

Budgeted Social Choice: From Consensus to Personalized Decision Making

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

We develop a general framework for social choice problems in which a *limited* number of alternatives can be recommended to an agent population. In our *budgeted social choice* model, this limit is determined by a budget, capturing problems that arise naturally in a variety of contexts, and spanning the continuum from pure consensus decision making (i.e., standard social choice) to fully personalized recommendation. Our approach applies a form of *segmentation* to social choice problems—requiring the selection of diverse options tailored to different agent types—and generalizes certain multi-winner election schemes. We show that standard rank aggregation methods perform poorly, and that optimization in our model is NP-complete; but we develop fast greedy algorithms with some theoretical guarantees. Experiments on real-world datasets demonstrate the effectiveness of our algorithms.

1 Introduction

The ease with which users now rate, compare or rank options has allowed an unprecedented degree of personalization in product recommendation, information retrieval, web search, and other domains. Despite this trend, tailoring options to specific users can be difficult because of privacy concerns, (actual or perceived), scarce data, or the infeasibility of complete personalization. For example, decisions about public projects (e.g., city parks) may require the choice of a single option: different projects cannot be built to meet the desires of different individuals. Similarly, companies may design a single product to maximize consumer satisfaction across its target market. In such settings, a single *consensus* recommendation must be made for the population as a whole. If recommendations are responsive to the preferences of individuals, we lie within the realm of *social choice*, an area receiving much attention in AI and computer science in recent years [?; 4; 9; 11; 14; 15; 23].

There is, of course, a middle ground between pure personalization and pure consensus recommendation. For example, if the company can configure its operations to produce three variants of the product in question, it must determine the three products that *jointly* maximize consumer satisfaction. With public projects, a city’s budget may allow small number of parks to be built. In domains like web search, if one has insufficient data about the user making a query, a limited number

of responses can be presented using available browser “real estate” to increase the odds that the user finds at least one result appealing. In such cases we fall somewhere between a single consensus decision and fully personalized recommendations for individuals. Some (perhaps implicit) aggregation of preferences must take place—fully personalized offerings to each individual are infeasible—placing us in the realm of social choice; but at the same time, we have an opportunity to tailor options to the preferences of the aggregated groups, and indeed, *explicitly design* the precise form of aggregation to optimize some social choice function.

In this paper, we develop a general model for such settings. The *budgeted social choice (BSC)* framework—unlike typical social choice models in which a single outcome (or ranking) is selected—allows multiple options to be offered, and assumes each user will *benefit from the best option, according to her own preferences*. However, the number of options is constrained by a *budget*, preventing pure personalization. This budget can take a variety of forms: a strict limit on the *number* of options (e.g., at most 3 products, or 10 web links on a page); or a maximum *total cost* (e.g., expenditure on city parks less than \$3M). We can also allow more nuanced tradeoffs between the cost of additional options and the increased benefit to the target population (e.g., add a fourth product option if increase in consumer satisfaction outweighs the cost of a fourth production line). Finally, we consider settings in which the budget is a function of both the options *and* their overall usage or uptake in the population: our framework allows for *fixed charges* (e.g., staffing an assembly line) and *unit costs* (e.g., marginal cost of one unit of product).

We can view budgeted social choice as applying *segmentation* [18] to social choice. Segmentation problems (defined formally below) seek k solutions to some combinatorial optimization problem which will be used by $n \geq k$ different “customers,” each with a different objective value for any of the candidate solutions. Optimization requires *segmenting* customers into k groups depending on which of the k solutions offers the greatest benefit. Bringing this view to social choice is illuminating when considering the tradeoff between the costs and benefits of “personalization” vs. consensus. Though our motivations are different, some multi-winner models in voting theory [5; 22] can also be viewed as a form of BSC, as we elaborate below.

We first present a general model of budgeted social choice:

given a set of alternative (or options) A and agents N , each agent having preferences over A , our goal is to assign an option from A to each $i \in N$. The assignment of an option to any i has a fixed cost (e.g., the cost of building a specific park) and a unit cost for each agent so assigned (e.g., per-person usage costs). Given a fixed budget, we seek an assignment that maximizes user satisfaction subject to the budget constraint. We also consider several special cases. When unit costs are zero, the problem is submodular and greedy optimization has theoretical guarantees; and when, in addition, all fixed costs are identical, we end up with proportional representation schemes from voting theory. We show the NP-hardness of even this simple case, prove that top- k methods based on typical rank aggregation methods can perform quite poorly, and examine our greedy algorithm empirically in both the general and special cases on two data sets. Empirical results suggest that greedy optimization is extremely fast and finds nearly optimal assignments (or option sets).

2 Background

We review some basic concepts from social choice (see [12] for more background) and segmentation problems. We assume a set of agents (or voters) $N = \{1, \dots, n\}$ and a set of alternatives $A = \{a_1, \dots, a_m\}$. Let Γ_A be the set of rankings (or votes) over A (i.e., permutations over A). Alternatives can represent any outcome space over which the voters have preferences (e.g., product configurations, restaurant dishes, candidates for office, public projects, etc.) and for which a single collective choice must be made. Agent ℓ 's preferences are represented by a ranking $v_\ell \in \Gamma_A$, where ℓ prefers a_i to a_j , denoted $a_i \succ_{v_\ell} a_j$, if $v_\ell(a_i) < v_\ell(a_j)$. The collection of votes $V = (v_1, \dots, v_n) \in \Gamma_A^n$ is a *preference profile*.

Given a preference profile, there are two main problems in social choice. The first is selecting a *consensus alternative*, requiring the design of a *social choice function* $f : \Gamma_A^n \rightarrow A$ which selects a “winner” given a profile. The second is selecting a *consensus ranking* [2], requiring a *rank aggregation function* $f : \Gamma_A^n \rightarrow \Gamma_A$. The consensus ranking can be used for many purposes; e.g., the top-ranked option may be taken as the consensus winner, or we might select the top k alternatives from the ranking in settings where multiple candidates are required (e.g., parliamentary seats [5; 22], or web search results [11]). *Plurality* is a common approach for selecting consensus alternatives: the option with the greatest number of “first place votes” wins (various tie-breaking schemes can be adopted). However, plurality fails to account for voter preferences for any alternative other than its top ranked (assuming sincere voting). Other schemes, such as the *Borda count* or *single transferable vote (STV)*, produce winners that are more sensitive to the relative preferences of voters. Among schemes that produce consensus rankings, *positional methods*, of which the *Borda ranking* is a special case, and the *Kemeny consensus* [16] are especially popular.

Definition 1. A positional scoring function (PSF) $\alpha : \{1, \dots, m\} \mapsto \mathbb{R}_{\geq 0}$ maps ranks onto scores such that $\alpha(1) \geq \dots \geq \alpha(m) \geq 0$. Given a ranking v_ℓ and alternative a , let $\alpha_\ell(a) = \alpha(v_\ell(a))$. The α -score of a , given profile V , is $\alpha(a, V) = \sum_{v_\ell \in V} \alpha_\ell(a)$. An α -ranking $r_\alpha^* = r_\alpha^*(V)$ is

any ranking that orders alternatives from highest to lowest α -score. The Borda score is the positional score given by score vector $\beta(i) = m - i$, and a Borda ranking $r_\beta^* = r_\beta^*(V)$ is defined using score vector β .

Definition 2. Let $\mathbf{1}$ be the indicator function, and r, v be two rankings. The Kendall-tau metric is $\tau(r, v) = \sum_{1 \leq i < j \leq m} \mathbf{1}[(v(a_i) - v(a_j))(r(a_i) - r(a_j)) < 0]$. Given a profile V , the Kemeny cost of a ranking r is $\kappa(r, V) = \sum_{v_\ell \in V} \tau(r, v_\ell)$. The Kemeny consensus is any ranking $r_\kappa^* = r_\kappa^*(V)$ that minimizes the Kemeny cost.

Intuitively, Kendall-tau distance measures the number of pairwise misorderings between an output ranking r and a vote v , while the Kemeny consensus minimizes total misorderings across profile V . While positional scoring is easy to implement, much work in computational social choice has focused on NP-hard schemes like Kemeny [11; 4].

Rank aggregation has interesting connections to work on *rank learning*, much of which concerns aggregating (possibly noisy) preference information from agents into full preference rankings. For example, Cohen et al. [8] focus on learning rankings from (multiple user) pairwise comparison data, while *label ranking* [14] considers constructing personalized rankings from votes. Often unanalyzed is why particular aggregation methods are suited to such settings.

Segmentation problems [18] formalize settings where one must solve an optimization problem for a number of different “customers,” each of whom has a different objective function, hence different values for a given solution. Full personalization (offering each customer her preferred solution) is often not feasible, but partial personalization can be realized by *segmenting* customers into groups, and determining the best solution for each segment. Assuming customers $i \leq n$, solutions $a \in A$, and customer “value” functions f_i over A for each i , a *fixed k -segmentation problem* asks for k solutions a_1, \dots, a_k maximizing:

$$\sum_{i \leq n} \max_{j \leq k} f_i(a_j). \quad (1)$$

From a customer-centric perspective, a customer benefits from the option made available that it favors most. A *variable segmentation problem* is identical except that the number of solutions k is variable, and each solution incurs a fixed cost γ . Thus we add $-\gamma k$ to the objective in Eq. 1.

3 General Budgeted Social Choice

While social choice techniques are used increasingly in applications like web search and recommender systems, the motivations for producing *consensus recommendations* for users with different preferences often varies. Consider, for instance, the motivation for “budgeted” consensus recommendation discussed in our introduction. If a decision maker can offer a limited set of K choices to a population of users to best satisfy their preferences, methods like Kemeny, Borda, etc. could be used to produce an aggregate ranking from which the top K alternatives are taken. However, there is little rationale for doing so without a deeper analysis of what it means to “satisfy” the preferences of the user population.

Rather than applying existing social choice schemes directly, we derive optimal consensus decisions from fundamental decision-theoretic principles and show how these differ from classic aggregation rules. Our model can be viewed as applying segmentation to problems in social choice, though our model differs from the variable segmentation formulation above [18]. We present our general model in this section, then analyze an important special case in the next.

We begin with the basic ingredients of social choice: voters N , alternatives A , and preference profile V . Rather than choosing a single alternative from A , we might allow a decision maker to propose a *recommendation set* $\Phi \subseteq A$, or slate of options, to improve voter *satisfaction*, but limit the choices by incorporating a *budget*. The precise formulation of voter satisfaction requires both a measure of individual voter satisfaction with Φ , and a measure of social welfare. Our framework can accommodate many measures of utility and welfare, but for concreteness we focus on one specific measure. First, we assume voters derive benefit from at most one $a \in \Phi$. This is consistent with many of the motivating scenarios discussed above. Models in which voters have preferences over the entire set [3] also fit into our budgeted framework, but will require a different style of analysis. Second, we use positional scoring rules (such as Borda) to quantify a voter’s satisfaction with her “selected” alternative (as defined below). Again, other models of voter satisfaction or utility can be incorporated in our model (these would require relatively minor changes to the analysis below). Finally we use the sum of such voter “utilities,” or *social welfare*, as our societal utility metric. Other models also fit within the framework (e.g., maxmin fairness) but we confine our analysis to social welfare maximization.

We assume a budget B that limits the total cost of a recommendation set $\Phi \subseteq A$: this is the key factor preventing full personalization in general. Thus, we must specify the *cost* of Φ . For each $a \in A$, let t_a be its *fixed cost* and u_a its *unit cost*. For instance, a company that decides to manufacture different product configurations must pay certain fixed production costs for each distinct configuration it offers (e.g., capital expenditures, product design costs); in addition, there are per-unit costs associated with producing each unit of the product (e.g., labor/material costs). The cost of Φ is then the fixed cost of each $a \in \Phi$, plus the unit-cost of a for each voter that derives benefit from a .

Intuitively, we’d like voters to derive benefit from their most preferred option in the recommendation set. However, since unit costs vary across $a \in \Phi$, a decision maker cannot simply propose a set Φ : allowing agents to choose their most preferred alternative freely may exceed the budget (e.g., all voters choose expensive options). Instead, the decision maker produces an *assignment* of alternatives to agents that maximizes social welfare. Putting this together gives:

Definition 3. A recommendation function $\Phi : N \rightarrow A$ assigns agents to alternatives. Given PSF α and profile V , the α -score of Φ is:

$$S_\alpha(\Phi, V) = \sum_{\ell \in N} \alpha_\ell(\Phi(\ell)). \quad (2)$$

Let $\Phi(N) = \{a : \Phi^{-1}(a) \neq \emptyset\}$ be the set of recommended

alternatives. The cost of Φ is:

$$C(\Phi) = \sum_{a \in A} \mathbf{1}[a \in \Phi(N)] \cdot t_a + \sum_{\ell \in N} u_{\Phi(\ell)}. \quad (3)$$

The first component in $C(\Phi)$ represents the fixed costs of recommended options, and second reflects total unit costs.

Definition 4. Given alternatives A , profile V , PSF α and budget $B > 0$, the budgeted social choice (BSC) problem is:

$$\max_{\Phi} S_\alpha(\Phi, V) \quad \text{subject to} \quad C(\Phi) \leq B. \quad (4)$$

We say that the problem is *infeasible* if every Φ has total cost exceeding B . While we define the problem using PSFs to measure social welfare, other variants of this problem are possible (as discussed above). Notice that the solution to a BSC problem not only chooses a set of alternatives, but explicitly—and optimally—segments voters into groups reflecting their assigned options. Though we adopt the intuitions of segmentation, our cost structure is more general than that of [18], as it allows per-unit charges and *variable* fixed costs. Our model has a few interesting special cases:

- If we wish to leave some voters “unassigned”, we use a *dummy item* d with $t_d = u_d = 0$. Each voter’s ranking for d can default to the bottom of its vote (i.e., positional score 0), or can reflect genuine preference for being unassigned. All such problems are feasible.
- Let $t_a = t$ (i.e., fixed charges are constant) and $u_a = 0$ for all $a \in A$, and $B = Kt$. Since unit costs are zero, we can select a *set* Φ of size K and let users select their most preferred option (hence Φ needn’t assign options to voters). This corresponds to the *limited choice model*, discussed in detail in the next section. If unit costs are non-zero but constant, $u_a = u$, we again can recommend a set of size $K = \lfloor (B - nu)/t \rfloor$.
- When fixed costs vary, but unit costs $u = u_a$ are constant, we generalize the limited choice model slightly: because unit costs are identical, agents can still select their preferred alternative from a slate (of varying size) whose total fixed cost does not exceed $B - nu$.
- If every recommendation function Φ satisfies $C(\Phi) \leq B$ (e.g., if all charges are zero), we are in a *fully personalizable setting*, and each agent is assigned their their most preferred alternative.

Our formulation can be modified in other ways. For instance, we may ignore budget, and instead allow an explicit tradeoff between social welfare (voter happiness) and costs, and simply maximize total score less total cost of Φ (as in [18] for specific forms of cost). In this way, unit cost would not prevent assignment of a more preferred option to a voter if the voter’s satisfaction outweighs the unit cost (once a fixed charge is incurred) or if it maximized surplus. This would better reflect a profit maximization motive in some settings (treating user satisfaction as a measure of willingness to pay). Our model as defined above is more appropriate in settings where users of a recommended alternative cannot be (directly) charged for its use (e.g., as in the case of certain public goods, corporate promotions or incentive programs, etc.).

Budgeted social choice is related to other optimization problems. When fixed costs vary but unit costs are constant, BSC is similar to *budgeted maximum coverage* [17] but differs in our use of scores for voter-alternative pairs (our “cover cost” varies with the assignment). More closely related is the *generalized maximum coverage* problem [7], but again differences exist (e.g., coverage is not required and unit costs are constant in [7]). Facility location problems are also related, especially to the limited choice model in the next section; and a form of (concave) unit costs for customers served at a facility is analyzed in [13].

General BSC can be written as an integer program (IP) with $m(n+1)$ variables and $1+mn+n$ constraints:

$$\begin{aligned} \max_{x_i, y_{\ell i}} \quad & \sum_{\ell \in N} \sum_{i=1}^m \alpha_{\ell}(a_i) \cdot y_{\ell i} \\ \text{subject to} \quad & \left[\sum_{i=1}^m t_{a_i} x_i \right] + \left[\sum_{\ell \in N} \sum_{i=1}^m u_{a_i} y_{\ell i} \right] \leq B, \quad (5) \\ & y_{\ell i} \leq x_i, \quad \forall \ell \leq n, i \leq m, \quad (6) \\ & \sum_{i=1}^m y_{\ell i} = 1, \quad \forall \ell \leq n. \quad (7) \end{aligned}$$

The variable x_i indicates whether alternative a_i appears in the recommendation assignment and $y_{\ell i}$ indicates whether agent ℓ has been assigned a_i . Constraint (5) is the budget limit, (6) and (7) ensure voters benefit only from alternatives in Φ , and benefit from exactly *one* such element.

Generally the solution to this IP will be computationally infeasible. Developing an approximation algorithm for BSC is complicated by the existence of unit costs. We need to limit the assignment of expensive alternatives despite “demand” from voters. Despite this, we develop a greedy heuristic algorithm called SweetSpotGreedy (or SSG); see Alg. 1. The main intuition is to successively “cover” or “satisfy” agents of a certain type by selecting their most preferred alternative. For a given $a \in A$, we sort voters based on their ranking of a and then compute the *bang-per-buck* ratio of assigning a to the first i voters—i.e., total score divided by total cost of assigning a to these i voters (much like knapsack heuristics). We pick the index i_a^* that maximizes the bang-per-buck ratio r_a^* . This is the *sweet spot*, where the marginal score improvement of assigning a to additional voters just fails to account for the incremental cost of offering more units of a . To the recommendation function Φ we add the a^* with the greatest ratio r_{a^*} and assign it to the $i_{a^*}^*$ agents who prefer it most. We repeat this procedure after removing the previously assigned a , each time selecting a new a^* and recommending it to the voters that maximize its bang per buck. This first phase may not produce a feasible assignment Φ : the budget may be exhausted before all agents are assigned an alternative. A second backtracking phase produces a feasible solution by rolling back the most recent updates to Φ from Phase 1. Each time an alternative is rolled back, we try to find an $a \in A$ that can be assigned to all unassigned agents without depleting the budget. If, after full backtracking, this can’t be achieved, the instance is infeasible.

Algorithm 1 The SweetSpotGreedy (SSG) algorithm.

Input: α, V, B , fixed costs t and unit costs u .

- 1: $\Phi \leftarrow \emptyset$ and $A^* \leftarrow \emptyset$
- 2: Let N_Φ denote $\{\ell : \Phi(\ell) \text{ is undefined}\}$
- 3: {PHASE 1 : ADD ITEMS WITH BEST SWEET SPOT}
- 4: **loop**
- 5: **for** $a \in A \setminus A^*$ **do**
- 6: $J \leftarrow \{\ell : a \succ_{\ell} \Phi(\ell) \text{ and } u_a \geq u_{\Phi(\ell)}\}$
- 7: $N_a = N_\Phi \cup J$
- 8: $R_a = \left[\frac{\alpha_{\ell}(a)}{u_a} \right]_{\ell \in N_\Phi} \cup \left[\frac{\alpha_{\ell}(a) - \alpha_{\ell}(\Phi(\ell))}{u_a - u_{\Phi(\ell)}} \right]_{\ell \in J}$
- 9: $SR_a \leftarrow \text{sort } R_a \text{ to get } (\beta_1/\gamma_1, \dots, \beta_{|R_a|}/\gamma_{|R_a|})$ {If $\gamma_i = 0$ then the “ratio” gets put in front of sorted list. For another denominator $\gamma_j = 0$ we then compare whether $\beta_i > \beta_j$.}
- 10: reorder N_a to $[\ell_1^a, \dots, \ell_{|N_a|}^a]$ so ℓ_i^a corresponds to β_i/γ_i
- 11: Let r_a^* and i_a^* be the max and argmax over i of $\{\frac{\sum_{j=1}^i \beta_j}{t_a + \sum_{j=1}^i \gamma_j} : i \in \{1, \dots, |SR_a|\}$ and $t_a + \sum_{j=1}^i \gamma_j \leq B - C(\Phi)\}$ if \emptyset then set to undefined.
- 12: **end for**
- 13: **if** $a^* \leftarrow \text{argmax}_{a \in A \setminus A^*} r_a^*$ is undefined **then**
- 14: **break** {all r_a^* is undefined—over budget}
- 15: **else**
- 16: append a^* to A^*
- 17: update Φ with $\{(\ell_i^{a^*}, a^*) : 1 \leq i \leq i_{a^*}^*\} \cup \{(\ell, a^*) : \ell \in N, a^* \succ_{\ell} \Phi(\ell) \text{ and } u_{a^*} \leq u_{\Phi(\ell)}\}$
- 18: **end if**
- 19: **end loop**
- 20: {PHASE 2: BACKTRACKING}
- 21: **while** Φ incomplete **do**
- 22: $a^* \leftarrow \text{pop } A^*$
- 23: remove $\{(\ell, a^*) : \ell \in N, \Phi(\ell) = a^*\}$ from Φ
- 24: $\tilde{A} \leftarrow \{a \in A : t_a + \sum_{\ell \in N_\Phi} u_a \leq B - C(\Phi)\}$
- 25: **if** $\tilde{A} \neq \emptyset$ **then**
- 26: $a^* \leftarrow \text{argmax}_{a \in \tilde{A}} \sum_{\ell \in N_\Phi} \alpha_{\ell}(a)$
- 27: update Φ with $\{(\ell, a^*) : \ell \in N_\Phi\}$ and **break**
- 28: **end if**
- 29: **end while**
- 30: **return** INFEASIBLE if $\Phi = \emptyset$, otherwise Φ

SSG has running time $O(m^2 n \log n)$. The intuition behind our algorithm is similar in spirit to the $1 - \frac{1}{e} - o(1)$ approximation algorithm for generalized maximum coverage [7]. However, that algorithm is impractical and of a more theoretical nature, requiring $O(m^2 n)$ calls to a fully polytime approximation scheme for maximum density knapsack.

As discussed above, when unit costs are zero our problem reduces to selecting a subset $\Phi \subseteq A$ with total fixed cost less than B . When fixed costs are constant, BSC reduces to the limited choice problem discussed below. In fact, SSG outputs the same recommendation function as the algorithm Greedy discussed below, which has an approximation guarantee:

Proposition 5. *If $u_a = 0$ and $t_a = 1$ for all $a \in A$ then SSG outputs the same recommendation as Greedy. Hence, it has an approximation ratio $1 - \frac{1}{e}$.*

We describe experiments with the SSG in Sec. 5. Full proofs of all results can be found in an preliminary, longer version of this paper [19].

4 The Limited Choice Model

We now consider the *limited choice problem*, a simple version of BSC in which one must choose a slate of K alternatives that maximizes voter satisfaction. As mentioned above, this can be viewed as BSC in which unit costs are zero and fixed costs are constant, $t_a = \lfloor B/K \rfloor$. For example, a company may be able to offer at most K products to its target market due to production constraints, where the products are substitutes (no consumer uses more than one); or a municipality may have budget for K new parks and citizens draw enjoyment from their most preferred park. Notice that we need not assign options to voters, but simply select a *recommendation set* Φ and allow voters to “use” their most preferred $a \in \Phi$:

Definition 6. A K -recommendation set is any $\Phi \subseteq A$ of size K . Given a PSF α , the α -score of Φ is:

$$S_\alpha(\Phi, V) = \sum_{\ell \in N} \max_{a \in \Phi} \alpha_\ell(a). \quad (8)$$

The optimal K -recommendation set with respect to α is:

$$\Phi_\alpha^* = \operatorname{argmax}_{|\Phi|=K} S_\alpha(\Phi, V). \quad (9)$$

$S_\alpha(\Phi, v)$ denotes the score with respect to a single vote v . We simply write S when α is clear from context, and S_β to denote Borda scoring.

Eq. 9 gives an identical model to the Chamberlin and Courant [5] scheme of proportional representation; thus results for that scheme apply directly to this version of the limited choice model, as we discuss below. Limited choice is related to the K -medians problem, treating alternatives as facilities, voters as customers, and (negated) positional scores as distances. Most work on K -medians focuses on metric rather than non-metric or ordinal settings (see [1] for an exception). Facility location is related as well (see Sec. 3).

Again, while positional scoring is our focus, other measures of utility and social desiderata fit within the limited choice framework. For example, we can use maxmin-fairness (with respect to positional scoring) encoded as:

$$\Phi_{\text{fair}}^* = \operatorname{argmax}_{|\Phi|=K} \min_{\ell \in N} S_\alpha(\Phi, v_\ell). \quad (10)$$

Similarly, setting $\alpha(i) = \mathbf{1}[i = 1]$ gives a satisfaction measure reflecting plurality scoring: the optimal Φ_α^* selects the K options with the greatest number of first-place “votes.” However, choosing the top K alternatives from a consensus ranking using positional scoring is, in general, not appropriate. For any ranking r , let $r|K$ denote the K top-ranked alternatives in r . The Borda ranking r_β^* can produce slates $r_\beta^*|K$ that are a factor of 2 from optimal using our limited-choice measure, while the α -ranking for arbitrary PSFs can be as much as a factor of K from optimal. We have, for any K :

Proposition 7. (a) $\inf_{(m,n,V)} \frac{S_\beta(r_\beta^*|K,V)}{S_\beta(\Phi_\alpha^*,V)} = 1/2$; and (b) $\inf_{(\alpha,m,n,V)} \frac{S_\alpha(r_\alpha^*|K,V)}{S_\alpha(\Phi_\alpha^*,V)} \leq 1/K$.

The full proof is omitted (see [19]), but the upper bound in (a) is demonstrated Fig. 1.

These results show that care must be taken in the application of rank aggregation methods to novel social choice

problems. In our limited choice setting, using positional scores (e.g., Borda) to determine the K most “popular” alternatives can perform very poorly.¹ We note too that STV, often used for proportional representation [23] can perform poorly: we can show that the slate produced by STV can also be a factor of 2 worse than optimal. Intuitively, the optimal slate appeals to the *diversity* of the agent preferences in a way that is not captured by “top K ” methods (this is one of the motivations for the proportional schemes [5; 22]). More importantly, our method is defined with respect to an explicit decision criterion.

Determining an optimal Φ in the limited choice model is NP-complete even in the specific case Borda scoring:²

Theorem 8. Deciding if there is a K -recommendation set Φ with (Borda) score $S_\beta(\Phi, V) \geq t$ is NP-complete, for $t \geq 0$,

We can formulate this NP-hard problem as an IP much like in BSC (see [22] who do so for proportional representation). Again, the IP will not scale to large problems. Fortunately, limited choice presents us with a constrained submodular maximization, which admits a simple greedy algorithm with approximation guarantees [21].

Algorithm Greedy. Inputs α, V and integer $K > 0$. Set $\Phi_0 \leftarrow \emptyset$. Update Φ K times by adding the option that increases score most, i.e., $\Phi_i \leftarrow \Phi_{i-1} \cup \{\operatorname{argmax}_{a \in A} S(\Phi_{i-1} \cup \{a\}, V)\}$. Output Φ_K .

Theorem 9. For a fixed V , $S(\cdot, V)$ defined over 2^A , with $S(\emptyset, V) = 0$, is submodular and non-decreasing. Thus, constrained maximization of Eq. (9) can be approximated within a factor of $1 - \frac{1}{e}$ by Greedy. That is, $\frac{S(\text{Greedy}, V)}{S(\Phi_\alpha^*, V)} \geq 1 - \frac{1}{e}$.

5 Experimental Evaluation

We evaluate our greedy algorithms, and compare them to top- k methods based on both Borda and Kemeny rankings, on two real-world data sets, *APA* and *Sushi*: Greedy is tested on both, and SweetSpotGreedy only on *Sushi* (since *APA* captures electoral data and is not suited to models with unit costs or non-uniform fixed costs).

Limited Choice: APA Dataset The American Psychological Association (APA) held a presidential election in 1980, where roughly 15,000 members expressed preferences for 5 candidates—5738 votes were full rankings. Members roughly divide into “academics” and “clinicians,” who are on “uneasy terms,” with voters in each group tending to favor like-minded candidates (candidate sets $\{1, 3\}$ and $\{4, 5\}$ appeal to different voters, with candidate 2 somewhere between) [10]. We applied our model to the full rankings with $K = 2$ and Borda scoring. We expected our model to favor “diverse” pairs (i.e., academic-clinician pairs score highest), and indeed, the optimal recommendation set is $\{3, 4\}$ with $S_\beta = 18182$. In fact, the four highest scoring pairs are

¹While the $\frac{1}{2}$ -approximation obtained in the case of Borda ranking is of interest theoretically, it does limit its practical use as an approximation method, since leaving 50% of societal value on the table is unlikely to be acceptable in practice.

²The NP-hardness of a *variant* of proportional representation [5] is known [23], but using a model with added flexibility that does not imply the NP-hardness of our limited choice model.




Fig. 1: Example showing that $r_\beta^*|K$ can be factor of 2 worse than optimal. Assume q items $\{0, 1, \dots, q - K - 1, \beta_1, \dots, \beta_K\}$, and $n = K(q - K - 2)$ votes. Votes are divided into K blocks, each with $q - K - 2$ votes. For each block $j \leq K$, item $j - 1$ is the top item in each vote, and $j \pmod{K}$ the worst. This means the optimal recommendation set is $\Phi^* = \{0, \dots, K - 1\}$, with $S_\beta(\Phi^*, V) = (q - 1)n$. The j th block of votes has a structure illustrated in the figure, with two example votes shown: the items j and $j \pmod{K}$ are fixed in the top/bottom spots and items β_1, \dots, β_K are fixed in positions $q/2 - K + 1, \dots, q/2$. (Fixed items are shaded.) Remaining items are arranged in the other positions in the first vote (unshaded). Starting with one such arrangement (e.g., the top vote in the figure), each alternative is “rotated downward” one unshaded position (with wrap around) to produce the next vote in the block. This is repeated until $q - K - 2$ votes are constructed for block j (i.e., one vote for each unshaded position). Thus, any non-fixed item occupies each unshaded rank in exactly once in block j . The average score of an unshaded item is $\sum_{i \in [q-2] \setminus \{q/2, \dots, q/2+K-1\}} i = \frac{-q^2+3q-2+qK+K^2-K}{-2q+2K+4} < q/2$ (since $q > K + 2$). Hence the average score of any item in $\{K, \dots, q - K - 1\}$ (which occupy only unshaded ranks) across all blocks is less than $q/2$. Also the average score of any item in Φ^* is less than $q/2$: item $j - 1$ has score $q - 1$ in block j but has score 0 in block $j - 2 \pmod{K}$, and average less than $q/2$ in all other blocks (it is unshaded in those). But the average score of β_i is at least $q/2$ (since its position is fixed in all blocks). Hence the top K items of the Borda ranking r_β^* are β_1, \dots, β_K . But $S_\beta(r_\beta^*|K, V) = (q/2 + K - 1)n$, so $S(r_\beta^*|K, V)/S(\Phi^*, V) = (q/2 + K - 1)/(q - 1)$, which approaches $1/2$ from above as $q \rightarrow \infty$.

diverse in this sense. Greedy outputs the diverse set $\{1, 5\}$ with score 17668, whereas the top 2-set from the Borda and Kemeny rankings is $\{1, 3\}$ with score 17352. Not only do Borda and Kemeny produce an inferior pair with respect to score, the pair is non-diverse, illustrating the key weakness of top- k methods based on standard consensus ranking methods. The quality of the Borda/Kemeny approximations is even worse with scoring functions that exaggerate score differences across rank positions (see below).

Limited Choice: Sushi Dataset The *sushi dataset* contains 5000 full preference orderings over 10 varieties of sushi [15]. In the limited choice setting, we might have a banquet in which only a small selection of sushi types can be provided to a large number of guests. Table 1 shows the approximation ratios of various algorithms for different slate sizes K , using an exponentially decreasing PSF $\alpha_{\text{exp}}(i) = 2^{m-i}$. CPLEX was used to solve an IP to determine optimal slates (computation times are shown). We evaluate Greedy, random sets of size K (averaged over 20 instances for each K), and Borda and Kemeny (recommending the top K options). We see that Greedy always finds the optimal slate (in fact, does so for all $K \leq 9$), and does so very quickly (under 1s.) relative to CPLEX optimization. Borda and Kemeny provide decent approximations, but are not optimal. Unsurprisingly, for large K (relative to $|A|$) random subsets do well, but perform poorly for small K . Results using Borda scoring are similar except that, unsurprisingly, random sets yield better approximations, since Borda count penalizes less for recommending lower-ranked alternatives than the exponential PSF.

General BSC: Sushi Dataset We tested SSG on the sushi dataset in the general BSC model. We first randomly generated fixed costs while holding unit costs at zero. This offers a slight generalization of limited choice. Integer fixed costs for each sushi variety were chosen uniformly at random from $[20, 50)$, with a budget of 100: so we typically recommend 2 to 5 items. We compared SSG to the optimal solution

K	Greedy	Borda	Kemeny	Random	CPLEX (sec.)
2	1.0	1.0	0.932	0.531	49.1
3	1.0	0.986	0.949	0.729	90.38
5	1.0	0.989	0.970	0.813	20.32
7	1.0	1.0	1.0	0.856	13.16

Table 1: Results using the limited choice model on the sushi dataset with 10 alternatives and 5000 full rankings. Four algorithms are shown in the columns 2–5 along with their approximation ratios for each K . CPLEX solution times are shown in the last column.

(using an IP solved with CPLEX) on 20 random cost profiles (the *preferences* are fixed by the data set).³ Borda scoring and the exponential PSF α_{exp} give similar results. With Borda, SSG is within 99% of the optimal recommendation function on average (it usually attains the optimum, and is never worse than 94% of optimal). Its running times lie in the range $[1.91s, 2.34s]$ (with a simple Python implementation). Meanwhile, CPLEX has an average solution time of 114s (and range $[69s, 176s]$), taking roughly two orders of magnitude longer to produce recommendations that improve of SSG by an average of only 1%.

In a second experiment, we varied both fixed and unit costs, with fixed much larger: integer unit costs were chosen uniformly from $[1, 4]$ and fixed costs from $[5000, 10000]$. A budget of 35000 allows roughly 3 unique alternatives to be recommended. We again compare SSG to the optimal recommendation function on 20 random instances. With Borda scores, SSG recommendations are on average within 98% of optimal, while taking 2–5s. to run. In contrast, CPLEX takes 458s. on average (range $[130s, 1058s]$) to produce an optimal solution. We achieve similar results using the exponential PSF, with SSG averaging 97% of optimal and taking 3–6s. while CPLEX averages 321s. (range $[131s, 614s]$). These ex-

³Top- k methods using consensus rankings like Borda or Kemeny cannot be used in general BSC problems, since they cannot reason about cost, nor can they assign voters to specific alternatives.

periments show that SSG has extremely strong performance, quickly finding excellent approximations to optimal recommendation functions.

6 Conclusion

We have introduced a new class of *budgeted social choice* problems that spans the spectrum from genuine consensus (or “one-size-fits-all”) decision making to fully personalized recommendation. Our key motivation, that some customization to the preferences of distinct groups of users may be feasible where complete individuation is not, holds of many real-world scenarios. Despite the diversity of user preferences, one must produce/recommend a limited number of alternatives, a problem best tackled by segmenting agents with similar preferences and selecting an alternative for each group. Our model includes certain schemes for proportional representation as special cases, and indeed motivates the possible application of proportional schemes to ranking and recommendation. Our framework often favors *diversity*, as opposed to *popularity*, of the chosen alternatives. Despite the theoretical complexity of optimization, our greedy methods perform extremely well on the data sets tested.

Extensions of this work include developing several variations of the budgeted model mentioned above: e.g., imposing separate budgets for fixed and unit costs; adopting different cost models; and incorporating more general utility metrics, including set-based voter preferences. If social welfare is a surrogate for decision maker revenue or return on investment (ROI), and other investment options are available (e.g. a government considering public projects) we may wish to relax the budget constraints and instead maximize ROI per unit cost. Deeper connections to proportional voting and other multi-winner schemes are also being explored, as is the performance of top- k methods using other rank aggregation schemes (such as Kemeny).

We are also developing probabilistic models and algorithmic methods for budgeted social choice with incomplete votes. Incomplete preferences are the norm in practical applications, such as recommender systems, where users offer preferences/ratings over only a small subset of alternatives. Budget-based slates can be constructed decision-theoretically exploiting distributional information over population preferences (e.g., those of some target market): recommendation sets or assignments are optimized with respect to the posterior over preferences given partial user data. In higher stakes settings, such as political or corporate voting, robust selection of alternatives based on *minimax regret* should prove very useful [20], particularly when used in conjunction with an elicitation process where users are asked only relevant queries. Preference distributions may prove useful in minimizing the total number of preference queries.

Acknowledgements: We thank Ariel Procaccia for useful discussions and the reviewers for their comments. This research was supported by the Natural Sciences and Engineering Research Council (NSERC).

References

- [1] M. R. Ackermann, J. Blömer, and C. Sohler. Clustering for metric and non-metric distance measures. *SODA-08*, pp.799–808, Philadelphia, 2008.
- [2] K. J. Arrow. A difficulty in the concept of social welfare. *J. Political Economy* 58:328–346, 1950.
- [3] S. Barberà, W. Bossert, and P. K. Pattanaik. Ranking sets of objects. In P.J. Hammond and C. Seidl (eds.), *Handbook of Utility Theory*, Ch. 17, pp.893–977. Kluwer, 2004.
- [4] J. Bartholdi III, C. Tovey, and M. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare* 6:157–165, 1989.
- [5] J. Chamberlin and P. Courant. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *Amer. Poli. Sci. Rev.*, 77(3):718–733, 1983.
- [6] Y. Chevaleyre, U. Endriss, J. Lang and Nicolas Maudet. A short introduction to computational social choice. *SOFSEM-07*, pp.51–69, Harrachov, CZ, 2007.
- [7] R. Cohen and L. Katzir. The generalized maximum coverage problem. *Inf. Process. Lett.*, 108(1):15–22, 2008.
- [8] W. W. Cohen, R. E. Schapire, and Y. Singer. Learning to order things. *JAIR*, 10:243–270, 1998.
- [9] V. Conitzer and T. Sandholm. Common voting rules as maximum likelihood estimators. *UAI-05*, Montreal, pp.145–152, 2005.
- [10] P. Diaconis. A generalization of spectral analysis with application to ranked data. *Annals of Stat.*, 17(3):949–979, 1989.
- [11] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. *WWW-01*, pp.613–622, 2001.
- [12] W. Gaertner. *A Primer in Social Choice Theory*. Oxford, 2006.
- [13] M. T. Hajiaghayi, M. Mahdian, and V. S. Mirrokni. The facility location problem with general cost functions. *Networks*, 42(1):42–47, 2003.
- [14] E. Hüllermeier, J. Fürnkranz, W. Cheng, and K. Brinker. Label ranking by learning pairwise preferences. *Art. Intel.*, 172:1897–1916, 2008.
- [15] T. Kamishima, H. Kazawa, and S. Akaho. Supervised ordering: An empirical survey. *IEEE ICDM-05*, pp.673–676, 2005.
- [16] J. G. Kemeny. Mathematics without numbers. *Daedalus*, 88(4):577–591, 1959.
- [17] S. Khuller, A. Moss, and J. Naor. The budgeted maximum coverage problem. *Inf. Process. Lett.*, 70(1):39–45, 1999.
- [18] J. Kleinberg, C. Papadimitriou, and P. Raghavan. Segmentation problems. *JACM*, 51:263–280, 2004.
- [19] T. Lu and C. Boutilier. Budgeted social choice: A framework for multiple recommendations in consensus decision making. *COMSOC-2010*, pp.55–66, Dusseldorf, 2010.
- [20] T. Lu and C. Boutilier. Robust approximation and incremental elicitation in voting protocols. *IJCAI-11*, Barcelona, 2011. To appear.
- [21] G. Nemhauser, L. Wolsey, and M. Fisher. An analysis of approximations for maximizing submodular set functions *Math. Prog.*, 14:265–294, 1978.
- [22] R. F. Potthoff and S. J. Brams. Proportional representation: Broadening the options. *J. Theor. Politics*, 10:147–178, 1998.
- [23] A. D. Procaccia, J. S. Rosenschein, and A. Zohar. On the complexity of achieving proportional representation. *Social Choice and Welfare*, 30:353–362, 2008.