

Vote Elicitation with Probabilistic Preference Models: Empirical Estimation and Cost Tradeoffs

Tyler Lu and Craig Boutilier

Department of Computer Science, University of Toronto
Toronto, Canada

{tl, cebly}@cs.toronto.edu
<http://www.cs.toronto.edu/~tl,~cebly>

Abstract. A variety of preference aggregation schemes and voting rules have been developed in social choice to support group decision making. However, the requirement that participants provide full preference information in the form of a complete ranking of alternatives is a severe impediment to their practical deployment. Only recently have incremental elicitation schemes been proposed that allow winners to be determined with partial preferences; however, while minimizing the amount of information provided, these tend to require repeated rounds of interaction from participants. We propose a probabilistic analysis of vote elicitation that combines the advantages of incremental elicitation schemes—namely, minimizing the amount of information revealed—with those of full information schemes—single (or few) rounds of elicitation. We exploit distributional models of preferences to derive the ideal *ranking threshold* k , or number of top candidates each voter should provide, to ensure that either a winning or a high quality candidate (as measured by max regret) can be found with high probability. Our main contribution is a general empirical methodology, which uses preference profile samples to determine the ideal ranking threshold for many common voting rules. We develop probably approximately correct (PAC) sample complexity results for one-round protocols with *any voting rule* and demonstrate the efficacy of our approach empirically on one-round protocols with Borda scoring.

Keywords: social choice, voting, preference elicitation, probabilistic rankings.

1 Introduction

Researchers in computer science have increasingly adopted preference aggregation methods from social choice, typically in the form of *voting rules*, for problems where a *consensus* decision or recommendation must be made for a group of users. The availability of abundant preference data afforded by search engines, recommender systems, and related artifacts, has accelerated the need for good computational approaches to social choice. One problem that has received little attention, however, is that of effective *preference elicitation* in social choice. Many voting schemes require users or *voters* to express their preferences over the entire space of options or *alternatives*, something that is not only onerous, but often extracts more information than is strictly necessary to determine a good consensus option, or *winner*. Reducing the amount of preference

information elicited is critical to easing cognitive and communication demands on users and mitigating privacy concerns.

Winners can't be determined in many voting schemes without a large amount of information in the worst case [2, 3]. Nonetheless, the development of elicitation schemes that work well in practice has been addressed very recently. Lu and Boutilier [10] use the notion of *minimax regret* for vote elicitation: this measure not only allows one to compute worst-case bounds on the quality of a proposed winner given *partial* voter preferences, it can also be used to drive incremental elicitation. Kalech et al. [6] develop several heuristic strategies for vote elicitation, including one scheme that proceeds in *rounds* in which voters provide larger “chunks” of information. This offers an advantage over the Lu-Boutilier schemes, where each voter query is conditioned on all previous responses of the other voters. Unfortunately, Kalech et. al's approach does not admit approximation (with quality guarantees), and no principles are provided to select an appropriate chunk size.

In this work, we develop an approach to vote elicitation that exploits distributional information over voter preferences to simultaneously reduce the amount of information elicited from voters and the number of *rounds* (a notion defined formally below) of elicitation. Indeed, these factors can be explicitly traded off against one another. Our model also supports approximation, using minimax regret, to further minimize the amount of information elicited, the number of rounds, or both. In this way, we provide the first framework that allows the design of vote elicitation schemes that address the complicated three-way tradeoff between approximation quality, total information elicited, and the number of rounds of elicitation.

Developing analytical bounds depends, of course, on the specific distributional assumptions about the preferences and the voting rule in question. While we make some suggestions regarding the types of results one might derive along these lines, our primary contribution is an empirical methodology that allows a designer to assess these tradeoffs and design elicitation schemes for *any preference distribution*, and *any voting rule* that can be interpreted using some form of scoring. To illustrate the use of both our general elicitation framework and our empirical methodology, we analyze one-round vote elicitation protocols. We develop general PAC sample complexity bounds for such one-round protocols. We then analyze these protocols empirically using Mallows models of preferences distributions [11, 12] and Borda scoring as the voting protocol. Our results suggest that good, even optimal, results can be obtained in one-round protocols even when only a small portion of the preferences of the voters is elicited.

2 Background

We begin with a brief overview of relevant background on social choice, vote elicitation, and preference distributions.

2.1 Voting Rules

We first define our basic social choice setting (see [5, 1] for further background). We assume a set of *agents* (or *voters*) $N = \{1, \dots, n\}$ and a set of *alternatives* $A =$

$\{a_1, \dots, a_m\}$. 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. Let Γ_A be the set of *rankings* (or *votes*) over A (i.e., permutations over A). Voter ℓ 's preferences are represented by a ranking $v_\ell \in \Gamma_A$. Let $v_\ell(a)$ denote the *rank* of a in v_ℓ . Then ℓ prefers a_i to a_j , denoted $a_i \succ_{v_\ell} a_j$, if $v_\ell(a_i) < v_\ell(a_j)$. We refer to a collection of votes $\mathbf{v} = \langle v_1, \dots, v_n \rangle \in \Gamma_A^n$ as a *preference profile*. Let \mathcal{V} be the set of all such profiles.

Given a preference profile, we consider the problem of selecting a *consensus alternative*, requiring the design of a *social choice function* or *voting rule* $r : \mathcal{V} \rightarrow A$ which selects a “winner” given voter rankings/votes. *Plurality* is one of the most common rules: the alternative with the greatest number of “first place votes” wins (various tie-breaking schemes can be adopted). Plurality does not require that voters provide rankings; however, this “elicitation advantage” means that it fails to account for relative voter preferences for any alternative other than its top choice. Other schemes produce winners that are more sensitive to relative preferences, among them, the Borda rule, Copeland, single-transferable vote (STV), the Kemeny consensus, maximin, Bucklin, and many others. We outline the *Borda rule* since we use it extensively below: let $B(i) = m - i$ be the Borda score for each rank position i ; the *Borda count* or *score* of alternative a given profile \mathbf{v} , is $s_B(a, \mathbf{v}) = \sum_\ell B(v_\ell(a))$. The winner is the a with the greatest Borda score.

Notice that both the Borda and plurality schemes *explicitly scores all alternatives* given voter preferences, implicitly defining “societal utility” for each alternative. Indeed, many (though not all) voting rules r can be interpreted as maximizing a “natural” *scoring function* $s(a, \mathbf{v})$ that defines some measure of the quality of an alternative a given a profile \mathbf{v} . We assume in what follows that our voting rules are *score-consistent* in this sense: $r(\mathbf{v}) \in \operatorname{argmax}_{a \in A} s(a, \mathbf{v})$. some “natural” scoring function $s(a, \mathbf{v})$.¹

2.2 Vote Elicitation

One obstacle to the widespread use of voting schemes that require full rankings is the informational and cognitive burden imposed on voters, and concomitant ballot complexity. Elicitation of sufficient, but still *partial* information about voter rankings could alleviate some of these concerns. We will assume in what follows that the partial information about any voter’s ranking can be represented as a collection of pairwise comparisons. Specifically, let the *partial vote* p_ℓ of voter ℓ be a partial order over A , or equivalently (the transitive closure of) a collection of *pairwise comparisons* of the form $a_i \succ a_j$. Let \mathbf{p} denote a partial profile, and $C(\mathbf{p})$ the set of consistent extensions of \mathbf{p} to full ranking profiles. Let \mathcal{P} denote the set of partial profiles.

If our aim is to determine the winner given a partial profile, theoretical worst-case results are generally discouraging, with the communication complexity of several common voting protocols (e.g., Borda) being $\Theta(nm \log m)$, essentially requiring communication of full voter preferences in the worst-case [3]. Despite its theoretical complexity, practical schemes for elicitation have been developed recently.

¹ We emphasize that *natural* measures of quality are the norm; trivially, any rule can be defined as score consistent using a simple indicator function.

Lu and Boutilier [10] use *minimax regret (MMR)* to determine winners given partial profiles, and also to guide elicitation. Intuitively, one measures the quality of a proposed winner a given \mathbf{p} by considering how far from optimal a could be in the worst case, given any completion of \mathbf{p} ; this is a 's *maximum regret* $MR(a, \mathbf{p})$. The *minimax optimal solution* is any alternative that is nearest to optimal in the worst case, i.e., with minimum max (minimax) regret. More formally:

$$Regret(a, \mathbf{v}) = \max_{a' \in A} s(a', \mathbf{v}) - s(a, \mathbf{v}) = s(r(\mathbf{v}), \mathbf{v}) - s(a, \mathbf{v}) \quad (1)$$

$$MR(a, \mathbf{p}) = \max_{\mathbf{v} \in C(\mathbf{p})} Regret(a, \mathbf{v}) \quad (2)$$

$$MMR(\mathbf{p}) = \min_{a \in A} MR(a, \mathbf{p}) ; \quad a_{\mathbf{p}}^* \in \operatorname{argmin}_{a \in A} MR(a, \mathbf{p}) . \quad (3)$$

This gives us a form of robustness in the face of vote uncertainty: every alternative has worst-case error at least as great as that of $a_{\mathbf{p}}^*$. Notice that if $MMR(\mathbf{p}) = 0$, then the minimax winner $a_{\mathbf{p}}^*$ is optimal in any completion $\mathbf{v} \in C(\mathbf{p})$. MMR can be computed in polytime for several common voting rules, including Borda [10].

MMR can also be used to determine (pairwise or top- k) queries that quickly reduce minimax regret; indeed, in a variety of domains, regret-based elicitation finds (optimal) winners with small amounts of voter preference information, and can find near-optimal candidates (with bounded maximum regret) with even less. However, these elicitation methods implicitly condition the choice of a voter-query pair on *all past responses*. Specifically, the choice any query is determined by first solving the minimax regret optimization (Eq. (3)) w.r.t. the responses to all prior queries. Hence each query must be posed in a separate *round*, making it impossible to “batch” multiple queries for a specific user.

Kalech et al. [6] develop two elicitation algorithms for winner determination with score-based rules (e.g., Borda, range voting) in which voters are asked for k th-ranked candidates in decreasing order of k . Their first method proceeds in fine-grained rounds much like the MMR-approach above, until a *necessary winner* [8, 16] is discovered. Their second method proceeds for a predetermined number of rounds, asking each voter at each stage for fixed number of positional rankings (e.g., the top k candidates, or the next k' candidates, etc.). Since termination is predetermined, necessary winners may not be discovered; instead possible winners are returned. Tradeoffs between the number of rounds and amount of information per round are explored empirically. One especially attractive feature of this approach is the explicit batching of queries: voters are only queried a fixed (ideally small) number of times (though each query may request a lot of information), thus minimizing interruption, waiting time, etc. However, no quality guarantees are provided, nor is a theoretical basis provided for selecting the amount of information requested at any round.

2.3 Probabilistic Models of Population Preferences

Probabilistic analysis in social choice has often focused on the *impartial culture model*, which asserts that all preference orderings are equally likely. However, the plausibility of this assumption, and the relevance of theoretical results based on it, have been seriously called into question by behavioral social choice theorists [14]. More realistic probabilistic models of preferences, or parameterized families of distributions over

rankings, have been proposed in statistics, econometrics and psychometrics. These models typically reflect some process by which people rank, judge or compare alternatives. Many models are unimodal, based on a “reference ranking” from which user rankings are seen as noisy perturbations. A commonly used model, adopted widely in machine learning—and one we exploit below—is the *Mallows ϕ -model* [11]. It is parameterized by a modal or *reference ranking* σ and a *dispersion parameter* $\phi \in (0, 1]$; and for any ranking r we define: $P(r; \sigma, \phi) = \frac{1}{Z} \phi^{d(r, \sigma)}$, where d is the Kendall-tau distance and Z is a normalization constant. When $\phi = 1$ we obtain the uniform distribution over rankings, and as $\phi \rightarrow 0$ we approach the distribution that concentrates all mass on σ . A variety of other models have been proposed that reflect different interpretations of the ranking process (e.g., Plackett-Luce, Bradley-Terry, Thurstonian, etc.); we refer to [12] for a comprehensive treatment. Mixtures of such models, which offer additional modeling flexibility (e.g., by admitting multimodal preference distributions), have also been investigated (e.g., [13, 9]).

Sampling rankings from specific families of distributions is an important task that we also rely on below. The *repeated insertion model (RIM)*, introduced by Doignon et al. [4], is a generative process that can be used to sample from certain distributions over rankings and provides a practical way to sample from a Mallows model. A variant of this model, known as the *generalized repeated insertion model (GRIM)*, offers more flexibility, including the ability to sample from conditional Mallows models [9].

3 A Regret-based Model of Probabilistic Vote Elicitation

We begin by developing a general model of vote elicitation that allows one to make explicit tradeoffs between the number of rounds of elicitation, the amount of information provided by each voter, and approximation quality. Let a *query* refer to a “single” request for information from a voter. Types of queries include simple pairwise comparisons (e.g., “Do you prefer a to b ?”); sets of such comparisons; more involved partial requests (e.g., “Who are your top k candidates?”); or requests for entire rankings. Different queries have different “costs”—both in terms of voter cognitive effort and communication costs (which range from 1 to roughly $m \log m$ bits)—and provide varying degrees of information.

Given a particular class of queries \mathcal{Q} , informally, a *multi-round voting protocol* selects, at each round, a subset of voters, and one query per selected voter. The *voter-query (VQ) pairs* selected at round t can be conditioned on the responses to all previous queries. More formally, let I_{t-1} be the *information set* available at round t (i.e., responses to queries at rounds $1, \dots, t-1$). We represent this information set as a partial profile \mathbf{p}^{t-1} , or a set of pairwise comparisons for each voter.² A *protocol* then consists of: (a) a *querying function* π , i.e., a sequence of mappings $\pi^t : \mathcal{P} \mapsto (N \mapsto \mathcal{Q} \cup \{0\})$, selecting for each voter a single query at stage t given the current information set; and (b) a *winner selection function* $\omega : \mathcal{P} \mapsto A \cup \{0\}$, where $\omega(\mathbf{p})$ denotes the winner

² Most natural constraints, including responses to many natural queries (e.g., pairwise comparison, top- k , etc.), can be represented in this way. One exception: arbitrary positional queries of the form “what candidate is in rank position k ?” induce disjunctive constraints, unless positions k are queried in (ascending or descending) order.

given partial profile \mathbf{p} . If $\omega(\mathbf{p}^t) = 0$, no winner is declared and the protocol proceeds to round $t + 1$; otherwise the protocol terminates with the chosen winner at round t . If $\pi^t(\mathbf{p}^{t-1})(\ell) = 0$, then no query is posed to voter ℓ at round t .

Suppose we have a distribution P over complete voter profiles. Given a protocol $\Pi = (\pi, \omega)$, we have an induced distribution over runs of Π , which in turn gives us a distribution over various properties reflecting the cost and performance of Π . There are three general properties of interest to us:

- (a) *Quality of the winner*: if Π terminates with information set \mathbf{p} and winner a , we can measure quality using either expected regret, $\sum_{\mathbf{v}} \text{Regret}(a, \mathbf{v})P(\mathbf{v}|\mathbf{p})$, or maximum regret, $MR(a, \mathbf{p})$. If Π is an *exact protocol* (always determining a true winner), both measures will be zero. We focus here on max regret, which provides worst-case guarantees on winner quality. In some settings, expected regret might be more suitable.
- (b) *Amount of information elicited*: this can be measured in various ways (e.g., equivalent number of pairwise comparisons or bits).
- (c) *Number of rounds* of elicitation.

There is a clear tradeoff between these factors. A greater degree of approximation in winner selection can be used to reduce informational requirements, rounds, or both [10]. For any fixed quality threshold, the number of rounds and the amount of information elicited can also be traded off against one another. At one extreme, optimal outcomes can clearly be found in one round if we ask each voter for full rankings. At the other extreme, optimal policies minimizing expected elicited information can always be constructed (though this will likely come at great computational expense) by selecting a single VQ-pair at each round, where each query carries very little information (e.g., a simple pairwise comparison), at a dramatic cost in terms of number of rounds. How one addresses these tradeoffs depends on the costs associated with each of these factors. For example, the cost of elicited information might reflect the number and type of queries asked of voters, while the cost associated with rounds might reflect interruption and delay experienced by voters as they “wait” for other voters to answer queries before receiving their own next query.³

Computing optimal protocols for specific voting rules, query classes, distributions over preferences, and cost models is a very important problem that can be addressed explicitly using our framework. The framework supports both Bayesian and PAC-style (probably approximately correct) analysis. We illustrate its use by considering a specific type of protocol using a PAC-style analysis in the next section.

4 Probably Approximately Correct One-round Protocols

Imagine we require a *one-round protocol*, where each voter can be asked, exactly once, to list their top- k candidates. A natural question is: what is the minimum value k^* for

³ We’re being somewhat informal, since some voters may only be queried at subset of the rounds. If a (conditional) sequence of queries is asked of a single voter ℓ without any interleaving queries to another voter j , we might count this as a single “session” or round for ℓ . These distinctions won’t be important in what follows.

which such top- k queries ensure that the resulting profile \mathbf{p} has low minimax regret, $MMR(\mathbf{p}) \leq \varepsilon$, with high probability, at least $1 - \delta$? We call ε and δ the *minimax regret accuracy* and *confidence* parameters, respectively. Obviously, such a k^* exists: with $k = m - 1$, we elicit each voter’s full ranking, always ensuring $MMR(\mathbf{p}) = 0$. This question is of interest when, for example, more than one round of elicitation is infeasible or very costly, an approximate solution (with tolerance ε) is suitable, and some small probability δ of a poor solution is acceptable.

Let $\mathbf{p}[k]$ denote the restriction of profile $\mathbf{v} = (v_1, \dots, v_n)$ to the subrankings consisting of each voter’s top k candidates. For any distribution P over voter preferences \mathbf{v} , $MMR(\mathbf{p}[k])$ is a random variable. Let $q_k = P(MMR(\mathbf{p}[k]) \leq \varepsilon)$. We would like to find $k^* = \min\{k : q_k \geq 1 - \delta\}$. Even if we assume P has a particular form, computing k^* might be analytically intractable; or the analytically derived upper bounds may be too loose to be of practical use. If one can instead sample vote profiles from the true distribution—without necessarily knowing what P is—a simple empirical methodology can be used to determine a small \hat{k} that, with high probability, has the desired MMR accuracy with near the desired MMR confidence (see Theorem 1 below). Specifically, we take the following steps:

- (a) Specify the following parameters: MMR accuracy $\varepsilon > 0$, MMR confidence $\delta > 0$, *sampling accuracy* $\xi > 0$, and *sampling confidence* $\eta > 0$.
- (b) Obtain t i.i.d. samples of vote profiles $S = (\mathbf{v}_1, \dots, \mathbf{v}_t)$ where

$$t \geq \frac{1}{2\xi^2} \ln \frac{2(m-2)}{\eta}. \quad (4)$$

- (c) Output \hat{k} , the smallest k for which

$$\hat{q}_k \equiv \frac{|\{i \leq t : MMR(\mathbf{p}_i[k]) \leq \varepsilon\}|}{t} > 1 - \delta - \xi.$$

The parameters ξ and η are required to account for sampling randomness, and are incorporated as part of the statistical guarantee on the algorithm’s success (see Theorem 1). In summary, the approach is to estimate q_k (which is usually intractable to derive analytically) using \hat{q}_k , and take the smallest \hat{k} that, accounting for sampling error, is highly likely to have the true probability, $q_{\hat{k}}$, lie close to the desired MMR confidence threshold $1 - \delta$. The larger the sample size t , the better the estimates, resulting in smaller ξ and η . Using a sample set specified as in the algorithm, one can obtain a PAC-style guarantee [15] on the quality of one-round, top- \hat{k} elicitation:

Theorem 1. *Let $\varepsilon, \delta, \eta, \xi > 0$. If the sample size t satisfies Eq. (4), then for any preference profile distribution P , with probability $1 - \eta$ over i.i.d. samples $\mathbf{v}_1, \dots, \mathbf{v}_t$, we have: (a) $\hat{k} \leq k^*$; and (b) $P[MMR(\mathbf{p}[\hat{k}]) \leq \varepsilon] > 1 - \delta - 2\xi$.*

Proof. For any $k \leq m - 2$ (for $k = 0$, minimax regret is $n(m - 1)$ and for $k \geq m - 1$ minimax regret is 0, so we are not interested in these cases), the indicator random variables $\mathbf{1}[MMR(\mathbf{p}_i[k]) \leq \varepsilon]$ for $i \leq t$ are i.i.d. By the Hoeffding bound, we have

$$\Pr_{S \sim P^t} [|\hat{q}_k - q_k| \geq \xi] \leq 2 \exp(-2\xi^2 t).$$

If we choose t such that $\frac{\eta}{m-2} \leq 2 \exp(-2\xi^2 t)$ we obtain Inequality (4) and

$$\begin{aligned}
& \Pr_{S \sim P^t} \left((|\hat{q}_1 - q_1| \leq \xi) \wedge (|\hat{q}_2 - q_2| \leq \xi) \wedge \dots \wedge (|\hat{q}_{m-2} - q_{m-2}| \leq \xi) \right) \\
&= 1 - \Pr_{S \sim P^t} \left[\bigcup_{k=1}^{m-2} |\hat{q}_k - q_k| > \xi \right] \\
&\geq 1 - (m-2) \cdot \frac{\eta}{m-2} \\
&= 1 - \eta,
\end{aligned} \tag{5}$$

where Inequality (5) follows from the union bound. Thus with probability at least $1 - \eta$, uniform convergence holds, and we have $\hat{q}_{k^*} > q_{k^*} - \xi > 1 - \delta - \xi$. Since \hat{k} is the smallest k with $\hat{q}_k > 1 - \delta - \xi$ we have $\hat{k} \leq k^*$. Furthermore, $q_{\hat{k}} > \hat{q}_{\hat{k}} - \xi > (1 - \delta - \xi) - \xi = 1 - \delta - 2\xi$, which shows part (2). ■

We note several significant features of this result. First, it is distribution-independent—we need t i.i.d. samples from P , where t depends only on ξ , η and m , and *not* on any property of P . Of course, depending on the nature of the distribution, the required sample size may be larger than necessary (e.g., if P is highly concentrated). Second, note that an algorithm that outputs $\hat{k} = m - 1$ guarantees $MMR = 0$, but is effectively useless to the elicitor; hence we desire an algorithm that proposes a \hat{k} that is not much larger than the optimal k^* . Our scheme guarantees $\hat{k} \leq k^*$. Third, while the true probability $q_{\hat{k}}$ of the estimated \hat{k} satisfying the regret accuracy requirement may not meet the confidence threshold, it lies within some small tolerance of that threshold. This is unavoidable in general. For instance, if we have $q_{k^*} = 1 - \delta$, there is potentially a significant probability that $\hat{q}_{k^*} < 1 - \delta$ for any finite sample; but our result ensures that there is only a small probability that $\hat{q}_{k^*} < 1 - \delta - \xi$. Fourth, part (b) of Theorem 1 remains valid if the sum $\delta + \xi$ is fixed (and in some sense, this sum can be interpreted as our ultimate confidence); but variation in δ and ξ does impact sample size (and part (a)). One can reduce the required sample size by making ξ larger and reducing δ correspondingly, maintaining the same “total” degree of confidence, but the guarantee in part (a) becomes weaker since k^* generally increases as δ decreases. This is a subtle tradeoff that should be accounted for in the design of an elicitation protocol.

We can provide no *a priori* guarantees on how small k^* might be, since this depends crucially on properties of the distribution; in fact, it might be quite large (relative to m) for, say, the impartial culture model (as we see below). But our theorem provides a guarantee on the size of \hat{k} w.r.t. the optimal k^* .

An analogous result can easily be obtained if one is interested in determining the smallest k for a one-round protocol that has small *expected MMR*. However, using expectation does not preclude MMR from being greater than a desired threshold with significant probability. Hence, expected MMR may be ill-suited to choosing k in many voting settings. The techniques above can also be used in a Bayesian fashion, where instead of using minimax regret to determine robust winners, one uses *expected regret* (i.e., expected loss relative to the optimal candidate given uncertainty over completions the partial profile). We defer treatment of expected regret to another article.

Our empirical methodology can also be used in a more heuristic fashion, *without* derivation of precise confidence bounds. One can simply generate random profiles, use

the empirical distribution over $MMR(\mathbf{p}[k])$ as an estimate of the true distribution, and select the desired k based directly on properties of the empirical distribution (e.g., represented as histograms, as we illustrate in the next section).

Finally, we note that samples can be obtained in a variety of ways, e.g., drawn from a learned preference model, such as a Mallows model or Mallows mixture (e.g., using RIM), or simply obtained from historical problem instances. In multi-round protocols, the GRIM model can be used to realize conditional sampling if needed. Our empirical methodology is especially attractive when k^* cannot easily be derived analytically (which may well be the case for Mallows, Plackett-Luce, and other common models).

5 Empirical Results

To explore the effectiveness of our methodology, we ran a suite of experiments, sampling voter preferences from Mallows models using a range of parameters, computing minimax regret for each sampled profile for various k , and estimating both the expected minimax regret and the MMR-distribution empirically. We also discuss experiments with two real-world data sets. Borda scoring is used in all experiments.

For the Mallows experiments, a preference profile is constructed by drawing n i.i.d. rankings, one per voter, from a fixed Mallows model. Each experiment varies the number of voters n , number of alternatives m , and dispersion ϕ , and uses 100 preference profiles. We simulate the elicitation of top- k preferences and measure both MMR and true regret (w.r.t. the true preferences and true winner) for $k = 1, \dots, m - 1$; results are “normalized” by reporting max regret and true regret *per voter*. Fig. 1 shows histograms reflecting the empirical distribution of both MMR and true regret for various k , ϕ , n , and m . That is, in each collection of histograms, as defined by particular (m, n, ϕ) parameter values, we generated 100 instances of random preference profiles. For each instance of a profile, and each k , we compute MMR of the partial votes when top- k preferences are revealed in the profile—this represents one data point along the horizontal axis, in the histogram corresponding to that particular k , and to parameters values (m, n, ϕ) . Note that (normalized) MMR per voter can range from 0 to 9 since we use Borda scoring.

Clearly MMR is always zero when $k = m - 1 = 9$. For small ϕ (e.g., 0.1–0.4), preferences across voters are reasonably similar, and values of $k = 1$ –3 are usually sufficient to find the true winner, or one with small max regret. But even with $m = 10$, $n = 100$ and $\phi = 0.6$, $k = 4$ results in a very good approximate winner: $MMR \leq 0.6$ in 90/100 instances. Even the most difficult case for partial elicitation—the uniform distribution with $\phi = 1$ —gives reasonable MMR guarantees with high probability with less than full elicitation ($k = 5$ –7, depending on one’s tolerance). The heuristic use of the empirical distribution in this fashion is likely to suffice in practice in a variety of settings; but we can apply the theoretical bounds above as well. Since we have a $t = 100$ (admittedly a small sample), by Eq. (4), we can set $\eta = 0.05$ and $\xi = 0.17$, and with $\delta = 0.9$, $\varepsilon = 0.5$, we obtain $\hat{k} = 4$. By Theorem 1, we are guaranteed with probability 0.95 that $\hat{k} \leq k^*$ and $q_{\hat{k}} > 0.56$. If we wanted $q_{\hat{k}}$ to be closer to 0.9, then requiring $t \geq 28842$ gives $\xi = 0.01$ and $q_{\hat{k}} > 0.88$.

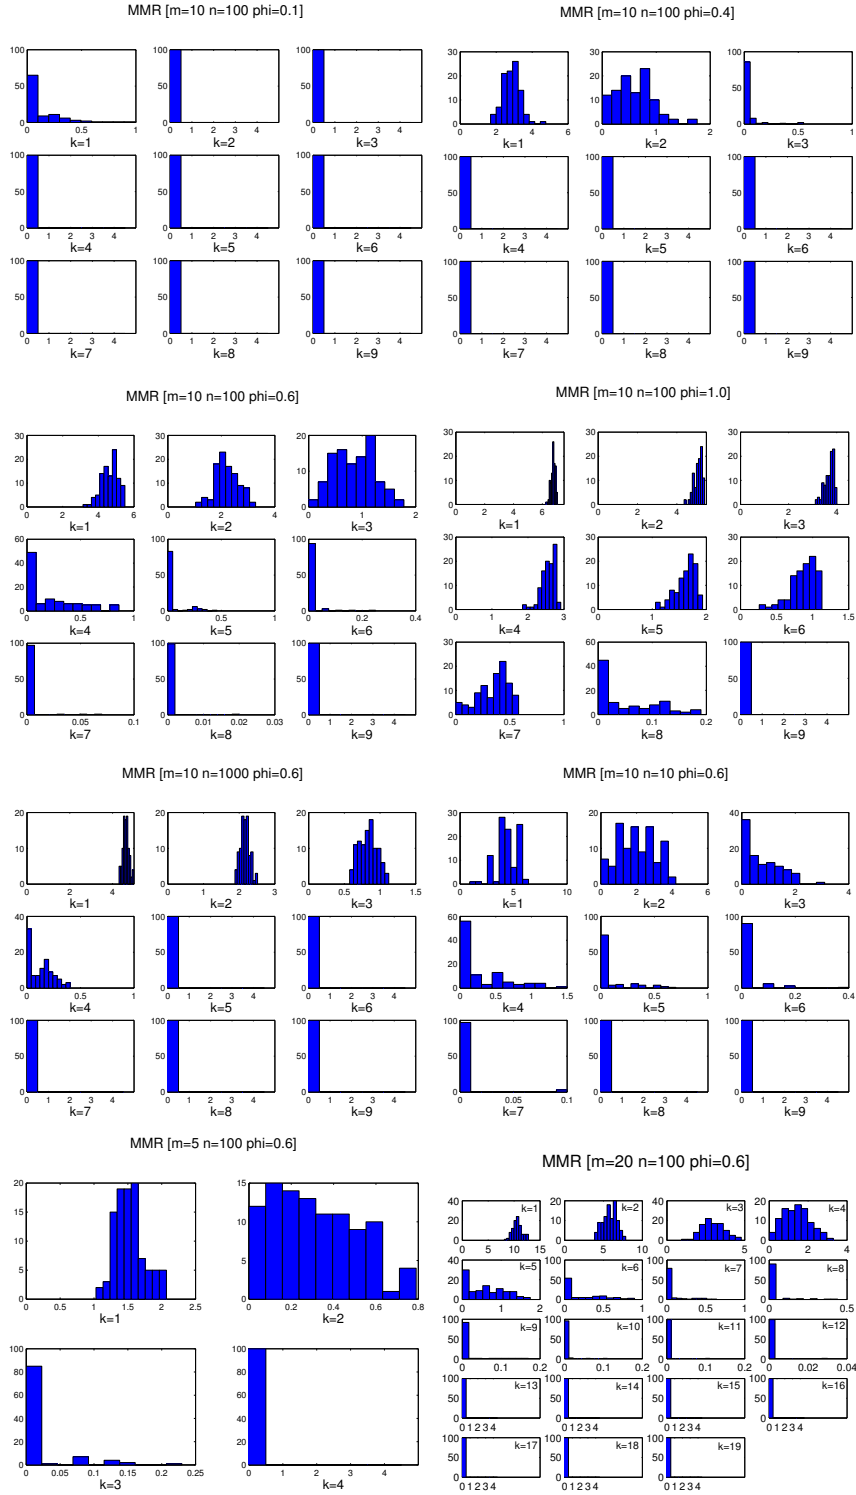


Fig. 1: MMR plots for various ϕ , n and m : for $m = 10$, $n = 100$ with $\phi \in \{0.1, 0.4, 0.6, 1.0\}$ and fixed $\phi = 0.6$ with $n \in \{10, 1000\}$; $m = 5$, $\phi = 0.6$; and $m = 20$, $\phi = 0.6$. Each histogram shows the distribution of MMR, normalized by n , after eliciting top- k .

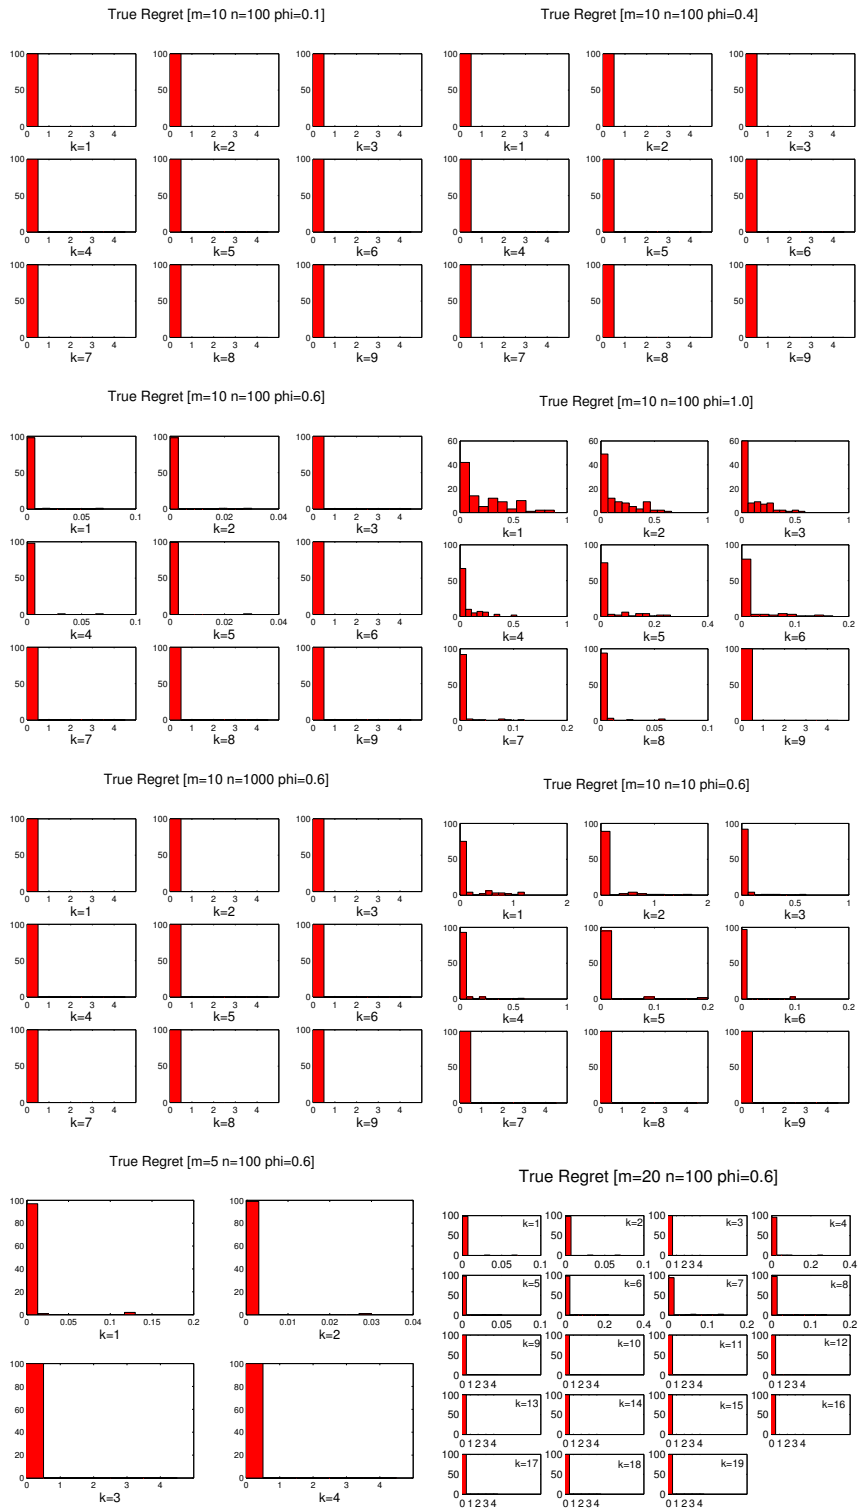


Fig. 2: The corresponding true regrets of experiments shown in Fig. 1.

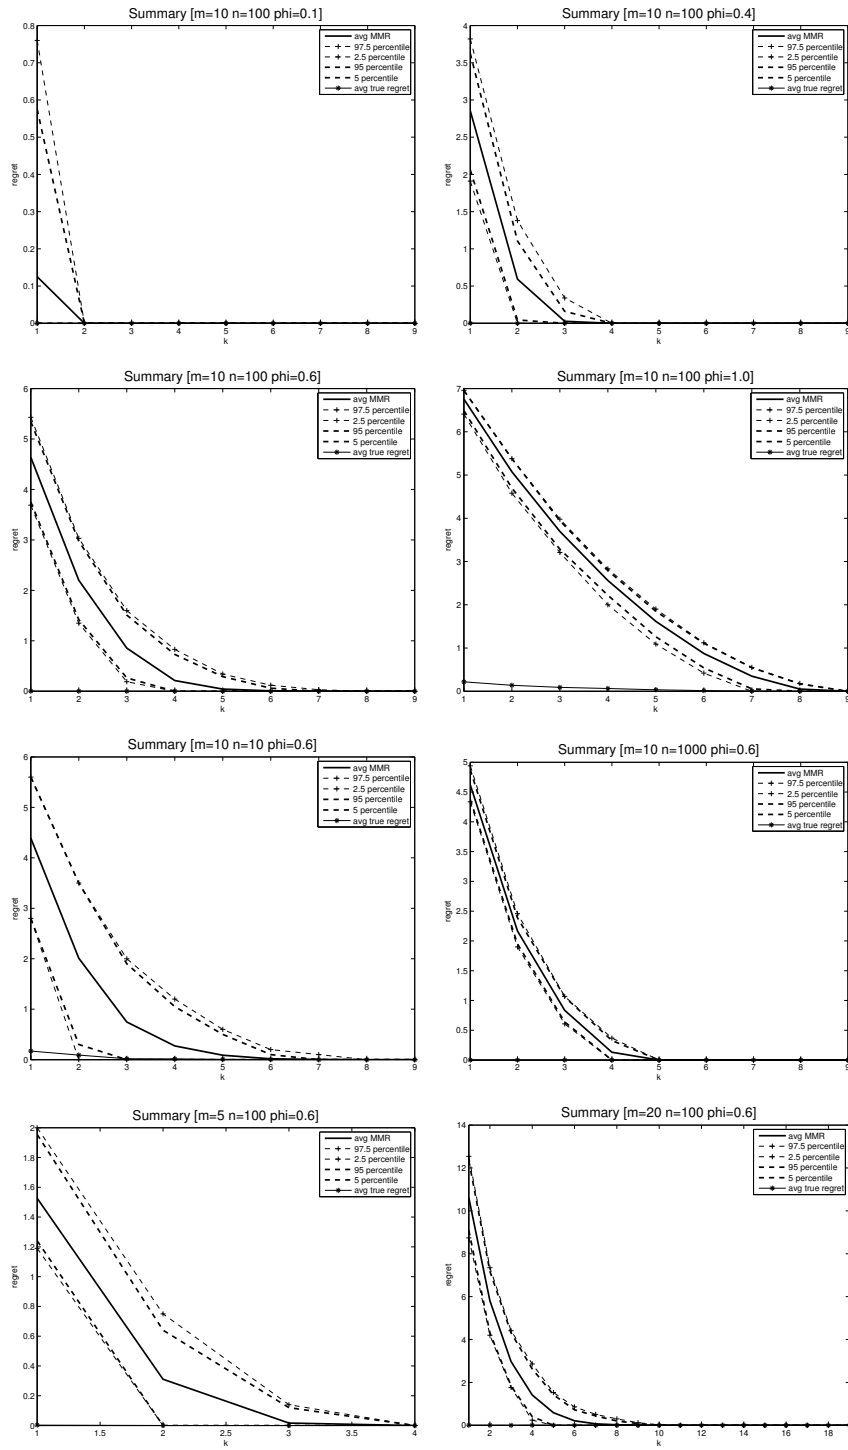


Fig. 3: Each plot corresponds to a summary of the experiments in Fig. 1, and shows the reduction in regret (avg. normalized (per voter) MMR and true regret over all instances) as k increases. Percentiles (.025, 0.05, 0.95, 0.975) for MMR are shown.

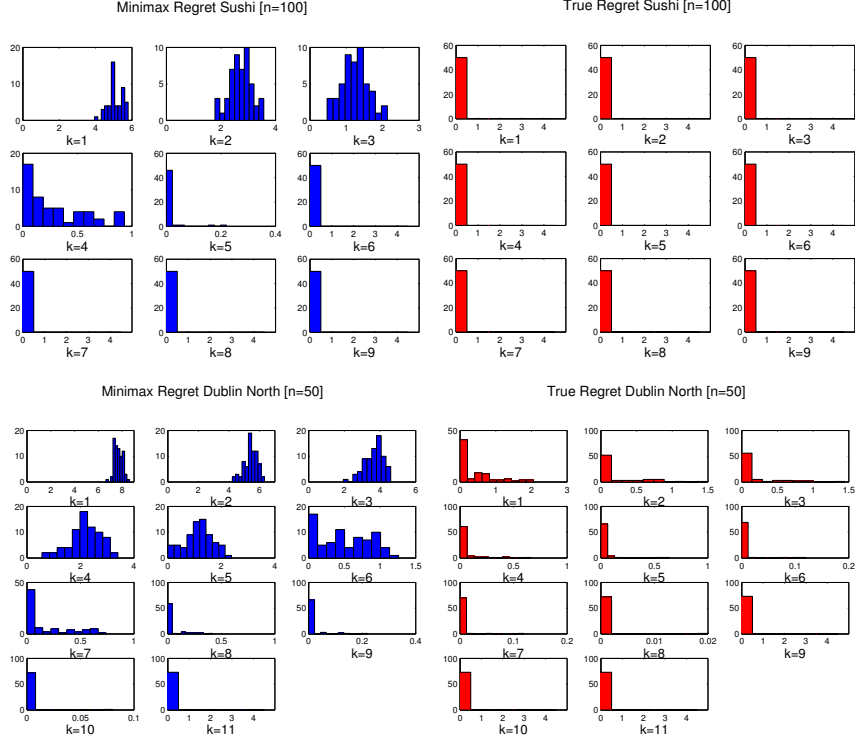


Fig. 4: Results on sushi rankings and Irish voting data.

True regret (see Fig. 2) is even more illuminating: with $\phi = 0.6$, the MMR solution *after only top-1 queries to each voter* is nearly always the true winner; and true regret never exceeds 2. Even for the uniform distribution with $\phi = 1$, true regret is surprisingly small: after top-2 queries, regret is less than 0.5 in 97/100 cases. As we increase the number of voters n , the MMR distribution becomes more concentrated around the mean (e.g., $n = 1000$), and often resembles a Gaussian. Roughly, this is because with Borda scoring, (normalized) MMR can be expressed as the average of independent functions of p_i through pairwise max regret $PMR_i(a_p^*, a') = \max_{v_i \in C(p_i)} B(v_i(a')) - B(v_i(a_p^*))$, where a' is the adversarial witness (see Eq. (1)).

Fig. 3 provides a summary of the above experiments, showing average MMR as a function of k , along with average true regret and several percentile bounds. As above, we see that a smaller ϕ requires a smaller k to guarantee low MMR. It also illustrates the desirable anytime property of MMR: regret drops significantly with the “first few candidates” and levels off before reaching zero. For example, with $m = 10, n = 100, \phi = 0.6$, top-3 queries reduce MMR to 0.8 per voter from the MMR of 9 obtained with no queries; but an additional 3 candidates (i.e., top-6 queries) are needed to reduce regret from 0.8 per voter to 0. If we fix $\phi = 0.6$ and increase the number of candidates m , the k required for small MMR decreases in relation to m : we see that for $m = 5, 10, 20$

we need top- k queries with $k = 3, 6, 8$, respectively, to reach MMR of zero. This is, of course, specific to the Mallows model.

Fig. 4 show histograms on two real-world data sets: *Sushi* [7] (10 alternatives and 5000 rankings) and *Dublin*, voting data from the Dublin North constituency in 2002 (12 candidates and 3662 rankings).⁴ With *Sushi*, we divided the 5000 rankings into 50 voting profile instances, each with $n = 100$ rankings, and plotted MMR histograms using the same protocol as in Fig. 1 and Fig. 2; similarly, *Dublin* was divided into 73 profiles each with $n = 50$. *Sushi* results suggest that with top-5 queries one can usually find a necessary winner; but top-4 queries are usually enough to obtain low MMR sufficient for such a low-stakes group decision (i.e., what sushi to order). True regret histograms show the minimax solution is almost always the true winner. With *Dublin*, top-5 queries virtually guarantee MMR of no more than 2 per voter; top-6, MMR of 1 per voter; and top-7, MMR of 0.5 per voter. True regret plots show minimax winner is either optimal or close to optimal in most profile instances.

6 Concluding Remarks

We have outlined a general framework for the design of multi-round elicitation protocols that are sensitive to tradeoffs between number of rounds of elicitation imposed on voters, the amount of information elicited per round, and the quality of the proposed winner. Our framework is probabilistic, allowing one to account for realistic distributions of voter preferences and profiles. We have formulated a probabilistic method for choosing the ideal threshold k for top- k elicitation in one-round protocols, and developed an empirical methodology that applies to *any voting rule* and *any preference distribution*. While the method can be used purely heuristically, our PAC-analysis provides our methodology with statistical guarantees. Experiments on random Mallows models, as well as real-world data sets (sushi preferences and Irish electoral data) demonstrate the practical viability and advantages of our empirical approach.

There are numerous opportunities for future research. We have dealt mainly with one-round elicitation of top- k candidates—developing algorithms for optimal multi-round instantiations of our framework is an important next step. Critically, we must deal with posterior distributions that are generally intractable, though GRIM-based techniques [9] may help. We are also interested in more flexible query classes such as batched pairwise comparisons. While the empirical framework is applicable to any preference distribution, we still wish to analyze the performance on additional distributions, including more flexible mixture models. On the theoretical side, we expect our PAC-analysis can be extended to different query classes and to multi-round protocols: we expect that probabilistic bounds on the amount of information required (e.g., k^* for top- k queries) will be significantly better than deterministic worst-case bounds [3] assuming, for example, a Mallows model. Bayesian approaches that assess candidate quality using expected regret rather than minimax regret are also of interest, especially in lower-stakes settings. We expect that combining expected regret and minimax regret might yield interesting solutions as well.

⁴ There are 43,942 ballots; 3662 are complete. See www.dublincountyreturningofficer.com

Acknowledgements: Thanks to Yann Chevaleyre, Jérôme Lang, and Nicolas Maudet for helpful discussions. This research was supported by NSERC.

References

1. Yann Chevaleyre, Ulle Endriss, Jérôme Lang, and Nicolas Maudet. A short introduction to computational social choice. In *Proceedings of the 33rd Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM-07)*, pages 51–69, Harrachov, Czech Republic, 2007.
2. Vincent Conitzer and Tuomas Sandholm. Vote elicitation: Complexity and strategy-proofness. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI-02)*, pages 392–397, Edmonton, 2002.
3. Vincent Conitzer and Tuomas Sandholm. Communication complexity of common voting rules. In *Proceedings of the Sixth ACM Conference on Electronic Commerce (EC’05)*, pages 78–87, Vancouver, 2005.
4. Jean-Paul Doignon, Aleksandar Pekec, and Michel Regenwetter. The repeated insertion model for rankings: Missing link between two subset choice models. *Psychometrika*, 69(1):33–54, 2004.
5. Wulf Gaertner. *A Primer in Social Choice Theory*. LSE Perspectives in Economic Analysis. Oxford University Press, USA, August 2006.
6. 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.
7. Toshihiro Kamishima, Hideto Kazawa, and Shotaro Akaho. Supervised ordering: An empirical survey. *Data Mining, IEEE International Conference on*, 0:673–676, 2005.
8. Jérôme Lang. Vote and aggregation in combinatorial domains with structured preferences. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, pages 1366–1371, Hyderabad, India, 2007.
9. Tyler Lu and Craig Boutilier. Learning Mallows models with pairwise preferences. In *Proceedings of the Twenty-eighth International Conference on Machine Learning (ICML-11)*, Bellevue, Washington, 2011.
10. 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)*, Barcelona, 2011. To appear.
11. Colin L. Mallows. Non-null ranking models. In *Biometrika*, volume 44, pages 114–130, 1957.
12. John I. Marden. *Analyzing and modeling rank data*. Chapman and Hall, 1995.
13. Thomas Brendan Murphy and Donal Martin. Mixtures of distance-based models for ranking data. *Computational Statistics and Data Analysis*, 41:645–655, January 2003.
14. Michel Regenwetter, Bernard Grofman, A. A. J. Marley, and Ilia Tsetlin. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge University Press, Cambridge, 2006.
15. Leslie G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
16. 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.