information.

Preference Profiles. We assume a set of *agents* (or *voters*) $N = \{1, \dots, n\}$ and a set of *alternatives* (or options) $A = \{a_1, \dots, a_m\}$. Let Γ be the set of *rankings* (or *votes*) over A (i.e., permutations over A). Alternatives represent any outcome space over which the voters have preferences (e.g., products, elected officials, public projects). Agent i 's preferences are represented by a ranking $v_i \in \Gamma$, where i prefers a to a' , denoted $a \succ_i a'$, if $v_i(a) < v_i(a')$. The collection of votes $\mathbf{v} = (v_1, \dots, v_n) \in \Gamma^n$ is a *preference profile*.

A *positional scoring function (PSF)* $\alpha : \{1, \dots, m\} \mapsto \mathbb{R}_{\geq 0}$ maps ranks onto scores such that $\alpha(1) \geq \dots \geq \alpha(m) \geq 0$. We write $\alpha_i(a) = \alpha(v_i(a))$, which can be interpreted as

a measure of i 's *satisfaction* with option a . In what follows, we focus on the *Borda score*, where $\alpha(i) = m - i$, but most of our results generalize to arbitrary scoring rules.

Proportional Representation. In single-winner social choice, the goal is to select a single option a that reflects some notion of consensus. For example, the *Borda winner* is the option a with the greatest total Borda score $\sum_i \alpha_i(a)$ in profile \mathbf{v} . Our focus here is on the selection of a *set of options* $\bar{a} \subseteq A$ where $|\bar{a}| \leq K$, allowing voters to derive satisfaction from their most preferred candidate. Specifically, given a preference profile \mathbf{v} , we define the *score* of a K -set and the *optimal K -set* as follows:

$$S(\bar{a}, \mathbf{v}) = \sum_{i \in N} \max_{a \in \bar{a}} \alpha_i(a) = \sum_i S_i(\bar{a}), \quad (1)$$

$$\bar{a}_{\mathbf{v}}^* = \operatorname{argmax}_{|\bar{a}| \leq K} S(\bar{a}, \mathbf{v}). \quad (2)$$

(We suppress dependence of S on α since the PSF will be fixed.) When α is the Borda PSF, this corresponds to the *proportional representation scheme* of Chamberlin and Courant [1983]. It can be viewed as a *segmentation problem* [Kleinberg *et al.*, 2004]; and it is also a special case of *budgeted social choice*, specifically, the *limited choice* form of the problem [Lu and Boutilier, 2011a]. More general forms of proportional representation [Monroe, 1995] and budgeted social choice [Lu and Boutilier, 2011a] allow for *assignment functions* that map voters to specific options (e.g., to ensure balanced representation, or budget feasibility); but here we assume that the only constraint is on the number of options selected. Sets of size less than K offer no advantage over those of size K in this case.

Partial Preferences. To represent partial information about the preferences of voters, we let a *partial ranking* for a voter i be a partially ordered set over A , or equivalently (the transitive closure of) a collection of consistent *pairwise comparisons* of the form $a \succ_i a'$. Most natural constraints on preferences, including the responses to natural queries (e.g., pairwise comparison, top- t preferences) can be represented this way. Let p_i be i 's partial ranking. We will simply write $a \succ_i a'$ to indicate that this preference can be inferred from p_i when the partial ranking is clear from context. A *completion* of p_i is any vote v_i that extends p_i . Let $C(p_i)$ denote the set of *completions* of p_i . When preferences are uncertain, we write $\alpha_i(a, v_i)$ to emphasize the dependence of voter i 's score on her vote. A *partial profile* is a collection of partial votes $\mathbf{p} = \langle p_1, \dots, p_n \rangle$. Let $C(\mathbf{p}) = C(p_1) \times \dots \times C(p_n)$ be the set of *completions* of \mathbf{p} .

Related Work. Multi-winner problems of the form Eq. 2 were proposed by Chamberlin and Courant [1983] as a form of *proportional representation (PR)*. Monroe [1995] considered a variant of this model in which representatives on the winning slate must have roughly balanced numbers of voters, requiring the use of assignment functions (so certain voters may not be “represented” by their most preferred option). Lu and Boutilier [2011a] generalize the model further by considering both fixed and unit costs. Computing optimal slates, even in the simple PR/limited choice

model given by Eq. 2, is NP-hard [Lu and Boutilier, 2011a; Meir *et al.*, 2008] (also see parameterized complexity results [Betzler *et al.*, 2011]). But optimal slates in the PR model (and also in Monroe’s scheme) can be computed greedily with a $1 - \frac{1}{e}$ approximation ratio [Lu and Boutilier, 2011a; Skowron *et al.*, 2013]; see [Potthoff and Brams, 1998; Procaccia *et al.*, 2008; Lucier and Oren, 2012] for additional computational results on the problem (including online variants). Segmentation [Kleinberg *et al.*, 2004], facility location [Hajiaghayi *et al.*, 2003] and maximum coverage problems [Cohen and Katzir, 2008] all bear close connection to the general budgeted problem (see [Lu and Boutilier, 2011a]).

While multi-winner problems with partial preferences have not been addressed in the literature to the best of our knowledge, winner determination in standard single-winner elections with partial preferences has received attention, both in the form of *possible and necessary winners* [Konczak and Lang, 2005; Xia and Conitzer, 2008], and using *minimax regret* [Lu and Boutilier, 2011b] to provide quality guarantees. We define the latter notion below and draw heavily on this approach in our work. Recent work on incremental vote elicitation has shown strong promise for reducing the informational requirements of various voting schemes in single-winner problems (including Borda and other PSF-based methods), despite seemingly discouraging worst-case communication complexity results [Conitzer and Sandholm, 2005]. Kalech *et al.* [2011] develop schemes based on the notion of possible and necessary winners. Lu and Boutilier [2011b; 2011c] use minimax regret to determine appropriate queries, an approach we generalize below.

3 Robust Multi-winner Optimization

We first consider the problem of selecting an optimal K -set of options when we have only a partial preference profile \mathbf{p} rather than a complete profile \mathbf{v} . Partial preferences may be common in applications involving limited surveying, partial customer purchase or rating data, or situations in which voters can complete partial ballots. Just as importantly, reasoning with partial profiles is vital for incremental vote elicitation.

Following Lu and Boutilier [2011b], we adopt *minimax regret* as a robustness criterion for making decisions with incomplete information. We first present definitions (for sets of fixed size K) then explain the intuitions:

$$\operatorname{Regret}(\bar{a}, \mathbf{v}) = \max_{|\bar{w}| \leq K} S(\bar{w}, \mathbf{v}) - S(\bar{a}, \mathbf{v}) \quad (3)$$

$$\operatorname{PMR}(\bar{a}, \bar{w}, \mathbf{p}) = \max_{\mathbf{v} \in C(\mathbf{p})} S(\bar{w}, \mathbf{v}) - S(\bar{a}, \mathbf{v}) \quad (4)$$

$$\operatorname{MR}(\bar{a}, \mathbf{p}) = \max_{\mathbf{v} \in C(\mathbf{p})} \operatorname{Regret}(\bar{a}, \mathbf{v}) = \max_{|\bar{w}| \leq K} \operatorname{PMR}(\bar{a}, \bar{w}, \mathbf{p}) \quad (5)$$

$$\operatorname{MMR}(\mathbf{p}) = \min_{|\bar{a}| \leq K} \operatorname{MR}(\bar{a}, \mathbf{p}) \quad (6)$$

$$\bar{a}_{\mathbf{p}}^* \in \operatorname{argmin}_{|\bar{a}| \leq K} \operatorname{MR}(\bar{a}, \mathbf{p}) \quad (7)$$

Given a vote profile \mathbf{v} , $\operatorname{Regret}(\bar{a}, \mathbf{v})$ describes the loss in satisfaction associated with offering set \bar{a} rather than the optimal K -set. Give a partial profile \mathbf{p} , the *pairwise max regret* $\operatorname{PMR}(\bar{a}, \bar{w}, \mathbf{p})$ is the worst-case loss that could be incurred, under all possible realizations of voter preferences, by offering \bar{a} rather than \bar{w} . Note that our definition of PMR does not

impose constraints on set sizes, a fact we exploit below. The *max regret* $MR(\bar{a}, \mathbf{p})$ of set \bar{a} is the worst-case loss relative to the *optimal* K -set under all preference realizations: this bounds the loss associated with \bar{a} given our preference uncertainty. Finally, a *minimax optimal set* $\bar{a}_\mathbf{p}^*$ is one with minimal max regret or *minimax regret* $MMR(\mathbf{p})$.

Observation 1. If $MMR(\mathbf{p}) = 0$, then $\bar{a}_\mathbf{p}^*$ is an optimal slate of alternatives for any $\mathbf{v} \in C(\mathbf{p})$.

3.1 Computing Minimax Optimal Sets

Before discussing computation of minimax regret, we begin with the simpler problem of computing pairwise max regret. From Eqs. 4–6, we see that the regret-optimal slate $\bar{a}_\mathbf{p}^*$ can be determined by first computing $PMR(\bar{a}, \bar{w}, \mathbf{p})$ for all pairs of K -sets \bar{a}, \bar{w} , maximizing over \bar{w} to determine $MR(\bar{a}, \mathbf{p})$, then minimizing over these terms to compute $MMR(\mathbf{p})$. If K is small, then robust optimization is efficient if PMR can be computed effectively, a problem to which we now turn. (We discuss an approach for large K in the next subsection.)

One can show PMR is additively decomposable:

$$\begin{aligned} PMR(\bar{a}, \bar{w}, p_i) &= \max_{v_i \in C(p_i)} S_i(\bar{w}, v_i) - S_i(\bar{a}, v_i) \\ PMR(\bar{a}, \bar{w}, \mathbf{p}) &= \sum_{i \in N} PMR(\bar{a}, \bar{w}, p_i) \end{aligned} \quad (8)$$

Thus we can compute the contributions of each voter i to PMR individually. When i is presented with slate \bar{a} , she will choose her most preferred element form \bar{a} and similarly for slate \bar{w} . Define the *undominated elements* in any set \bar{a} to be:

$$u_i(\bar{a}) = \{a \in \bar{a} : \neg \exists a' \in \bar{a} \text{ s.t. } a' \succ_i a\}.$$

Only undominated elements in a set can be chosen in any completion of i 's preferences. In the ranking $v_i \in C(p_i)$ that maximizes pairwise regret, only one element in \bar{w} will be chosen by i (the most preferred), which defines PMR :

$$PMR(\bar{a}, \bar{w}, p_i) = \max_{w \in u_i(\bar{w})} PMR(u_i(\bar{a}), \{w\}, p_i). \quad (9)$$

Given this, there are two cases to consider in determining the adversarial completion $v_i \in C(p_i)$ that maximizes $PMR(u_i(\bar{a}), \{w\}, p_i)$.

Case 1. Suppose there is an $a \in \bar{a}$ such that $a \succ_i w$. This means in there is no completion in which i would choose w , so PMR is negative. Maximizing pairwise regret requires reducing the ‘‘gap’’ between the most preferred $a^* \in \bar{a}$ and w . The only options that must lie between the most preferred a and w are undominated elements $u_i(\bar{a})$ of \bar{a} that dominate w , or those b known to lie between such an a and w . Define

$$\begin{aligned} u_i(\bar{a})_w^+ &= \{a \in u_i(\bar{a}) : a \succ_i w\}, \\ B_i(\bar{a}, w) &= \{b \in A : \exists a \in u_i(\bar{a})_w^+, a \succ_i b \succ_i w\}. \end{aligned}$$

$B_i(\bar{a}, w)$ includes all options that must lie between the best $a \in \bar{a}$ and w (the specific choice or placement of the elements in these two sets has no impact on PMR). Every other option can consistently be ordered above the best a or below

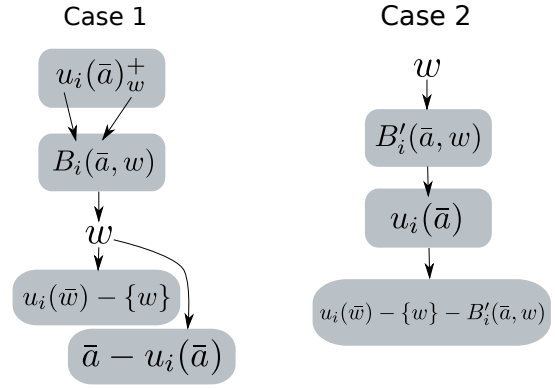


Figure 1: Adversarial completions of p_i regret $PMR(u_i(\bar{a}), \{w\}, p_i)$, Cases 1 and 2.

w depending on constraints in p_i .¹ Thus we have:

$$PMR(u_i(\bar{a}), \{w\}, p_i) = -|u_i(\bar{a})_w^+| - |B_i(\bar{a}, w)|. \quad (10)$$

See Fig. 1(case 1) for an illustration.

Case 2. Now suppose that for voter i , no element in $a \in \bar{a}$ is known to be preferred to w . If $w \in u_i(\bar{a})$ then $PMR(u_i(\bar{a}), \{w\}, p_i) = 0$, since any adversarial completion can place w above all items in $u_i(\bar{a}) \setminus \{w\}$ (otherwise regret would be negative). Otherwise the desired completion must maximize the gap between w and any item in $u_i(\bar{a})$. The following options can be placed between w and \bar{a} :

$$B'_i(\bar{a}, w) = \{b \in A \setminus \bar{a} : b \not\succeq_i w \text{ and } \forall a \in u_i(\bar{a}), a \not\succeq_i b\}.$$

The relative ordering of these items does not impact regret. With $B'_i(\bar{a}, w)$ placed below w , some item from $u_i(\bar{a})$ must lie immediately below the last element of this set (becoming the most preferred $a \in \bar{a}$). Thus, we have:

$$PMR(u_i(\bar{a}), \{w\}, p_i) = \begin{cases} 1 + |B'_i(\bar{a}, w)| & \text{if } w \notin u_i(\bar{a}), \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

See Fig. 1(case 2) for an illustration (where $w \notin u_i(\bar{a})$).

In both cases, the undominated sets $u_i(\bar{a})$ and $u_i(\bar{w})$ can be computed in $O(K^2)$ time. In case 1, $u_i(\bar{a})_w^+$ can be computed in $O(K)$ time once $u_i(\bar{a})$ is known, and $B_i(\bar{a}, w)$ can be computed in $O(mK)$ time by checking if each $b \in A$ satisfies the constraints w.r.t. $u_i(\bar{a})_w^+$ and w . For case 2, $B'_i(\bar{a}, w)$ can be found in time $O(mK)$ by checking each $b \in A$ with w and the options in $u_i(\bar{a})$. Using Eqs. 9 and 8, $PMR(\bar{a}, \bar{w}, \mathbf{p})$ can be computed in $O(n(K^2 + K(K + mK))) = O(nmK^2)$ time. Note that for $K = 1$, the approach is identical to PMR computation proposed by Lu and Boutilier [2011b]. Putting this together we have:

Theorem 2. $PMR(\bar{a}, \bar{w}, \mathbf{p})$ is given by:

$$\sum_{i \in N} \max_{w \in u_i(\bar{w})} \begin{cases} -|u_i(\bar{a})_w^+| - |B_i(\bar{a}, w)| & \text{if } \exists a \in \bar{a} : a \succ_i w, \\ 1 + |B'_i(\bar{a}, w)| & \text{o.w., and } w \notin u_i(\bar{a}), \\ 0 & \text{o.w.,} \end{cases}$$

and is computable in $O(nmK^2)$ time.

¹For ‘‘nonlinear’’ scoring rules, where the score difference for two options depends not just on relative, but also absolute rank position, placement of options above or below a and w requires more care, but is straightforward in most cases.

The minimax optimal slate \bar{a}_p^* can be computed by computing max regret $MR(\bar{a}, \mathbf{p})$ for each size K slate \bar{a} (and selecting the slate \bar{a}_p^* that minimizes regret); and in turn $MR(\bar{a}, \mathbf{p})$ can be computed by determining its pairwise max regret for each size K witness set \bar{w} . Hence:

Proposition 3. *The minimax regret optimal slate \bar{a}_p^* can be found in time $O(nm^{2K+1}K^2)$.*

The additive decomposability of PMR has the nice computational consequence that, during the course of incremental elicitation (see Sec. 4), one need only update the contributions to PMR of those agents who have their partial preferences updated by responding to a query.

For slates of small bounded size K , enumeration of option sets may be practical—with bounded K , slates can be computed in polynomial time (in n, m). However, since this form of proportional representation and budgeted social choice is an NP-hard optimization problem [Lu and Boutilier, 2011b] (as are related forms [Procaccia *et al.*, 2008]), finding the minimax optimal slate is also NP-hard (simply let \mathbf{p} be a full preference profile). Indeed, even simply computing $MR(\bar{a}, \mathbf{p})$ is NP-hard:

Theorem 4. *Given threshold $r \geq 0$, partial profile \mathbf{p} , set size K , and set \bar{a} of size at most K , deciding if $MR(\bar{a}, \mathbf{p}) \geq r$ (i.e., does some set \bar{w} of size at most K satisfy $PMR(\bar{a}, \bar{w}, \mathbf{p}) \geq r$) is NP-hard.*

Proof. This is easy to see using a simple reduction from the limited choice (LC) problem (with Borda scoring), which is NP-hard [Lu and Boutilier, 2011a]. Given any LC instance with budget K and complete profile \mathbf{v} over m items, we transform it into a partial profile \mathbf{p} with $m + K$ items (the m original items plus K “dummy” items). Set \bar{a} to be the dummy items. Define each partial p_i to be identical to v_i on the top- m ranked items, while the rest of the ranking is unspecified. Computing whether some slate \bar{w} of size at most K has $PMR(\bar{a}, \bar{w}, \mathbf{p}) \geq r$ can then be used to determine if there is a slate \bar{w}' with score above a threshold in the LC instance. \square

3.2 A Greedy Algorithm

Due to the lack of a (general) efficient algorithm to compute \bar{a}_p^* , we develop a heuristic approach that will be practical for large K . It turns out that a relatively simple greedy optimization procedure can be used for this purpose. To this end, define the following problem, which we call the *additional option problem*: assume a partial profile \mathbf{p} and a fixed set \bar{a} of $k - 1$ options; if one can add a k th option to the set, which next option minimizes maximum regret under the PR/limited choice model? We define this problem in the obvious way:

$$\begin{aligned} PMR(a, w, \mathbf{p}|\bar{a}) &= PMR(\bar{a} \cup \{w\}, \bar{a} \cup \{a\}, \mathbf{v}) \\ MR(a, \mathbf{p}|\bar{a}) &= \max_{w \in A} PMR(a, w, \mathbf{p}|\bar{a}) \\ MMR(\mathbf{p}|\bar{a}) &= \min_{a \in A} MR(a, \mathbf{p}|\bar{a}) \\ a_{\bar{a}, \mathbf{p}}^* &\in \operatorname{argmin}_{a \in A} MR(a, \mathbf{p}|\bar{a}). \end{aligned} \quad (12) \quad (13)$$

(Note that setting $k = 1$ gives the (single-winner) robust voting problem of [Lu and Boutilier, 2011b].)

The additional option problem can be solved in polynomial time. We can explicitly compute the pairwise max regret $PMR(a, w, \mathbf{p}|\bar{a})$ of all m^2 pairs of alternatives (a, w) (where a is a proposed additional option and w is an adversarial witness), using intuitions very similar to those above (as we discuss below). Note that while we can apply our previous algorithm for finding the pairwise max regret for arbitrary pairs of slates (\bar{a}, \bar{w}) , the algorithm and analysis for $PMR(a, w, \mathbf{p}|\bar{a})$ that we provide below provides a factor k speedup. With PMR in hand, we can readily determine minimax regret using Eqs. 12 and 13. We need only show that PMR can be computed in polynomial time. As above, we can compute each voter’s i ’s contribution $PMR(a, w, p_i|\bar{a})$ independently. We again consider two cases.

Case 1. If $a \succ_i w$, then $PMR(a, w, p_i|\bar{a}) \leq 0$ since adding w to \bar{a} cannot improve i ’s score any more than adding a . Assuming $a \succ_i w$, if there is some $b \in \bar{a}$ such that $a \not\succeq_i b$, then b can be ordered over a and $PMR(a, w, p_i|\bar{a}) = 0$. However, if $a \succ_i b$ for all $b \in \bar{a}$, regret must be negative. $PMR(a, w, p_i|\bar{a})$ is then maximized (or negative regret minimized) by placing as few options as possible between a and the best element of $\bar{a} \cup \{w\}$. For any $a \succ_i b$, define

$$T_i(a, b) = \{b' : a \succ_i b' \succ_i b\}.$$

Then regret is maximized by ordering the options in $u_i(\bar{a} \cup \{w\})$ such that the element with the fewest possible options between it and a is ranked first. This gives:

$$PMR(a, w, p_i|\bar{a}) = \max_{b \in u_i(\bar{a} \cup \{w\})} -|T_i(a, b)| - 1.$$

Case 2. If $a \not\succeq_i w$, then $PMR(a, w, p_i|\bar{a}) \geq 0$. In this case, if there is some $b \in \bar{a}$ such that $b \succ_i w$, then w can never be selected; but since w can be ordered over a , $PMR(a, w, p_i|\bar{a}) = 0$. However, if there is no $b \in \bar{a}$ with $b \succ_i w$, then regret is maximized by maximizing the gap between w and the best element of $\bar{a} \cup \{a\}$. In particular, the options $B'_i(\bar{a} \cup \{a\}, w)$, as defined above, can all be ordered between w and the best such option. This gives us:

$$PMR(a, w, p_i|\bar{a}) = |B'_i(\bar{a} \cup \{a\}, w)| + 1.$$

Taken together, this shows that $PMR(a, w, p_i|\bar{a})$ for the additional option problem can be computed in polynomial time. This gives rise to a very simple *greedy algorithm* for approximating a minimax optimal K -set: starting with the empty slate $\bar{a}_0 = \emptyset$, at each of iteration $k \leq K$ we add option $a_k^* = a_{\bar{a}_{k-1}, \mathbf{p}}^*$, i.e., the option with least max regret given the prior items, to slate \bar{a}_{k-1} . The method is detailed in Algorithm 1. While this algorithm comes with no strong approximation guarantees (but see below), we show in Sec. 5 that it works extremely well in practice.

In terms of run time, Case 1 takes $O(mk)$ time (at the k th iteration) and Case 2, as discussed previously, takes $O(k^2 + mk) = O(mk)$ time. Computing pairwise max regret for all pairs (a, w) , across all agents, and finding the next best item $a_{\bar{a}_k, \mathbf{p}}^*$ for each of the K spots on slate \bar{a} results in a total running time of $O(nm^3K^2)$.

There are two reasons the greedy algorithm cannot guarantee that we find the minimax optimal slate. The first is

Algorithm 1 Greedy algorithm

```
1:  $\bar{a} \leftarrow \emptyset$ 
2: for  $k = 1$  to  $K$  do
3:    $\text{MMR} \leftarrow \infty$ 
4:   for  $a \in A$  do
5:      $\text{MR} \leftarrow -\infty$ 
6:     for  $w \in A : w \neq a$  do
7:       for  $i \in N$  do
8:         if  $a \succ_i w$  then
9:            $\text{PMR} \leftarrow \text{PMR} + \max_{b \in u_i(\bar{a} \cup \{w\})} -|T_i(a, b)| - 1$ 
10:        else
11:           $\text{PMR} \leftarrow \text{PMR} + |B'_i(\bar{a} \cup \{a\}, w)| + 1$ 
12:        end if
13:      end for
14:      if  $\text{PMR} > \text{MR}$  then
15:         $\text{MR} \leftarrow \text{PMR}$ 
16:      end if
17:    end for
18:    if  $\text{MR} < \text{MMR}$  then
19:       $\text{MMR} \leftarrow \text{MR}$ 
20:       $a^* \leftarrow a$ 
21:    end if
22:  end for
23:   $\bar{a} \leftarrow \bar{a} \cup \{a^*\}$ .
24: end for
```

unrelated to preference uncertainty: *even with complete preference information*, the greedy algorithm is unable to provide an optimal K -slate. In other words, it can produce a slate that has positive max regret. However, the greedy algorithm does provide a $1 - \frac{1}{e}$ approximation in the full information setting. It is not hard to see that if we have sufficient information to make the “optimal greedy choice” at each stage, then the regret-based approach will correspond to the exact greedy algorithm described by Lu and Boutilier [2011a]:

Proposition 5. *If $\text{MR}(a_k^*, \mathbf{p}|\bar{a}_{k-1}) = 0$ for all $k \leq K$, then the greedy-MMR set \bar{a}_K is identical to the set produced by the (full-information) greedy algorithm given any $\mathbf{v} \in C(\mathbf{p})$.*

If the last element added has non-zero max-regret, we are assured that true minimax regret is also nonzero:

Observation 6. *If $\text{MR}(a_K^*, \mathbf{p}|\bar{a}_{K-1}) > 0$, then $\text{MMR}(\mathbf{p}) > 0$.*

However, we cannot be sure that if only the last element has zero regret that we have found the greedy-optimal slate. But even if the minimax-optimal option a_i^* does not have zero max regret, we can still obtain bounds on the quality of the solution. Based on known results for “approximate” greedy optimization [Goundan and Schulz, 2007], we obtain:

Proposition 7. *Let \mathbf{v} be an (unknown) complete preference profile. Let m_k be any lower bound on the marginal value of the k th item added to the slate by the full information greedy algorithm. That is, $m_k \leq S(\bar{a}_{k-1} \cup \{a_k^*\}, \mathbf{v}) - S(\bar{a}_{k-1}, \mathbf{v})$ where \bar{a}_{k-1} consists of the first $k-1$ greedily selected items. If $\text{MR}(a_k^*, \mathbf{p}|\bar{a}_{k-1}) \leq m_k / (1 - \alpha)$ for all $k \leq K$, for some $\alpha \geq 1$, then the greedy-MMR set \bar{a}_K is within a factor of $1 - \frac{1}{e^{1/\alpha}}$ of the optimal (full-information) option slate.*

4 Preference Elicitation

We now turn our attention to the question of incremental elicitation of voter preferences. With a partial preference profile

\mathbf{p} , we cannot guarantee that an optimal slate can be obtained (regardless of whether we resort to greedy or exact optimization): specifically, if $\text{MMR}(\mathbf{p}) > 0$, no slate can be guaranteed to be optimal. To improve the quality of the slate, further information must be elicited from one or more voters. Our goal is to reduce *relevant* uncertainty, i.e., find those queries that have the greatest potential to reduce minimax regret. To do this we adopt a general technique known as the *current solution strategy* (CSS). This has been used effectively in both (single-agent) robust optimization [Boutilier *et al.*, 2006; Brazianas and Boutilier, 2010], voting [Lu and Boutilier, 2011b] and stable matching [Drummond and Boutilier, 2013] settings, and relies on eliciting information that helps assess the relative degree of preference between the minimax optimal solution and the adversarial witness. We focus here on pairwise comparison queries, in which some voter i is asked whether $a \succ_i a'$; but the basic principles can be applied to other forms of queries (e.g., top- t queries, where a voter is asked to reveal their t th-most preferred alternative if they have already provided their top $t-1$).

In our model, the use of CSS differs depending on whether we are using the greedy heuristic or optimal MMR computation to determine K -sets. We present our approach in the context of greedy MMR computation since it is the more practical method for problems involving slates of reasonable size. The general principles can be readily adapted to optimal MMR computation as well. Our elicitation scheme works by using the greedy algorithm to compute an (approximately) minimax optimal slate $\bar{a}_{\mathbf{p}}^* = \langle a_1^*, \dots, a_K^* \rangle$ given the current partial profile \mathbf{p} . If $\text{MR}(a_K^*, \mathbf{p}|\bar{a}_{K-1}^*) = 0$, we treat this as an (approximately) minimax optimal slate and stop. Otherwise, we know that $\text{MMR}(\mathbf{p}) > 0$, so we select a voter i and pairwise comparison query $a \succ_i a'$ with the greatest potential to reduce $\text{MR}(a_K^*, \mathbf{p}|\bar{a}_{K-1}^*)$, using CSS. Let a_K^* be the last item added to the slate and w_K be the witness option (i.e., where $\text{MR}(a_K^*, \mathbf{p}|\bar{a}_{K-1}^*) = \text{PMR}(a_K^*, w_K, \mathbf{p}|\bar{a}_{K-1}^*)$). CSS identifies the appropriate query (and its potential) for a particular voter i based on several specific cases/sub-cases.

Case 1: Suppose $a_K^* \succ_i w_K$. Then i 's contribution PMR_i to $\text{PMR}(a_K^*, w_K, \mathbf{p}|\bar{a}_{K-1}^*)$ must be $\text{PMR}_i \leq 0$. If $\text{PMR}_i = 0$, then either: (i) a_K^* is dominated in i 's partial order p_i by some $a_j \in \bar{a}_{K-1}^*$, or (ii) a_K^* is not dominated by any such a_j . In case (i), no query can reduce $\text{MR}(a_K^*, \mathbf{p}|\bar{a}_{K-1}^*)$ since voter i would never select either of a_K^* or w_K given the rest of the slate \bar{a}_{K-1}^* , so *no query* is asked of i . In case (ii), the adversary can set $\text{PMR}_i = 0$ by ordering some option $a_j \in \bar{a}_{K-1}^*$ over a_K^* (if no such option were possible, PMR_i would have to be negative). In this case, any query that prevents a_j from being orderable above a_K^* can reduce PMR_i (by making it negative). Specifically, any query of the form $b \succ_i c$ for $b \in a_K^* \cup \{a : a_K^* \succ_i a\}$ and $c \in a_j \cup \{a : a \succ_i a_j\}$ will suffice. Since the degree of PMR_i is determined by the relationship of a_K^* to w_K and *not the gap* with a_j , we choose query $a_K^* \succ_i a_j$ since it is implied by any other.

If $\text{PMR}_i < 0$, then (iii) a_K^* dominates each $a_j \in \bar{a}_{K-1}^*$ as well as w_K . Queries that can extend the advantage over the best item of $w^* \in u_i(\{w_K\} \cup \bar{a}_{K-1}^*)$ (i.e., make PMR_i even smaller) take two forms: learning that $a_K^* \succ_i c$ for some ancestor c of w^* (since we do not know w^* , we choose an




Figure 2: Performance of Greedy/CSS on (a) *Sushi* (20 trials), (b) *Irish* (20 trials) and (c) *Mallows* datasets (20 trials).

arbitrary item in $u_i(\{w_K\} \cup \bar{a}_{K-1}^*)$; or learning that $d \succ_i w^*$ for some descendant d of a_K^* . The query with the greatest

potential is that with the largest number of descendants (resp. ancestors) lying between it and c (resp. d) in p_i .

Case 2: Suppose $a_K^* \not\succeq_i w_K$. Then i 's contribution PMR_i to $PMR(a_K^*, w_K, \mathbf{p}|\bar{a}_{K-1}^*)$ must be $PMR_i \geq 0$. (i) If $PMR_i = 0$, then we must have w_K dominated by some $a_j \in \bar{a}_{K-1}^*$. In this case we ask no query. (ii) If $PMR_i > 0$, then w_K is not dominated by any $a_j \in \bar{a}_{K-1}^*$ (i.e., by none of the K items). Then regret can be reduced only by asking a query that removes elements from the set $B'_i(\bar{a}_{K-1}^*, w_K)$ by either placing w_K (or one of its ancestors) below some a_j (it is only necessary to consider those in $u_i(\bar{a}_{K-1}^*)$); by placing some $a_j \in u_i(\bar{a}_{K-1}^*)$ (or one of its descendants) above w_K ; or placing some element that is incomparable to both a_K^* and w_K either above w_K (hence placing its ancestors above as well) or below a_K^* (hence placing its descendants below as well). In the case that w_K dominates all of \bar{a}_{K-1}^* , one can ask queries that either (a) move a descendant d of w_K , where d is not an ancestor of some $a_j \in \bar{a}_{K-1}^*$, below such an a_j ; or (b) move an ancestor c of some a_j , where c is not a descendant of w_K , above w_K . The potential of a query to reduce PMR_i is measured by the number of elements it removes from the set $B'_i(\bar{a}_{K-1}^*, w_K)$.

We note that if $PMR(a_K^*, w_K, \mathbf{p}|\bar{a}_{K-1}^*) > 0$, then one of the query-generating cases above must hold for at least one voter. As a consequence, CSS cannot “stall” as long as the last item added to the slate has non-zero max regret.

5 Experimental Evaluation

We present a set of experiments designed to test the ability of greedy slate optimization method described in Sec. 3.2, coupled with the CSS elicitation strategy, to find good (or even optimal) slates of options with few voter queries. We evaluate the approach on two real datasets as well as on more systematically generated random data. The *Sushi* data set [Kamishima *et al.*, 2005] consists of 5000 complete user preference rankings over 10 varieties of sushi. We consider aggregation problems (or *elections*) of 100 users, drawing full preference profiles at random from this set. The *Irish* dataset consists of partial (top- t) preferences of voters from the 2002 Dublin West constituency elections, involving nine candidates and 29,988 voters; we consider elections of 100 voters by drawing random profiles from among the 3800 full rankings in the dataset. In the *Sushi* experiment we set slate size to be $K = 3$ and in *Irish*, $K = 4$. To test our methods more systematically, we generate random profiles of 100 voter rankings over 10 items, using the *Mallows distribution* to generate the preferences of each voter [Mallows, 1957; Marden, 1995]. The Mallows model is quite flexible: the dispersion parameter $\phi \in (0, 1]$ controls how concentrated preferences are: when $\phi = 1$, we obtain uniformly random preferences, or *impartial culture*, while we approach a distribution where all mass is placed on a single-ranking (i.e., all voters have identical rankings) as $\phi \rightarrow 0$. In the *Mallows* experiments, we set $K = 3$ and analyze elicitation performance as we vary ϕ .

Experimentally, each instance consists of a full profile. We start with no vote information then, using CSS to generate queries, elicit pairwise comparisons from voters (who respond accurately based on their underlying preference). After

each query or round, we use the greedy algorithm to compute an (approximately) optimal K -slate. Elicitation terminates once the conditional max regret (CondMR) of the K th item added to the slate, $MR(a_K^*, \mathbf{p} | \bar{a}_{K-1}^*)$, is zero. We also compare this to the use of exact minimax regret computation at each iteration to determine the truly optimal K -slate.

Results in Fig. 2 (a) and (b) show performance for the *Sushi* and *Irish* data. The plots show, for the slate produced by Greedy after each query, its conditional max regret (CondMR), its true max regret and two dotted lines of one standard deviation above and below its max regret (results are averaged over 20 randomly drawn profiles). We also show the performance of CSS when exact MMR is computed and the optimal slate is generated at each iteration. We also compare CSS with a baseline random strategy that randomly picks a voter and pairwise query (ensuring this query is not implied by that voter’s partial ranking), using Greedy to compute a slate at each round and measuring its max regret. Dotted lines above and below indicate a range of one standard deviation. We see that CSS works very well: it finds a slate with zero max regret “per voter” with only about 20 queries per user in *Sushi* (resp., 15 in *Irish*), even with the relatively large ratio of K to m in each setting (30% and 44%, resp.). CSS also reaches near-zero regret in about 10 queries (resp., 8) per user; thus, its *anytime profile* is very encouraging for settings where approximately optimal solutions that reduce elicitation burden are permissible. Note that the true regret may be zero even if max regret is not. We contrast the number of queries needed by CSS with the demands of complete sorting to provide a full ranking, which requires $O(m \log(m))$ pairwise comparisons using methods with good average case performance, or equivalently 34 (resp., 29) queries. Random requires 25 (resp., 22) queries per user to reach zero regret, and has a much worse anytime profile.

Notice that the greedy algorithm itself works extremely well: it almost always finds the minimax optimal slate—the MR Greedy and MMR curves coincide almost exactly—and in the rare cases that it does not, Greedy MR is very close to true MMR. MR may not decrease monotonically, as preference updates may “mislead” Greedy into choosing an inferior slate (by contrast, true MMR is non-increasing). CondMR is also a good proxy for true max regret: in *Sushi*, the per-voter difference is at most 0.41 and in *Irish* at most 0.24. Thus, CondMR—which can be computed efficiently—is an excellent surrogate for MR—which is NP-hard—as a quality measure and a stopping criterion for elicitation.

Mallows results in Fig. 2 (c) show how the same quantities change as a function of the total number of queries, for different dispersion values ϕ . The results show that, unsurprisingly, more concentrated preference distributions (smaller ϕ) require fewer queries to find good slates. This is consistent with observations in single winner voting [Lu and Boutilier, 2011b].

Table 1 shows wall clock runtimes of Greedy for different values of m (candidates) and K (slate size) on a 3.0GHz Intel Xeon processor. Results are averaged over “complete” CSS elicitation runs (i.e., elicitation proceeds until CondMR is zero), on random profiles of $n = 40$ voters drawn from a *Mallows* distribution with $\phi = .7$. Average runtime increases sig-

m	$K = 2$	$K = 3$	$K = 4$	$K = 6$	$K = 8$
10	0.015	0.020	0.023	0.028	0.033
20	0.105	0.152	0.194	0.275	0.345
30	0.342	0.508	0.642	0.987	1.282
50	1.577	2.042	2.247	4.439	6.344

Table 1: Avg. Greedy runtime (sec.), on random *Mallows* profiles.

nificantly with the number of candidates m , but less dramatically with K . This is consistent with the “quadratic in K ” and “cubic in m ” computational results above. Still, Greedy is very practical, taking only 6.3s. to find optimal slates for $m = 50$ and $K = 8$.

6 Conclusion

We have provided a new model for analyzing the quality of slates of options with incomplete preference profiles, and developed algorithms for the exact and greedy optimization of K -slates with minimax regret. The greedy method has been shown to be quite practical. We also adapted the CSS elicitation heuristic to this setting and demonstrated that—when coupled with the greedy method for producing slate, and the efficiently computable conditional max-regret criterion as a proxy for max regret—it finds optimal slates while requiring relatively little preference information from voters. Moreover, it has an excellent anytime profile, finding slates with very low max regret very quickly, and important property in low-stakes domains where decision quality may be sacrificed for increased informational, cognitive, and communication efficiency or privacy.

This work raises a number of interesting questions for future research. Developing regret-based methods for additional scoring/voting rules and other multi-winner criteria would be of significant value. Theoretical analysis of the communication complexity for multi-winner social choice would advance our understanding of elicitation methods. Finally, developing elicitation and optimization methods that exploit preference distributions could further reduce information requirements.

Acknowledgements: Lu was supported by a Google Ph.D. Fellowship. Lu and Boutilier acknowledge the support of the the Natural Sciences and Engineering Research Council (NSERC).

References

- [Betzler *et al.*, 2011] Nadja Betzler, Arkadii Slinko, and Johannes Uhlmann. On the computation of fully proportional representation. SSRN: 1952497, 2011.
- [Boutilier *et al.*, 2006] Craig Boutilier, Relu Patrascu, Pascal Poupard, and Dale Schuurmans. Constraint-based optimization and utility elicitation using the minimax decision criterion. *Artificial Intelligence*, 170(8–9):686–713, 2006.
- [Braziunas and Boutilier, 2010] Darius Braziunas and Craig Boutilier. Assessing regret-based preference elicitation with the UTPREF recommendation system. In *Proceed-*

- ings of the Eleventh ACM Conference on Electronic Commerce (EC'10), pages 219–228, Cambridge, MA, 2010.
- [Chamberlin and Courant, 1983] John R. Chamberlin and Paul N. Courant. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *The American Political Science Review*, 77(3):718–733, 1983.
- [Cohen and Katzir, 2008] Reuven Cohen and Liran Katzir. The generalized maximum coverage problem. *Inf. Process. Lett.*, 108(1):15–22, 2008.
- [Conitzer and Sandholm, 2005] Vincent Conitzer and Thomas Sandholm. Communication complexity of common voting rules. In *Proceedings of the Sixth ACM Conference on Electronic Commerce (EC'05)*, pages 78–87, Vancouver, 2005.
- [Drummond and Boutilier, 2013] Joanna Drummond and Craig Boutilier. Elicitation and approximately stable matching with partial preferences. In *Proceedings of the Twenty-third International Joint Conference on Artificial Intelligence (IJCAI-13)*, Beijing, 2013. to appear.
- [Goundan and Schulz, 2007] Pranava R. Goundan and Andreas S. Schulz. Revisiting the greedy approach to submodular set function maximization. *Optimization Online*, 2007.
- [Hajiaghayi et al., 2003] Mohammad Taghi Hajiaghayi, Mohammad Mahdian, and Vahab S. Mirrokni. The facility location problem with general cost functions. *Networks*, 42(1):42–47, 2003.
- [Kalech et al., 2011] Meir Kalech, Sarit Kraus, Gal A. Kaminka, and Claudia V. Goldman. Practical voting rules with partial information. *Journal of Autonomous Agents and Multi-Agent Systems*, 22(1):151–182, 2011.
- [Kamishima et al., 2005] Toshihiro Kamishima, Hideto Kazawa, and Shotaro Akaho. Supervised ordering: An empirical survey. In *IEEE International Conference on Data Mining*, pages 673–676, Los Alamitos, CA, USA, 2005. IEEE Computer Society.
- [Kleinberg et al., 2004] Jon Kleinberg, Christos Papadimitriou, and Prabhakar Raghavan. Segmentation problems. *Journal of the ACM*, 51:263–280, 2004.
- [Konczak and Lang, 2005] Kathrin Konczak and Jérôme Lang. Voting procedures with incomplete preferences. In *IJCAI-05 Workshop on Advances in Preference Handling*, pages 124–129, Edinburgh, 2005.
- [Lu and Boutilier, 2011a] Tyler Lu and Craig Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the Twenty-second International Joint Conference on Artificial Intelligence (IJCAI-11)*, pages 280–286, Barcelona, 2011.
- [Lu and Boutilier, 2011b] 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-11)*, pages 287–293, Barcelona, 2011.
- [Lu and Boutilier, 2011c] Tyler Lu and Craig Boutilier. Vote elicitation with probabilistic preference models: Empirical estimation and cost tradeoffs. In *Proceedings of the Second International Conference on Algorithmic Decision Theory (ADT-11)*, pages 135–149, Piscataway, NJ, 2011.
- [Lucier and Oren, 2012] Brendan Lucier and Joel Oren. On-line (budgeted) social choice. In *Proceedings of the Fourth International Workshop on Computational Social Choice (COMSOC-2012)*, Kraków, Poland, 2012.
- [Mallows, 1957] Colin L. Mallows. Non-null ranking models. *Biometrika*, 44:114–130, 1957.
- [Marden, 1995] John I. Marden. *Analyzing and Modeling Rank Data*. Chapman and Hall, London, 1995.
- [Meir et al., 2008] Reshef Meir, Ariel D. Procaccia, Jeffrey S. Rosenschein, and Aviv Zohar. Complexity of strategic behavior in multi-winner elections. *Journal of Artificial Intelligence Research (JAIR)*, 33:149–178, 2008.
- [Monroe, 1995] Burt L. Monroe. Fully proportional representation. *The American Political Science Review*, 89(4):925–940, 1995.
- [Oren et al., 2013] Joel Oren, Yuval Filmus, and Craig Boutilier. Efficient vote elicitation under candidate uncertainty. In *Proceedings of the Twenty-third International Joint Conference on Artificial Intelligence (IJCAI-13)*, Beijing, 2013. to appear.
- [Potthoff and Brams, 1998] Richard F. Potthoff and Steven J. Brams. Proportional representation: Broadening the options. *Journal of Theoretical Politics*, 10(2):147–178, 1998.
- [Procaccia et al., 2008] Ariel D. Procaccia, Jeffrey S. Rosenschein, and Aviv Zohar. On the complexity of achieving proportional representation. *Social Choice and Welfare*, 30:353–362, 2008.
- [Skowron et al., 2013] Piotr Skowron, Piotr Faliszewski, and Arkadii Slinko. Achieving proportional representation is easy in practice. In *Proceedings of the Twelfth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-13)*, St. Paul, MN, 2013. To appear.
- [Xia and Conitzer, 2008] Lirong Xia and Vincent Conitzer. Determining possible and necessary winners under common voting rules given partial orders. In *Proceedings of the Twenty-third AAAI Conference on Artificial Intelligence (AAAI-08)*, pages 202–207, Chicago, 2008.