

Who Is in Your Top Three? Optimizing Learning in Elections with Many Candidates

Nikhil Garg, Lodewijk L. Gelauff, Sukolsak Sakshuwong, Ashish Goel
Stanford University
{nkgarg, lodewijk, sukolsak, ashishg}@stanford.edu

Abstract

Elections and opinion polls often have many candidates, with the aim to either rank the candidates or identify a small set of winners according to voters' preferences. In practice, voters do not provide a full ranking; instead, each voter provides their favorite K candidates, potentially in ranked order. The election organizer must choose K and an aggregation rule.

We provide a theoretical framework to make these choices. Each K -Approval or K -partial ranking mechanism (with a corresponding positional scoring rule) induces a *learning rate* for the speed at which the election recovers the asymptotic outcome. Given the voter choice distribution, the election planner can thus identify the rate optimal mechanism. Earlier work in this area provides coarse order-of-magnitude guarantees which are not sufficient to make such choices. Our framework further resolves questions of when randomizing between multiple mechanisms may improve learning for arbitrary voter noise models.

Finally, we use data from 5 large participatory budgeting elections that we organized across several US cities, along with other ranking data, to demonstrate the utility of our methods. In particular, we find that historically such elections have set K too low and that picking the right mechanism can be the difference between identifying the ultimate winner with only a 80% probability or a 99.9% probability after 400 voters.

1 Introduction

Elections and opinion polls with many candidates and multiple winners are common. In participatory budgeting (PB), for example, people directly determine a part of the government's budget (Alós-Ferrer and Granić 2012; Goel et al. 2016). These elections often contain many candidate projects (up to 70, cf. Gelauff et al. (2018)) and only a few thousand voters, with potentially millions of dollars on the line (Public Agenda 2016). Similarly, polls may compare tens of candidates and yet only sample hundreds of voters.

Unfortunately, the number of voters required to recover the asymptotic ranking or set of winners often scales, potentially exponentially, with the number of candidates (Caragiannis and Micha 2017). Thus with many candidates, it

is essential to use a voting mechanism that most efficiently elicits information from each voter.

In this work, we analyze *positional scoring rules* (de Borda 1781; Young 1975), mechanisms in which each position in each voter's personal ranking maps to a score given to the candidate that occupies that position. We focus on the special cases of such rules implied by K -Approval elicitation, in which each voter is asked to select their favorite K candidates, as they the most commonly used such mechanisms in practice. Section 3 formalizes our model. Then:

Section 4. For a given election, we show how the particular scoring rule used affects the rate at which the final outcome (asymptotic in the number of voters) is learned. These rates, based on large deviation bounds, extend and tighten the results of Caragiannis and Micha (2017), and are precise enough to determine, for example, which of 3-Approval and 4-Approval is better in a particular context. We focus on the goals of learning both a ranking over all candidates and identifying a subset of winners.

Section 5.1. Leveraging these rates, we study when randomization between scoring rules can improve learning, extending previous results to general positional scoring rules, the goal of selecting a set of winners, and arbitrary noise models. In particular, we find that randomizing between scoring rules can never speed up learning, for arbitrary noise models. This contrasts to the case when one is restricted to K -Approval mechanisms.

Section 5.2. For the Mallows model, we study how the optimal K in K -Approval scales with the noise parameter, the number of candidates, and the number of winners desired. We find that, in contrast to design choices made in practice, one should potentially ask voters to identify their favorite *half* of candidates, even if the goal is to identify a *single* winner.

Section 6. We apply our approach to experimental ballots attached to real participatory budgeting elections across several US cities, as well as other ranking data from a host of domains. We find that the exact mechanism used matters: in one setting, for example, asking voters to identify their favorite candidate results in only a 80% chance of identifying the best candidate after 400 voters, while

asking voters for their favorite 2 candidates identifies the same best candidate 99.9% of the time. Extending our theoretical insights, we find that, historically across elections, K has been set too low for effective learning. We further identify real-world examples in which randomization would have sped up learning.

Our work bridges a gap between coarse theoretical analyses of voting rules and the fine-grained design questions a practitioner wishes to answer. Proofs are in the Appendix.

2 Related Work

Our work is part of several strands of research on mechanisms that elicit peoples’ preferences. Aggregating voter rankings has a long history (de Borda 1781; marquis de Condorcet 1785; Copeland 1951; Kemeny 1959; Young 1988).

Learning properties of voting rules. Most related are works that study the learning properties of voting rules, assuming that a “true” ranking. One approach is to specify a noise model under which voter preferences are drawn (e.g., Mallows, Plackett-Luce) and then derive error rates by the number of voters for maximum likelihood or similar estimators under the model (Maystre and Grossglauser 2015; Zhao, Piech, and Xia 2016; Lu and Boutilier 2011; Guiver and Snelson 2009; Procaccia and Shah 2015; de Weerd, Gerding, and Stein 2016; Chierichetti and Kleinberg 2014).

Caragiannis, Procaccia, and Shah (2013) ask similar questions to us: under what voter noise models do certain voting rules asymptotically recover the true underlying ranking, and how quickly do they do so. They define a class of voting rules and voter noise models under which a “true” ranking of candidates is eventually recovered. They further show that for a subset of this class (that does not contain positional scoring rules) and under the Mallows model, only a number of voters that is logarithmic in the number of candidates is required, where each voter provides a full ranking. Lee et al. (2014) develop an algorithm that can approximate the Borda rule, given a number of comparisons by each voter that is logarithmic in the number of candidates.

Most similar is that of Caragiannis and Micha (2017). They show that under the Mallows model, K -Approval with any fixed K takes exponentially many voters (in the number of candidates) to recover the underlying ranking; on the other hand, K -approval with K chosen uniformly at random for each voter takes only a polynomial number of voters.

These works provide order estimates for the learning rate, *asymptotic in the number of candidates*; fine-grained differentiation between different rules or K -Approval mechanisms for a given election is not possible. We provide the latter and show that it matters.

Other approaches to comparing mechanisms. Many works take an axiomatic and computational approach, comparing mechanisms that may produce different outcomes even given asymptotically many votes (Fishburn and Gehrlein 1976; Fishburn 1978; Staring 1986; Tataru and Merlin 1997; Wiseman 2000; Ratliff 2003; Elkind et al. 2017; Aziz et al. 2015; Caragiannis et al. 2017; Aziz et al. 2017; Lackner and Skowron 2018a; 2018b; Faliszewski and

Talmon 2018). Caragiannis et al. (2019) for example show how to find a scoring rule that most agrees with a given partial ground truth ranking. In contrast, we compare mechanisms’ learning rates under a condition (formalized in Section 3.2) in which they produce the same asymptotic outcome.

Benade et al. (2018) and Gelauff et al. (2018) experimentally compare different mechanisms across several dimensions, including ease of use and consistency with another mechanism; the latter leverages data from a participatory budgeting election at a university.

Large deviation analysis of elicitation mechanisms.

Theoretically, we leverage large deviation rates and Chernoff bounds to derive how quickly a given scoring rule learns its outcome; see work of Dembo and Zeitouni (2010) for an introduction to large deviations. This work is thus conceptually similar to work on elicitation design for rating systems (Garg and Johari 2018; 2019). In those works, the authors derive large deviation-based learning rates that depend on the questions that are asked to buyers as they review an item, where the goal is to accurately rank items; they further run an experiment on an online labor platform. In that setting, however, buyers rate a *single* item, and mechanisms are distinct based on the *behavior* they induce; in this work, voters see all the candidates and provide a partial ordering, and different designs (e.g., 3-Approval vs 4-Approval) *constrain* the types of orderings voters can provide.

3 Model

We now present our model and a condition under which different positional scoring rules induce the same asymptotic outcome.

3.1 Model primitives

We begin with the model primitives: candidates and voters, the election goal, and elicitation and aggregation.

Candidates and Voters. There is a set of M candidates $C = \{1, \dots, M\}$, typically indexed by $i, j \in C$. There are N voters $V = \{1, \dots, N\}$. Each voter $v \in V$ has a strict ranking of candidates σ_v , drawn independently and identically from probability mass function over strict rankings $F(\sigma)$. Let $i \succ_\sigma j$ denote that i is preferred over j in σ , and $\sigma(i) = k$ denote that candidate i is in the k th position in σ .

A special case for F is the *Mallows* model (Mallows 1957), in which there is a “true” societal preference σ^* from which each voter’s ranking is a noisy sample. In particular,

$$F_{\text{Mallows}}(\sigma) \propto \phi^{d(\sigma, \sigma^*)}$$

Where $d(\sigma, \sigma^*)$ is the Kendall’s τ distance between rankings σ, σ^* , and $\phi \in [0, 1]$ is the noise parameter: the smaller it is, the more concentrated F is around σ^* .

Election goal. We assume that the *goal* G is to divide the candidates into T disjoint, ordered *tiers* $G = \{C_1, \dots, C_T\}$, such that $C = \cup_{t=1}^T C_t$, where candidate $i \in C_s$ is deemed societally preferable over $j \in C_t$ if $s < t$. The size of each tier is fixed before the election. For example, recovering a strict ranking over all candidates corresponds to

$G = \{C_1, \dots, C_M\}$, where $|C_t| = 1$. Alternatively, identifying a set of W winners, without distinguishing amongst the winners, corresponds to $G = \{C_1, C_2\}$, with $|C_1| = W$.

In the main text and especially the empirics, we will focus on the task of selecting W winners as it is the most common task in practice. However, this general notation allows comparison of the learning properties of different settings, and for example ask how much more expensive is it (in terms of the number of voters needed) to identify a strict ranking as opposed to just a set of winners.

Elicitation and Aggregation. Voters vote using an elicitation mechanism. Their votes are then aggregated using a positional scoring rule, parameterized as $\beta : \{1, \dots, M\} \mapsto \mathbb{R}$. We consider the following mechanisms:

K -Ranking Voter v ranks her favorite K candidates, i.e., reveals $\{(i, \sigma_v(i)) : \sigma_v(i) \leq K\}$. Candidate i then receives a score $s_{iv} = \beta(\sigma_v(i))$ if ranked, 0 otherwise. For example, $\beta(k) = M - k$ for the Borda count.¹

K -Approval Voter v selects her favorite K candidates, i.e., reveals $\{i : \sigma_v(i) \leq K\}$. A candidate receives a score $s_{iv} = 1$ for being selected, 0 otherwise.

β encodes both elicitation and aggregation. For example, K -Approval is equivalent to K -ranking with score function $\beta(k) = \mathbb{I}[k \leq K]$. Furthermore, note that given K -ranking data, one can simulate K' -ranking elicitation for $K' \leq K$ with a β s.t. $\beta(k) = 0$ for $k > K'$.

The scoring rule β is a design choice made by the election organizer, and so we will refer to β as the election’s *design*. We restrict ourselves to non-constant, non-increasing scoring rules, i.e., $\beta \in \mathcal{B} = \{\beta : \forall k < \ell \in 1, \dots, M, \beta(k) \geq \beta(\ell), \text{ and } \exists k < \ell, \beta(k) > \beta(\ell)\}$.

Outcome. After N voters, candidate i ’s cumulative score is $s_i^N = \frac{1}{N} \sum_{v=1}^N s_{iv}$. Candidates are ranked in descending order of score, to form ranking σ^N , with ties broken uniformly at random. We denote the *outcome* after N voters, corresponding to the goal G , as $O^N(M, F, \beta, G)$. For example, for the goal of selecting W winners, $O^N(M, F, \beta, G)$ is simply the top W candidates in σ^N . When (M, F, β, G) is clear from context, we will refer to the outcome as O^N .

As the number of voters $N \rightarrow \infty$, candidate scores $s_i^N \rightarrow \mathbb{E}_F[s_{iv}] \triangleq s_i$ by the law of large numbers; when such expected scores are distinct, i.e., $s_i \neq s_j$ for $i \neq j$, then $\sigma^N \rightarrow \sigma^*$ for some ranking σ^* . However, note that there may exist an asymptotic outcome $O^N \rightarrow O^*$ even without an asymptotic ranking $\sigma^N \rightarrow \sigma^*$, as long as expected scores s_i and goal G are such that candidates with identical expected scores are sorted into the same tier.

3.2 Asymptotic design invariance

The asymptotic outcome O^* of an election may vary with the scoring rule β . For example, there may be a different winner if voters are asked to identify their favorite two candidates than if they identify their single favorite candidate, if the winner in the latter case is a polarizing candidate. As an

¹In Borda, candidates not ranked receive a score $(M - K - 1)/2$, consistent with assuming they are all tied in position $(K + 1)$.

axiomatic comparison between outcomes is out of the scope of this paper, we restrict our attention to cases where all “reasonable” choices of different β asymptotically result in the same outcome (where “reasonable” corresponds to the set of scoring rules \mathcal{B} defined above).

Definition 1. A setting (M, F) is asymptotically design-invariant for goal G if any reasonable β induces the same outcome asymptotically. $\exists O^* : \forall \beta \in \mathcal{B}$,

$$\lim_{N \rightarrow \infty} O^N(M, F, \beta, G) = O^*, \text{ with probability } 1$$

Such design invariance only occurs under a fairly strong condition on the voter preference distribution: that the candidates can be separated into tiers (according to goal G) such that candidates in higher tiers are *strictly* more likely to be ranked by a voter in the top k positions, for all $k < M$, than are candidates in lower tiers.

Proposition 1. A setting (M, F) for goal G is asymptotically design-invariant if and only if there exist candidate tiers $O^* = \{C_1^* \dots C_T^*\}$ (corresponding to G) s.t. $\forall s < t : i \in C_s^*, j \in C_t^* \implies \Pr_F(\sigma_v(i) \leq k) > \Pr_F(\sigma_v(j) \leq k)$, $\forall k \in \{1 \dots M - 1\}$.

Note that this condition is stronger than stochastic dominance as the inequality is strict for *every* position k .

This proposition connects to Caragiannis, Procaccia, and Shah (2013) as follows: they prove that many rules (including all positional scoring rules and the Bucklin rule) asymptotically recover the base ranking σ^* of a generalization of the Mallows model in which the probability $F(\sigma)$ of a ranking σ is monotonic in the distance $d(\sigma, \sigma^*)$, where distance function d is itself in some general class that contains the Kendall’s τ distance. Their results directly imply that such noise models, including the standard Mallows model, are asymptotically design-invariant for any goal G .

However, for goals G where recovering a full ranking is unnecessary, the condition in Proposition 1 is weaker than the assumptions of Caragiannis, Procaccia, and Shah (2013); there need not even be a single base ranking σ^* . For example, when G such that we wish to select a set of W winners, F corresponding to a mixture of Mallows models – with all possible permutations of the W candidates in the top W positions in the base rankings – would still be design-invariant. Constructing a general class of ranking noise models that satisfies this property is an avenue for future work.

Assuming asymptotic design-invariance on voter preferences F may seem restrictive. However, absent axioms – that are precise enough for design purposes – to prefer one scoring rule β over another, the assumption allows us to proceed in a principled manner. We believe it is unlikely that such precise, satisfactory axioms exist generally. In the Appendix, we provide a simple example (similar to that of Starling (1986)) where 1-Approval and 2-Approval select disjoint sets of 2-Winners, and such examples can be adapted more generally to selecting W winners from either K -Approvals or K' -Approvals. In participatory budgeting with the goal of identifying 6-10 winning projects out of over twenty projects, it is unclear whether there is a principled reason to prefer 4-Approval over 8-Approval. However, such axioms would be an interesting avenue for future work.

Furthermore, in Section 6.2 we show that design invariance is often approximately satisfied in practice, especially for identifying a small set of winners, using data from a wide range of participatory budgeting and other elections.

4 Learning Rates and Optimal Design

Different elicitation and aggregation mechanisms may take different amounts of voters to learn the asymptotic outcome. For example, suppose we want to identify the worst candidate out of 100, where the voter’s rankings are drawn from a Mallows model with $\phi > 0$. Then, asking each voter to identify their single favorite candidate will eventually identify the worst candidate, but after many more voters than if we ask each voter to identify their least favorite candidate. We make such learning rates precise in this section. Our results in this section extend those of Caragiannis and Micha (2017) as discussed above, both to arbitrary positional scoring rules and by providing tighter bounds for how a scoring rule affects the convergence rate. These rates are precise enough to *design* scoring rules, for example comparing 4-Approval and 8-Approval in the above example.

4.1 Learning rates

We begin by deriving rates for how quickly a given positional scoring rule β learns its asymptotic outcome O^* (given it exists), as a function of the voter preference model F . In particular, we use *large deviation rates* at which a scoring rule learns (Dembo and Zeitouni 2010).

Definition 2. Consider a sequence $\{A_N \geq 0\}_{N \in \mathbb{N}}$, where $A_N \rightarrow 0$. Value $r > 0$ is the large deviation rate for A_N if

$$r = - \lim_{N \rightarrow \infty} \frac{1}{N} \log A_N$$

When $r > 0$ exists, $A_N \rightarrow 0$ exponentially fast, with exponent r asymptotically, i.e., A_N is $e^{-rN \pm o(N)}$. These rates provide us both upper and lower bounds for the probability of an error or the number of errors in an outcome after N voters, up to polynomial factors. In particular, in the propositions below, we will calculate the large deviation rate of errors in the outcome. We will also then provide (loose) upper bounds for such errors after N voters that hold without any missing polynomial factors, for any N . These upper bounds are equivalent to Chernoff bounds.

The particular forms for these rates, derived below for general noise models F , may seem complex. However, they are useful both for theoreticians and practitioners. For example, in Section 5.1, we use the structure of such rates to resolve open questions regarding when randomization between mechanisms can help learn the outcome from votes drawn from an *arbitrary* noise model. In Section 6, we show that learning rates – even when empirically calculated – reflect the true behavior of errors in real elections with a small number of voters; we then use empirically calculated learning rates to draw design insights across elections.

Rates for separating two candidates. We now derive the large deviation learning rates for recovering the true ordering between a pair of candidates i, j , given noise model F .

These rates will directly translate to the learning rate for the overall election, given some goal G .

Proposition 2. Fix scoring rule $\beta \in \mathcal{B}$, voter distribution F , and consider candidates i, j such that $s_i > s_j$. Then, the probability of making a mistake in ranking these two candidates after N voters, $\Pr(\sigma^N(i) > \sigma^N(j))$, goes to zero with large deviation rate

$$r_{ij}(\beta) = - \inf_{z \in \mathbb{R}} \log \mathbb{E}_F [\exp(z(\beta(\sigma_v(i)) - \beta(\sigma_v(j))))]$$

Further, the following upper bound holds for any N .

$$\Pr(\sigma^N(i) > \sigma^N(j)) \leq \exp(-r_{ij}(\beta)N)$$

The proof follows directly from writing a random variable for the event of making a mistake after N voters and then applying known large deviation rates. This simplicity emerges because positional scoring rules are additive across voters.

The proposition establishes that – for a fixed number of candidates M and voter noise model F – the probability of making a mistake on any single pair of candidates i, j decreases exponentially with the number of voters, at a rate governed by the scoring rule β and the candidates’ relative probabilities of appearing at each position of a voter’s preference ranking. The rate $r_{ij}(\beta)$ is non-negative, and larger values correspond to faster learning of the relative ranking of i, j . Note that for notational convenience, we suppress F in the argument for the rate.

For general β , we cannot find a closed form for $r_{ij}(\beta)$. However, the structure of this rate, in particular that of the argument in the $\log(\cdot)$, will directly let us show that randomization cannot help learning outcomes among positional scoring rules, for *arbitrary* noise models F .

For K -Approval voting, further, the rate simplifies.

Proposition 3. Consider β consistent with K -Approval voting for some fixed K , and candidates i, j such that $s_i > s_j$. Then the large deviation rate $r_{ij}(\beta)$ in Proposition 2 is

$$r_{ij}(K) = - \log \left(2\sqrt{t_{ij}^i(K)t_{ij}^j(K)} + 1 - t_{ij}^i(K) - t_{ij}^j(K) \right)$$

Where $t_{ij}^i(K) \triangleq \Pr_F(\sigma_v(i) \leq K, \sigma_v(j) > K)$, i.e., the probability that a voter approves i but not j .

The proof follows directly from the structure of β for K -Approval, $\beta(k) = \mathbb{I}[k \leq K]$; for each pair of candidates, the sufficient statistics are how often each candidate appears in a voter’s top K list but the other candidate does not.

We overload notation and use K directly in the argument for $r_{ij}(K)$. This rate function $r_{ij}(K)$ is convex in the probabilities $t_{ij}^i(K), t_{ij}^j(K)$; this fact will let us show that randomization, even among K -Approval mechanisms, cannot help learning the relationship of any pair of candidates.

Rates for learning the outcome. In general, the rates at which one learns each pair of candidates immediately translate to rates for learning the entire outcome O^* .

Proposition 4. Consider goal G and $\beta \in \mathcal{B}$ such that $O^N \rightarrow O^*$. Let Q^N be the expected number of errors in

the outcome after N voters, $\sum_{i \in C_s^*, j \in C_t^*, s < t} \Pr(\sigma^N(i) > \sigma^N(j))$. Then Q^N goes to zero with large deviation rate

$$r(\beta) = \min_{i \in C_s^*, j \in C_t^*, s < t} r_{ij}(\beta)$$

Further, the following upper bound holds for any N .

$$Q^N \leq M^2 \exp(-rN)$$

The large deviation rate $r(\beta)$ thus provides a tight characterization for how many voters it takes to (with high confidence) recover the asymptotic outcome of an election. Note that the goal plays an important role: for selecting W winners, for example, it is not important to learn the exact relationship among candidates $\{1, \dots, W\}$, speeding up outcome learning. Design β also matters; e.g., even amongst approval voting mechanisms, $K = 1$ vs $K = 5$ will produce substantially different $t_{ij}^i(K)$. To derive learning rates for K -Approval for any given noise model or using real-world data, one simply needs to calculate these values. We do so numerically for the Mallows model and empirically with real world data in Sections 5.2 and 6, respectively.

4.2 Optimal design and discussion

Now that we can quantify how quickly a given scoring rule β learns its asymptotic outcome, we apply our framework to *designing* elections, i.e., choosing an optimal scoring rule β . For the rest of this work, we assume that the setting (M, F) is asymptotically design-invariant for the goal G , i.e., there exists an outcome that is asymptotically induced by every reasonable scoring rule. Then, the design of an election β only affects the *rate* at which the election converges to the asymptotic outcome O^* , as calculated above. With no other constraints, then, the design challenge is simple: find the rate optimal β .

Definition 3. A scoring rule $\beta^* \in \mathcal{B}$ is rate optimal if it maximizes the rate in Proposition 4. K^* -Approval is Approval rate optimal if it maximizes the rate among K -Approval mechanisms.

Rate optimal designs β learn the outcome faster than others in the number of voters, and so are preferable to other designs. What influences how quickly a design β learns? $\mathbb{E}_F[\exp(z(\beta(\sigma_v(i)) - \beta(\sigma_v(j))))]$ must be small (near zero) for negative z , and so $\beta(k) - \beta(k')$ must be large when $\Pr(\sigma_v(i) = k, \sigma_v(j) = k')$ is large. In other words, a scoring rule must reward a candidate achieving a position in a voter’s ranking that is only achieved by asymptotically high-ranking candidates. For example, if it is common for worse candidates to be ranked second in a given voter’s ranking but not to be ranked first, then $\beta(1) \gg \beta(2)$ would be beneficial.

Note that finding such designs requires knowledge of the voter noise model F , which in many settings may not be available before the election. However, next in Sections 5 and 6, we show that there are valuable insights that apply *across* elections, including how our approach has informed participatory budgeting deployments.

5 Theoretical Design Insights

The learning rates derived in the previous section provide election design insights, even before our approach is applied to real-world data. In particular, in this section, we first extend the previous literature on the (potential) benefits of randomizing between mechanisms. Then, we study the task of selecting W winners using K -Approval voting.

5.1 When does randomization help?

We now consider the question of whether *randomizing* between mechanisms in an election may speed up learning. By randomization, we mean: consider a set of scoring rules $B = \{\beta_1, \dots, \beta_P\} \subseteq \mathcal{B}$; elicitation and aggregation for a given voter is done according to a scoring rule picked at random from B , where β_p is selected with probability d_p .

Note that the learning rate of such randomized schemes can be calculated as before, by summing across β_p inside the $\mathbb{E}[\cdot]$ of $r_{ij}(\beta)$ or – for B consisting only of K -Approval votes – directly through the resulting probability that the voter approves i but not j . We use $r_{ij}(B, D), r(B, D)$ to denote the candidate pairwise and overall outcome learning rates, respectively, for randomized mechanism (B, D) , where $B = \{\beta_1, \dots, \beta_P\} \subseteq \mathcal{B}$ and $D = \{d_1, \dots, d_P\}$.

It is known that in some settings randomization improves learning, asymptotically in the number of candidates. Caragiannis and Micha (2017) provide an example in which randomizing uniformly between all possible K -Approval mechanisms outperforms any static K -Approval elicitation, when the goal is to rank all the candidates. Their insight is that, under the Mallows model and under a fixed K , either the first two candidates will be hard to distinguish from each other, or the last two will, and randomizing between mechanisms balances learning each pair.

We now study randomization for the goal of selecting W winners and for arbitrary positional scoring rules and voter noise models. Our first result is that randomizing between scoring rules does not help, for *any* voter noise model, in contrast to the case when restricted to approval votes.

Theorem 1. Randomization does not improve the outcome learning rate for any asymptotically design-invariant noise model F or goal G . For any randomized scoring rule mechanism (B, D) , where $B \subset \mathcal{B}$, for any F, G , the scoring rule $\beta^*(k) = \sum_p d_p \beta_p(k)$ satisfies $r(\beta^*) \geq r(B, D)$.

The result follows from the fact that $\mathbb{E}_F[\exp(z(\beta(\sigma_v(i)) - \beta(\sigma_v(j))))]$ is convex in $\beta(k)$, for all i, j, z, F . Then, given a randomization over β_1, \dots, β_P , we can increase $-\inf_z \log(\cdot)$ by decreasing its argument, by instead using the static scoring rule defined by the corresponding convex combination of β_1, \dots, β_P . Note that such a negative result cannot be obtained via analysis that is asymptotic in the number of candidates; we need learning rates for a given election.

Next, we further refine the result of Caragiannis and Micha (2017), by showing that the “pivotal pair” feature of their example – where different pairs of candidates dominate the learning rate for different mechanisms – is key. In particular, our next result establishes, again for any noise model,

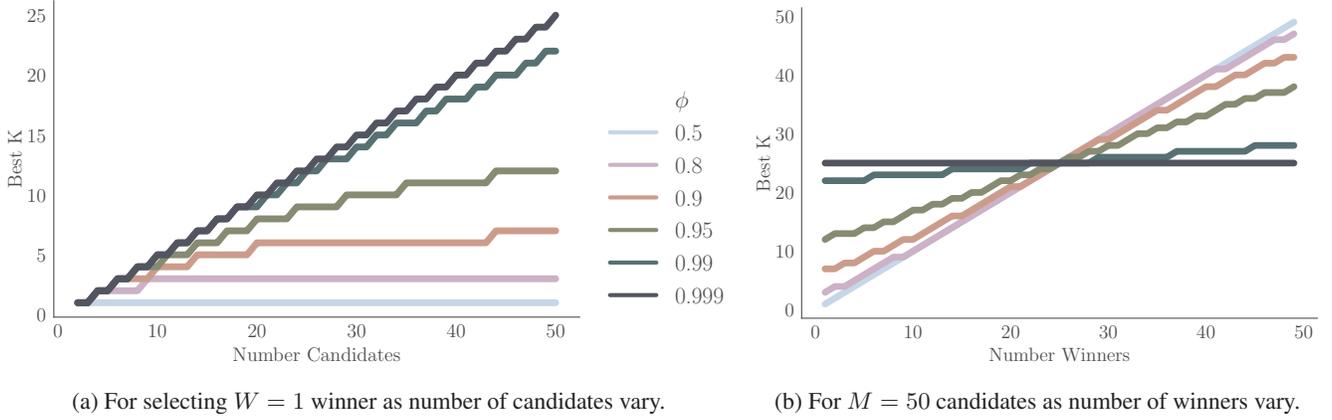


Figure 1: K -Approval rate optimal mechanism for the Mallows model as ϕ , number of candidates, and number of winners vary.

that randomization amongst K -Approval mechanisms cannot help separate any given pair of candidates.

Theorem 2. *Randomization amongst K -Approval mechanisms does not improve the learning rate for separating a given pair of candidates i, j for any asymptotically design-invariant noise model F or goal G . For any randomized K -Approval mechanism (B, D) , where $\beta_p \in B$ corresponds to p -Approval, for any F, G , there exists a mechanism K_{ij}^* -Approval such that $r_{ij}(K_{ij}^*) \geq r_{ij}(B, D)$.*

The proof relies on the pairwise rate function $r_{ij}(K)$ being convex in the approval probabilities $t_{ij}^i(K), t_{ij}^j(K)$.

This theorem directly implies that, for the Mallows model, randomization among K -Approval voting cannot speed up learning when the goal is to identify a set of W winners, as opposed to when the goal is to rank.

Corollary 1. *Randomization among K -Approval mechanisms does not improve the learning rate for selecting W winners from the Mallows model. For any randomized K -Approval mechanism (B, D) , where $\beta_p \in B$ corresponds to p -Approval, for selecting W winners from the Mallows model, there exists an Approval rate optimal mechanism K^* -Approval such that $r(K^*) \geq r(B, D)$.*

The proof simply notes that under the Mallows model with this goal, the candidate pair $W, W+1$ (when candidates are indexed according to reference distribution σ^*) is pivotal regardless of the K -Approval mechanism used. This corollary does not extend to arbitrary noise models, where randomization amongst K -approval mechanisms may improve the learning rate.

Theorem 3. *Randomization among K -Approval mechanisms may improve the learning rate for the goal of selecting W winners. There exist asymptotically design-invariant settings (M, F) for the goal of selecting W winners such that a randomized K -Approval mechanism (B, D) , where $\beta_p \in B$ corresponds to p -Approval, satisfies*

$$r(B, D) > \max_K r(K)$$

We prove the result two ways: (1) we construct an example in which candidate h is asymptotically selected, and

candidates i, j are not. Which of $h \succ i$ or $h \succ j$ is the pivotal pair (determines the overall rate function) depends on the K -Approval mechanism used, and randomizing between two mechanisms improves the overall rate; (2) perhaps more interestingly, we find many examples in our real PB elections and other ranking data in which randomization would have sped up learning for the task of selecting a set of winning candidates (see Section 6.4).

5.2 K -Approval for selecting W winners

One of the most common voting settings is identifying a set of W winners using K -Approval, whether in representative democracy elections (typically $K = W = 1$), polling for such elections (where the goal often is to identify the top few candidates out of many, especially in primary races), or crowd-sourcing labels (where one wants one or a few labels for an item out of many possible ones). Here, we study how to design such elections, i.e., how to choose the best K , i.e., the one that maximizes the learning rate. For simplicity, we work with the Mallows model, extending the resulting insights to real-world data in the next section.

Recall that in a Mallows model, each voter’s ranking is a noisy sample from a reference distribution σ^* . With this symmetric model, one may believe that setting $K = W$ is always optimal. For example, when noise parameter $\phi = 0$ and so each voter’s ranking is exactly σ^* , $K = W$ is optimal; in fact, any other design $K \neq W$ fails to correctly identify the set of winners even asymptotically: it would not distinguish among the first K candidates in σ^* or among the last $M - K$ candidates. However, our next result establishes that the cases with $\phi > 0$ are different.

Theorem 4. *Under the Mallows model and the goal of selecting W winners, W -Approval may not be Approval rate optimal.*

We prove the theorem by example. To find this example and to generate the plots discussed next, we use an efficient dynamic program to exactly calculate the joint distributions of the locations $\sigma_v(i), \sigma_v(j)$ of pairs of candidates i, j in a voter’s ranking, given the Mallows noise parameter; we can then directly calculate $t_{ij}^i(K), t_{ij}^j(K)$ and thus

the learning rate for each K -Approval mechanism. This program leverages Mallows repeated insertion probabilities (Lu and Boutilier 2014; Diaconis 1988) and may be of independent interest for numerical analyses of the Mallows model.

Numerical analysis. We now numerically analyze, for the Mallows model, how the Approval rate approval K -Approval mechanism varies with the Mallows noise parameter ϕ , the number of candidates M , and the number of winners W . Recall that the Mallows model is asymptotically design invariant, so different mechanisms only differ in *how quickly* they learn the asymptotic outcome.

In Figure 1a, the goal is to select $W = 1$ winner, and ϕ and M are varied. With low noise, $\phi \lesssim .5$, it is rate optimal to use 1-Approval, i.e., ask each voter to select their favorite candidate, regardless of how many candidates there are. However, with higher noise ϕ , as the number of candidates in the election increases, so does the K in the optimal K -Approval mechanism. For $\phi = .999$, $M = 50$, for example, it is best to ask each voter to select their favorite 25 candidates, even if the task is to identify the single best candidate according to the reference distribution σ^* .

Similarly, Figure 1b shows how the rate optimal K -Approval mechanism changes with the number of winners desired and the noise parameter, fixing the number of candidates at $M = 50$. Again with high noise, it is best to ask voters to identify their favorite half of candidates, regardless of how many winners need to be identified. With low noise, however, W -Approval is optimal to select W winners.

Overall, the analysis suggests that with higher noise in the voter model, one should tend toward asking voters to rank their favorite half of candidates, regardless of M and W .

The high-noise setting may seem unrealistic; however, as we will see in the next section, which K -Approval mechanism is rate optimal in practice often scales like the high noise settings, consistent with the idea that voting distributions in practice do not look like they are drawn from a low-noise Mallows model. We now turn to such empirical analyses.

6 Empirics and PB deployments

We now apply our insights to practice. We focus on K -Approval voting, as opposed to general scoring rules. This section is organized as follows. In Section 6.1, we describe our data sources. We validate our model in Section 6.2; first, we demonstrate that large deviation rates effectively describe how quickly various mechanisms learn; next, we show that in practice voter noise models are approximately design invariant. In Section 6.3, we show that the insights from Section 5.2 regarding optimal approval mechanisms extend to practice. Finally in Section 6.4 we note that we find many examples in practice where randomizing between K -Approval mechanisms improves learning.

6.1 Data description

We leverage two data sources (detailed dataset information is in Appendix Table 1). First, we have partnered with dozens of local governments to help run participatory budgeting (PB) elections in the last five years. These elections

have used a variety of methods, primarily K -Approval; our data in this work comes from 5 elections where K -Ranking was used, including 3 recent elections where $K = 10$. This data is particularly useful as PB is among the most common types of elections with many candidates and several winners, with several theoretical analyses (Goel et al. 2016; Garg et al. 2019; Freeman et al. 2019).

Second, we use data available on PrefLib (Mattei and Walsh 2013; O’Neill 2013; Regenwetter et al. 2007; 2008; Popov, Popova, and Regenwetter 2014), limiting ourselves to 28 elections with at least 5 candidates and 700 voters who provided full rankings. This ranking data spans many domains, from people’s sushi preferences to Glasgow City Council elections. This domain breadth supports the broad applicability of the design insights explored in this section.

We focus on ranking data to be able to simulate counterfactuals for the same election: with K -Ranking data, we can simulate what would have occurred with any K' -Approval elicitation mechanism, for $K' \leq K$ (assuming no behavioral quirks). With approval data, on the other hand, one cannot compare the mechanism to any other for that given election.

One challenge is that ranking many candidates is onerous, and so voters rank at most 14 candidates in our dataset. For the data we use from on PrefLib, full rankings (rankings up to the number of candidates) are available. In the PB elections in our partner cities, typically each voter ranks or selects her favorite $K \ll M$ candidates.

6.2 Model validation

Our model and design approach has two components that must be validated: (1) that learning rates can effectively be used to compare different mechanisms, and (2) that design invariance (approximately) holds in practice.

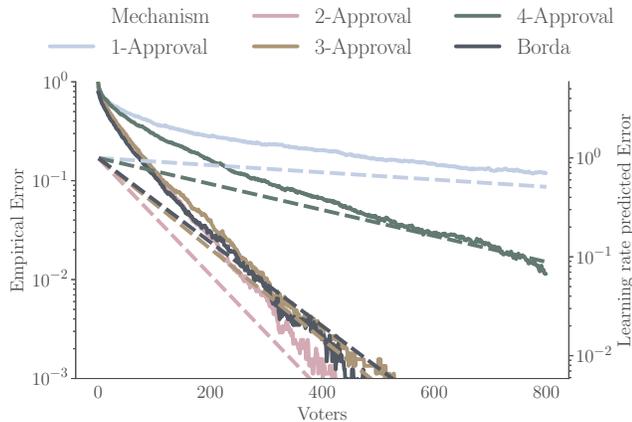
Large deviation rates as effective proxies for learning.

We now confirm that, for a given election, empirically calculated large deviation learning rates are effective proxies for the rate at which the error in recovering the asymptotic output decreases as the number of voters increases (even though large deviation learning rates are only asymptotically valid in the number of voters). As examples, we first identify three elections and goals for which many of the potential K -Approval mechanisms return exactly the same asymptotic outcome. Then, we bootstrap voters from the available data of voters and empirically calculate the errors made in identifying the winning set of candidates. We further calculate the large deviation learning rates for these mechanisms, using F implied by the voting data and the formula in Proposition 4.²

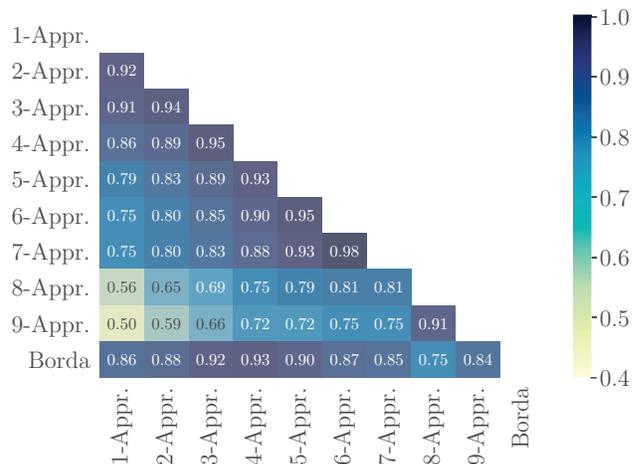
Figure 2a shows the resulting errors over time for one such election where 4-Rankings are available. We further plot e^{-rN} for each mechanism, i.e., the error over time implied by the learning rate (up to polynomial factors). This plot, along with Appendix Figure 3, yields several insights:

- (1) The mechanism matters: when selecting 1 winner from the election in Figure 2a after 400 votes, there is 20% chance of not picking the ultimate winner if 1-Approval

²Given an empirical \hat{F} , learning rates can be numerically calculated: the $\inf_z [\cdot]$ is a convex minimization problem.



(a) Boston 2016 PB election, selecting 1 winner: Average empirical bootstrapped error – i.e., fraction of times the asymptotic winner is selected (solid lines, left axis), compared to such errors over time implied by the (empirically calculated) learning rates – i.e., e^{-rN} (dashed lines, right axis). The right axis is a vertically shifted (in log scale) version of the left axis, reflecting that the learning rate errors are asymptotically valid up to polynomial factors. All mechanisms return the same winner when all votes are counted. “Borda” is the Borda count for the 4 candidates ranked, and all others are assumed to be tied at rank 5 for each voter.



(b) Approximate design invariance across elections. For the task of selecting $W = 4$ winners, this plot shows the average overlap in the top 4 candidates identified by different mechanisms across all the elections in our dataset, if all voters with complete rankings are counted. For example, of the top 4 candidates identified by 1-Approval across elections, 92% are also identified as top 4 candidates by 2-Approval. For each K -Approval mechanism, we include all elections where there were at least $K + 1$ candidates.

Figure 2: Validating model: comparing learning rates to empirical error, and showing approximate design invariance.

is used. With 2 or 3-Approval, this number is 0.1%. The winner appears often in a voter’s top two or three positions (but not necessarily first), while the ultimate second place candidate often falls outside the top three. Scoring rules that reward top three placements thus perform well.

- (2) The learning rates effectively capture the behavior of the empirical error: both comparatively across mechanisms, as well as the asymptotic rate (slope of the line in log scale). This property enables use of large deviation learning rates as proxies for learning even in elections with a small number of voters.
- (3) Ranking K candidates rather than selecting K candidates is more onerous for voters. However, it does not always provide more information in terms of learning rates, as in the examples in Appendix Figure 3.

Design invariance in practice. Design invariance does not strictly hold in any election in our dataset (as expected as the condition is strong). However, it *approximately* holds. Similar mechanisms produce the same asymptotic outcome for many tasks. Figure 2b shows, for example, the average overlap across elections in the top 4 candidates identified by each mechanism. (Appendix Figure 4 shows the same plot for the top 1 and 3 candidates, as well as the average Kendall’s τ rank correlation between the full rankings identified by different mechanisms). Furthermore, we find many elections and goals where most mechanisms return the same asymptotic answer, as in the elections we leverage for the plots showing learning rates are effective proxies. This rela-

tive consistency, especially for similar mechanisms, enables us to compare different mechanisms by their learning rates.

6.3 K -Approval for selecting W winners

In Section 5.2, we showed for the Mallows model how the rate optimal K -Approval mechanism changes with the noise parameter ϕ , the number of candidates, and the number of winners. We now show this scaling in practice.

For every election in our dataset, we find the Approval rate optimal mechanism (among K we can simulate) for every goal of selecting W winners, for $1 \leq W \leq M$. We then run a regression across all the elections for which K is rate optimal, versus the number of winners desired and the number of candidates; see Table 2 in the Appendix for the regression table. While there is some variation across elections, the number of candidates and winners proves a reasonable metric across elections for the rate approval K -Approval mechanism ($R^2 \approx .27$).

The regression confirms the idea that in practice, one should regularize toward asking voters to choose their favorite half the candidates. For picking a small subset of winners $W \approx 4$ out of more than 10 candidates, for example, one should ask voters to provide their favorite $K \approx 6$ candidates, with $K > W$. This suggestion directly counters common practice. In the PB elections that we have helped run, for example, 4 or 5-Approval is most typical, even though ultimately 6-10 projects may be funded (out of ≈ 15 -20).

Then, in Figure 5 in the Appendix, we plot the line induced from the regression coefficients with the Mallows rate optimal lines, for $M \leq 10$ candidates. Comparing to the

rate optimal mechanisms for the Mallows model with various ϕ (within the candidate range for which we have empirical data), we find that empirical data behaves most closely to a Mallows model with noise parameter $\phi \in [.8, .9]$. (We are not claiming that empirical data is drawn from a Mallows model; it most certainly is not, with factors such as polarizing projects important in practice). This coarse comparison provides an approximate expected scaling behavior for elections with many candidates.

6.4 Randomization in practice

We find 16 examples in which randomizing between two K -Approval mechanisms leads to faster learning than using either mechanism separately, including 8 examples where such randomization beats the Approval rate optimal mechanism. Table 3 in the Appendix contains details.

7 Discussion

We show that in elections with many candidates, the elicitation mechanism and corresponding scoring rule used affect how quickly the final outcome is learned. The learning speed differential between mechanisms can be the difference between identifying the ultimate winner with only a 80% probability or a 99.9% probability after 400 voters, for example. We then provide design decisions that emerge when our framework is applied to data from real elections. When using K -Approval to select a small number of W winners, for example, it is often better to ask voters to identify their favorite $K > W$ candidates. The insights from this work should be applicable in a variety of such settings, from elections to crowdsourcing labeling tasks.

There are several important, open research avenues. Most importantly, in real elections maximizing the rate at which the final outcome is identified is not the only goal, and future work should seek to balance such multiple objectives.

For example, there may be axiomatic reasons to prefer one elicitation mechanism over another, e.g., that the final outcome corresponds to the candidate(s) that the most voters indicate is their first choice. Another objective may be to minimize the cognitive load imposed on voters. Asking voters to provide a full ranking over the candidates and then using a rate-optimal scoring rule trivially provides faster learning than any other mechanism. However, asking voters to rank 20 candidates is prohibitive in many settings. Future empirical work, in line with that of Benade et al. (2018) and Gelauff et al. (2018), should study the cognitive load various mechanisms impose on voters, to better understand the trade-off between the objectives.

Acknowledgments

We thank our Participatory Budgeting city partners, especially those in Boston, Durham, and Rochester. We also thank anonymous reviewers for their comments. This work was supported in part by the Stanford Cyber Initiative, the Office of Naval Research grant N00014-19-1-2268, and National Science Foundation grants 1544548 and 1637397.

References

- Alós-Ferrer, C., and Granić, D.-G. 2012. Two field experiments on approval voting in germany. *Social Choice and Welfare* 39(1):171–205.
- Aziz, H.; Gaspers, S.; Gudmundsson, J.; Mackenzie, S.; Mattei, N.; and Walsh, T. 2015. Computational aspects of multi-winner approval voting. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, 107–115. International Foundation for Autonomous Agents and Multiagent Systems.
- Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48(2):461–485.
- Benade, G.; Itzhak, N.; Shah, N.; and Procaccia, A. D. 2018. Efficiency and Usability of Participatory Budgeting Methods. 8.
- Caragiannis, I., and Micha, E. 2017. Learning a Ground Truth Ranking Using Noisy Approval Votes. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, 149–155. Melbourne, Australia: International Joint Conferences on Artificial Intelligence Organization.
- Caragiannis, I.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research* 58:123–152.
- Caragiannis, I.; Chatzigeorgiou, X.; Krimpas, G. A.; and Voudouris, A. A. 2019. Optimizing positional scoring rules for rank aggregation. *Artificial Intelligence* 267:58–77.
- Caragiannis, I.; Procaccia, A. D.; and Shah, N. 2013. When do noisy votes reveal the truth? In *Proceedings of the fourteenth ACM conference on Electronic commerce*, 143–160. ACM.
- Chierichetti, F., and Kleinberg, J. 2014. Voting with limited information and many alternatives. *SIAM Journal on Computing* 43(5):1615–1653.
- Copeland, A. H. 1951. A reasonable social welfare function. Technical report, mimeo, 1951. University of Michigan.
- de Borda, J. C. 1781. Mémoire sur les élections au scrutin.
- de Weerd, M. M.; Gerding, E. H.; and Stein, S. 2016. Minimising the rank aggregation error. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, 1375–1376. International Foundation for Autonomous Agents and Multiagent Systems.
- Dembo, A., and Zeitouni, O. 2010. *Large Deviations Techniques and Applications*, volume 38 of *Stochastic Modelling and Applied Probability*. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Diaconis, P. 1988. Group representations in probability and statistics. *Lecture notes-monograph series* 11:i–192.
- Elkind, E.; Faliszewski, P.; Skowron, P.; and Slinko, A. 2017. Properties of Multiwinner Voting Rules. *Social Choice and Welfare* 48(3):599–632.
- Faliszewski, P., and Talmon, N. 2018. A framework

- for approval-based budgeting methods. *arXiv preprint arXiv:1809.04382*.
- Fishburn, P. C., and Gehrlein, W. V. 1976. Borda’s rule, positional voting, and Condorcet’s simple majority principle. *Public Choice* 28(1):79–88.
- Fishburn, P. C. 1978. Axioms for approval voting: Direct proof. *Journal of Economic Theory* 19(1):180–185.
- Freeman, R.; Pennock, D. M.; Peters, D.; and Vaughan, J. W. 2019. Truthful aggregation of budget proposals. *arXiv preprint arXiv:1905.00457*.
- Garg, N., and Johari, R. 2018. Designing informative rating systems: Evidence from an online labor market. *arXiv preprint arXiv:1810.13028*.
- Garg, N., and Johari, R. 2019. Designing optimal binary rating systems. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*.
- Garg, N.; Kamble, V.; Goel, A.; Marn, D.; and Munagala, K. 2019. Iterative local voting for collective decision-making in continuous spaces. *Journal of Artificial Intelligence Research* 64(1):315–355.
- Gelauff, L.; Sakshuwong, S.; Garg, N.; and Goel, A. 2018. Comparing voting methods for budget decisions on the ASSU ballot. Technical report.
- Goel, A.; Krishnaswamy, A. K.; Sakshuwong, S.; and Aitamurto, T. 2016. Knapsack Voting: Voting mechanisms for Participatory Budgeting.
- Guiver, J., and Snelson, E. 2009. Bayesian inference for plackett-luce ranking models. In *proceedings of the 26th annual international conference on machine learning*, 377–384. ACM.
- Kemeny, J. G. 1959. Mathematics without numbers. *Daedalus* 88(4):577–591.
- Lackner, M., and Skowron, P. 2018a. Consistent Approval-Based Multi-Winner Rules. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC ’18, 47–48. New York, NY, USA: ACM.
- Lackner, M., and Skowron, P. 2018b. A quantitative analysis of multi-winner rules. *arXiv preprint arXiv:1801.01527*.
- Lee, D. T.; Goel, A.; Aitamurto, T.; and Landemore, H. 2014. Crowdsourcing for participatory democracies: Efficient elicitation of social choice functions. In *Second AAAI Conference on Human Computation and Crowdsourcing*.
- Lu, T., and Boutilier, C. 2011. Learning Mallows Models with Pairwise Preferences. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, ICML’11, 145–152. USA: Omnipress.
- Lu, T., and Boutilier, C. 2014. Effective sampling and learning for mallows models with pairwise-preference data. *Journal of Machine Learning Research* 15:3963–4009.
- Mallows, C. L. 1957. Non-null ranking models. i. *Biometrika* 44(1/2):114–130.
- marquis de Condorcet, M. J. A. 1785. *Essai sur l’application de l’analyse a la probabilité des décisions: rendues a la pluralité de voix*. De l’Imprimerie royale.
- Mattei, N., and Walsh, T. 2013. Preflib: A library of preference data [HTTP://PREFLIB.ORG](http://preflib.org). In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT 2013)*, Lecture Notes in Artificial Intelligence. Springer.
- Maystre, L., and Grossglauser, M. 2015. Fast and accurate inference of plackett–luce models. In *Advances in neural information processing systems*. 172–180.
- O’Neill, J. 2013. Open STV.
- Popov, S. V.; Popova, A.; and Regenwetter, M. 2014. Consensus in organizations: Hunting for the social choice conundrum in apa elections. *Decision* 1(2):123.
- Procaccia, A. D., and Shah, N. 2015. Is Approval Voting Optimal Given Approval Votes? In *Advances in Neural Information Processing Systems* 28, 1801–1809.
- Public Agenda. 2016. Public Spending By The People: Participatory Budgeting in the United States and Canada in 2014–15. Technical report, The Yankelovich Center for Public Judgment.
- Ratliff, T. C. 2003. Some startling inconsistencies when electing committees. *Social Choice and Welfare* 21(3):433–454.
- Regenwetter, M.; Kim, A.; Kantor, A.; and Ho, M.-H. R. 2007. The unexpected empirical consensus among consensus methods. *Psychological Science* 18(7):629–635.
- Regenwetter, M.; Grofman, B.; Popova, A.; Messner, W.; Davis-Stober, C. P.; and Cavagnaro, D. R. 2008. Behavioural social choice: a status report. *Philosophical Transactions of the Royal Society B: Biological Sciences* 364(1518):833–843.
- Staring, M. 1986. Two paradoxes of committee elections. *Mathematics Magazine* 59(3):158–159.
- Tataru, M., and Merlin, V. 1997. On the relationship of the Condorcet winner and positional voting rules. *Mathematical Social Sciences* 34(1):81–90.
- Wiseman, J. 2000. Approval voting in subset elections. *Economic Theory* 15(2):477–483.
- Young, H. P. 1975. Social choice scoring functions. *SIAM Journal on Applied Mathematics* 28(4):824–838.
- Young, H. P. 1988. Condorcet’s theory of voting. *American Political science review* 82(4):1231–1244.
- Zhao, Z.; Piech, P.; and Xia, L. 2016. Learning Mixtures of Plackett-Luce Models. In *International Conference on Machine Learning*, 2906–2914.