The Smoothed Complexity of Computing Kemeny and Slater Rankings

Lirong Xia[1] **and Weiqiang Zheng**[2]
[1] RPI
[2] CFCS, Computer Science Dept., Peking University
xial@cs.rpi.edu, weiqiang.zheng@pku.edu.cn

Abstract

The computational complexity of *winner determination* under common voting rules is a classical and fundamental topic in the field of computational social choice. Previous work has established the NP-hardness of winner determination under some commonly-studied voting rules, such as the Kemeny rule and the Slater rule. In a recent position paper, Baumeister, Hogrebe, and Rothe (2020) questioned the relevance of the worst-case nature of NP-hardness in social choice and proposed to conduct smoothed complexity analysis (Spielman and Teng 2009) under Bläser and Manthey's (2015) framework.

In this paper, we develop the first smoothed complexity results for winner determination in voting. We prove the smoothed hardness of Kemeny and Slater using the classical smoothed runtime analysis, and prove a parameterized typical-case smoothed easiness result for Kemeny. We also make an attempt of applying Bläser and Manthey's (2015) smoothed complexity framework in social choice contexts by proving that the framework categorizes an always-exponential-time brute force search algorithm as being smoothed poly-time, under a natural noise model based on the well-studied Mallows model in social choice and statistics. Overall, our results show that smoothed complexity analysis in computational social choice is a challenging and fruitful topic.

1 Introduction

The computational complexity of *winner determination* under common voting rules is a classical and fundamental topic in the field of computational social choice (Brandt et al. 2016, Section 1.2.3). A low computational complexity of winner determination is desirable and is indeed the case for many commonly-studied and widely-applied voting rules. On the other hand, winner determination has been proved to be NP-hard for some classical voting rules such as the Kemeny rule, the Slater rule, the Dodgson rule, and the Young rule (Bartholdi, Tovey, and Trick 1989; Conitzer 2006; Rothe, Spakowski, and Vogel 2003).

To address the worst-case nature of NP-hardness, average-case analysis was conducted to provide a more realistic analysis of algorithms. However, average-case analysis is sensitive to the choice of the distribution over input instances, which may itself be unrealistic. To tackle this problem, Spielman and Teng (2004) introduced *smoothed complexity analysis* to generalize and combine the worst-case analysis and the average-case analysis. The idea is that the input $\vec{x}$ of an algorithm Alg is often a noisy perception of the ground truth input $\vec{x}^*$. Consequently, let $\text{Time}_{\text{Alg}}(\vec{x})$ denote the runtime of Alg when the input is $\vec{x}$, the worst-case is analyzed by assuming that an adversary chooses a ground truth $\vec{x}^*$ and then Nature adds a noise $\vec{\epsilon}$ (e.g. a Gaussian noise) to it, so that the algorithm's input becomes $\vec{x} = \vec{x}^* + \vec{\epsilon}$. Then, the expected runtime is evaluated according to the noise introduced by Nature. Formally, the smoothed runtime of Alg is defined as:

$$\max_{\vec{x}^*} \{\mathbb{E}_{\vec{\epsilon}}[\text{Time}_{\text{Alg}}(\vec{x}^* + \vec{\epsilon})]\}$$

This generalizes the worst-case runtime $\max_{\vec{x}^*}[\text{Time}_{\text{Alg}}(\vec{x}^*)]$ and the average-case runtime $\mathbb{E}_{\vec{x}^*\sim\pi}[\text{Time}_{\text{Alg}}(\vec{x}^*)]$ under a distribution π over inputs.

Smoothed complexity analysis has been applied to a wide range of problems in mathematical programming, machine learning, numerical analysis, discrete math, combinatorial optimization, and equilibrium analysis and price of anarchy, see the survey by Spielman and Teng (2009). In a recent position paper, Baumeister, Hogrebe, and Rothe (2020) proposed to conduct smoothed complexity analysis in computational social choice under Bläser and Manthey's (2015) framework and proposed a natural noise model that leverages the celebrated Mallows (1957) model. However, we are not aware of a technical result on the smoothed complexity of winner determination in voting. The following question remains open.

What is the smoothed complexity of winner determination under commonly-studied voting rules?

As illustrated in the following example, the question is highly relevant not only in the theory of computational social choice, but also in AI-aided group decision-making.

Example 1. Suppose the developer of an intelligent system is planning to implement a voting rule for group decision making. Kemeny is being considered, but the system needs to estimate the practical runtime of computing Kemeny to decide how much computational resource is sufficient for a timely decision. The system can learn and predict agents' preferences from their past behavior, but is only able to do it probabilistically—the agents' preferences can be modeled as their ground truth preferences plus some random noise. Is there an algorithm for Kemeny whose expected runtime is low, no matter what the "ground truth" is? □

Our Model. In this paper we answer the question for the Kemeny rule and the Slater rule, for which winner determination means computing an optimal consensus ranking. Successfully addressing the question requires appropriate choices of (1) a noise model for social choice scenarios, and (2) a notion of expected runtime.

(1) Noise model. We adopt the *smoothed social choice framework* (Xia 2020), which covers a wide range of models, including the Mallows-based model proposed by Baumeister, Hogrebe, and Rothe (2020). In the framework, for any number of alternatives and any number of agents, denoted by $m \geq 3$ and $n \geq 1$ respectively, the adversary chooses a distribution π_j for each agent j from a set of distributions Π_m over all rankings.

(2) Expected runtime. Let $\vec{\pi} = (\pi_1, \ldots, \pi_n) \in \Pi_m^n$ denote the vector of distributions chosen by the adversary, one per agent. Then, given an algorithm Alg, the adversary aims at choosing $\vec{\pi}$ to maximize the expected runtime of Alg on the profile P where each ranking is generated independently from $\vec{\pi}$, formally defined as follows.

$$\widetilde{\text{RT}}_{\Pi_m}(\text{Alg}, m, n) = \sup_{\vec{\pi} \in \Pi_m^n} \mathbb{E}_{P \sim \vec{\pi}} \text{Time}_{\text{Alg}}(P) \qquad (1)$$

In (1), the expectation is taken over randomly-generated profile P according to the distributions $\vec{\pi}$ set by the adversary. Because the size of the input of Alg is the size of P, i.e., $\Theta(nm \log m)$, we desire $\widetilde{\text{RT}}_{\Pi_m}(\text{Alg}, m, n)$ to be polynomial in m and n. Following the convention in statistics and the notation in (Xia 2020), for any $m \geq 3$, we use a *single-agent preference model* $\mathcal{M}_m = (\Theta_m, \mathcal{L}(\mathcal{A}_m), \Pi_m)$ to model the adversary's capability, where Θ_m is the parameter space and $\mathcal{L}(\mathcal{A}_m)$ is the set of all rankings over m alternatives.

Note that when m is a constant, many commonly-studied voting rules, including Kemeny and Slater, are easy to compute. Therefore, to meaningfully analyze the smoothed complexity, we will consider an infinite series of models $\vec{\mathcal{M}} = \{\mathcal{M}_m = (\Theta_m, \mathcal{L}(\mathcal{A}_m), \Pi_m) : m \geq 3\}$, following the convention in average-case complexity theory (Bogdanov and Trevisan 2006) and the smoothed complexity theory proposed by Bläser and Manthey (2015).

Our Contributions. We conduct the smoothed analysis according to (1) and prove smoothed hardness results for Kemeny (Theorem 1) and Slater (Theorem 2), respectively. Both theorems state that for a large class of models, if a smoothed poly-time algorithm exists, then RP=NP, which is considered very unlikely to hold.

Then, we consider parameterized typical-case smoothed complexity and prove a mildly positive result in Theorem 3, which implies that for a large class of Mallows-based models, if the average Kendall Tau's distance in the central rankings and the average dispersion parameters are not too large, then the dynamic programming algorithm for Kemeny proposed by Betzler et al. (2009) runs in poly-time with high probability.

Finally, we make an attempt of applying Bläser and Manthey's (2015) framework in our model in Proposition 1. According to our understanding, the always-exponential-time brute force search algorithm for Kemeny and Slater is smoothed poly-time (in the sense of (Bläser and Manthey 2015), see Definition 8) w.r.t. a large class of models including the model proposed by Baumeister, Hogrebe, and Rothe (2020).

Related Work and Discussions. We are not aware of a previous technical result on the smoothed complexity of social choice problems. As discussed above, Baumeister, Hogrebe, and Rothe (2020) proposed to conduct smoothed complexity analysis in computational social choice and proposed a natural Mallows-based model for such analysis. The authors also suggested that smoothed analysis can be done for analyzing ties and paradoxes in social choice.

Xia (2020) independently proposed to conduct smoothed analysis in social choice, provided a general framework for doing so, and proved dichotomous characterizations for Condorcet's paradox and the ANR impossibility theorem to vanish in the smoothed sense. Our model adopts the noise model by Xia (2020) combined with formulation of the worst average-case runtime (1) as done by Spielman and Teng (2009). We emphasize that (Spielman and Teng 2009) used a different noise model from the one used in this paper.

Our paper aims at making the first technical attempt of smoothed complexity analysis in computational social choice, and overall our results show that the topic is highly challenging and fruitful. The seemingly paradoxical smoothed efficiency of brute force search under Bläser and Manthey's (2015) framework (Proposition 1) is indeed not technically surprising and is deliberately allowed, as Bläser and Manthey (2015) commented. See more technical discussions after Proposition 1. Therefore, this result does not mean that Bläser and Manthey's (2015) theory is wrong or inconsistent, but instead, it can be viewed as a call for future research in the theory of smoothed complexity analysis and statistical models in social choice contexts. The smoothed hardness of Kemeny (Theorem 1) and Slater (Theorem 2) are negative news and the parameterized typical-case smoothed efficiency (Theorem 3) is positive news. Proof techniques developed for these theorems may be useful in future work.

There is a large body of literature on the computational complexity of Kemeny. The corresponding optimization problem, KEMENY RANKING, was proved to be NP-hard (Bartholdi, Tovey, and Trick 1989) and $\text{P}_{\|}^{\text{NP}}$-complete (Hemaspaandra, Spakowski, and Vogel 2005). Approximation algorithms (Ailon, Charikar, and Newman 2008; van Zuylen and Williamson 2007), PTAS (Kenyon-Mathieu and Schudy 2007), and fixed-parameter efficient algorithms (Betzler et al. 2009; Karpinski and Schudy 2010; Cornaz, Galand, and Spanjaard 2013) for KEMENY RANKING have been developed. Practical algorithms have been proposed (Davenport and Kalagnanam 2004; Conitzer, Davenport, and Kalagnanam 2006) and Ali and Meila (2012) compared the performance of 104 algorithms. Conitzer (2006) proved that the decision variant of Slater is NP-hard and proposed an efficient heuristic algorithm for computing

SLATER RANKING.

There is a large body of literature on smoothed complexity of algorithms (Spielman and Teng 2009). Bläser and Manthey (2015) established a complexity theory for smoothed complexity analysis by defining the counterparts to P and NP, called SMOOTHED-P and DIST-NP$_{\text{para}}$ respectively, together with a smoothed reduction and complete problems. Their definitions are closely related to the average-case complexity theory established by Levin (1986). Our result in Section 5 suggests it may not be suitable for computational social choice. Our Theorem 1 and 2 illustrate the hardness of KEMENY RANKING and SLATER RANKING using the same pattern in (Huang and Teng 2007), which states that the existence of a smoothed poly-time algorithm would lead to a surprise in complexity theory.

2 Preliminaries

Basics of Voting. For any $m \geq 3$, let $\mathcal{A}_m = \{a_1, \ldots, a_m\}$ denote the set of m alternatives. A *(preference) profile* $P \in \mathcal{L}(\mathcal{A}_m)^n$ is a collection of n rankings (linear orders). Let $\text{WMG}(P)$ denote the *weighted majority graph* of P, which is a directed weighted graph whose vertices are $\mathcal{A}_m$ and for each pair of alternatives a, b, the weight on edge $a \to b$, denoted by $w_P(a,b)$, is the winning margin of a over b in their pairwise competition. That is, $w_P(a,b) = -w_P(b,a) = \#\{R \in P : a \succ_R b\} - \#\{R \in P : b \succ_R a\}$. Let $\text{UMG}(P)$ denote the *unweighted majority graph* of P, which is the unweighted directed graph obtained from $\text{WMG}(P)$ by removing edges whose weights are ≤ 0.

The *Kendall's Tau* distance between two linear orders $R, W \in \mathcal{L}(\mathcal{A})$, denoted by $\text{KT}(R,W)$, is the number of pairwise disagreements between R and W. Given a profile P and a linear order R, the *Kemeny score* of R in P is $\sum_{W \in P} \text{KT}(R,W)$ and the *Slater score* of R in P is $\text{KT}(R, \text{UMG}(P))$, where KT is extended to measure the distance between a linear order and an unweighted graph in the natural way—any pair of alternatives $\{a,b\}$ such that $a \succ b$ in the linear order but $b \to a$ in the graph contribute one to KT. The Kemeny rule (respectively, the Slater rule) aims at selecting the linear order with the minimum Kemeny score (respectively, Slater score) in P. The corresponding winner determination problems are defined as follows.

Definition 1 (KEMENY RANKING **and** SLATER RANKING**).** *Given $m \geq 3$, $n \in \mathbb{N}$, and $P \in \mathcal{L}(\mathcal{A}_m)^n$, in* KEMENY RANKING *(respectively,* SLATER RANKING*), we are asked to compute a ranking with minimum Kemeny score (respectively, Slater score).*

Definition 2 (Single-agent preference model (Xia 2020)). *A* single-agent preference model *for m alternatives is denoted by $\mathcal{M}_m = (\Theta_m, \mathcal{L}(\mathcal{A}_m), \Pi_m)$, where Π_m is the set of distributions over $\mathcal{L}(\mathcal{A}_m)$ indexed by the parameter space Θ_m. $\mathcal{M}_m$ is* neutral *if for any $\theta \in \Theta_m$ and any permutation σ over $\mathcal{A}_m$, there exists $\theta' \in \Theta_m$ such that for all $V \in \mathcal{L}(\mathcal{A}_m)$, we have $\pi_\theta(V) = \pi_{\theta'}(\sigma(V))$. $\mathcal{M}_m$ is* P-samplable *if there exists a poly-time sampling algorithm for each distribution in Π_m.*

Technically $\mathcal{M}_m$ is completely determined by Π_m. Following the convention in statistics, we still keep the parameter space Θ_m and sample space $\mathcal{L}(\mathcal{A}_m)$ in the definition. For example, let us recall the definition of single-agent Mallows model as follows.

Definition 3. *In a* single-agent Mallows model $\mathcal{M}_{Ma,m}$, *$\Theta_m = \mathcal{L}(\mathcal{A}_m) \times (0,1]$, where in each $(R,\varphi) \in \Theta_m$, R is called the* central ranking *and φ is called the* dispersion parameter. *For any $W \in \mathcal{L}(\mathcal{A}_m)$, we have $\pi_{(R,\varphi)} = \varphi^{KT(R,W)}/Z_\varphi$, where $Z_\varphi = \frac{\prod_{i=2}^{m}(1-\varphi^i)}{(1-\varphi)^{m-1}}$ is the normalization constant. For any $0 < \underline{\varphi} \leq \overline{\varphi} \leq 1$, we let $\mathcal{M}_{Ma,m}^{[\underline{\varphi},\overline{\varphi}]}$ denote the sub-model whose parameter space is $\mathcal{L}(\mathcal{A}_m) \times [\underline{\varphi},\overline{\varphi}]$.*

It is not hard to verify that $\mathcal{M}_{\text{Ma},m}^{[\underline{\varphi},\overline{\varphi}]}$ is neutral and P-samplable (Doignon, Pekeč, and Regenwetter 2004).[1]

When there are $n \geq 2$ agents, the adversary chooses $\vec{\pi} = (\pi_1, \ldots, \pi_n) \in \Pi_m^n$, and then agent j's ranking will be independently (but not necessarily identically) generated from π_j.

Example 2. *Suppose $m = 3$ and $n = 2$, and the model is $\mathcal{M}_{Ma,3}^{[0.3,1]}$. Then, the adversary can set the first (respectively, second) agent's distribution to be the Mallows distribution given ground truth $(a_1 \succ a_2 \succ a_3, 0.4)$ (respectively, $(a_3 \succ a_2 \succ a_1, 0.8)$). Then, the probability of generating profile $(a_2 \succ a_1 \succ a_3, a_1 \succ a_3 \succ a_2)$ is $\Pr(a_2 \succ a_1 \succ a_3|(a_1 \succ a_2 \succ a_3, 0.4)) \times \Pr(a_1 \succ a_3 \succ a_2|(a_3 \succ a_2 \succ a_1, 0.8)) = \frac{0.4}{Z_{0.4}} \times \frac{0.8^2}{Z_{0.8}} \approx 0.026$.*

As another example, the Mallows-based model proposed by Baumeister, Hogrebe, and Rothe (2020) corresponds to the single-agent Mallows model with fixed φ, i.e., $\mathcal{M}_{Ma,m}^{[\varphi,\varphi]}$.

Because KEMENY SCORE is in P when m is bounded above by a constant, the smoothed complexity analysis ought to be done for variable m. Therefore, following (Bläser and Manthey 2015), we are given a series of single-agent preference models $\vec{\mathcal{M}} = \{\mathcal{M}_m = (\Theta_m, \mathcal{L}(\mathcal{A}_m), \Pi_m) : m \geq 3\}$. In particular, we will focus on the *Mallows series*, defined as follows.

Definition 4. *For any $0 < \underline{\varphi} \leq \overline{\varphi} \leq 1$, we let $\vec{\mathcal{M}}_{Ma}^{[\underline{\varphi},\overline{\varphi}]} = \{\mathcal{M}_{Ma,m}^{[\underline{\varphi},\overline{\varphi}]} : m \geq 3\}$ denote the* Mallows series.

In most part of this paper (except Section 5), we focus on the classical smoothed poly-time algorithms proposed by Spielman and Teng (2009) w.r.t. $\vec{\mathcal{M}}$.

Definition 5 (Smoothed poly-time). *Given a series of single-agent preference models $\vec{\mathcal{M}} = \{\mathcal{M}_m = (\Theta_m, \mathcal{L}(\mathcal{A}_m), \Pi_m) : m \geq 3\}$, an algorithm Alg is smoothed poly-time, if for any $m \geq 3$ and $n \geq 1$, $\widetilde{RT}_{\Pi_m}(Alg, m, n)$ as defined in (1) is polynomial in m and n.*

[1] We note that in this paper $\underline{\varphi}$ and $\overline{\varphi}$ are constants that do not depend on m. We believe that this is a natural model in social choice, and the case of variable $\underline{\varphi}$ and $\overline{\varphi}$ is an interesting direction for future work.

3 Smoothed Hardness of Kemeny and Slater

In this section, we follow the classical smoothed runtime analysis (Definition 5) to analyze KEMENY RANKING and SLATER RANKING. We first recall the *orthogonal decomposition* of weighted directed graphs (Young 1974; Zwicker 2018).

A WMG G_{cyc} is called a *cycle*, if the absolute weight of any edge is 0 or 1, and the edges with positive weights forms a cycle. Let $a \in \mathcal{A}_m$, a WMG G_a is called a *co-cycle* centered at a, if for any $b \in \mathcal{A}_m$ such that $b \neq a$, $w_{G_a}(a,b) = -w_{G_a}(b,a) = 1$ and all other edges have weight 0.

Note that any WMG G can be viewed as a vector in $\mathbb{R}^{\frac{m(m-1)}{2}}$ whose components are indexed by (i_1, i_2), where $1 \le i_1 < i_2 \le m$, with value $w_G(a_{i_1}, a_{i_2})$. Given any pair of WMGs G_1 and G_2, we define their dot product as

$$G_1 \cdot G_2 = \sum_{1 \le i_1 < i_2 \le m} w_{G_1}(a_{i_1}, a_{i_2}) \times w_{G_2}(a_{i_1}, a_{i_2})$$

Let $\mathcal{V}_{cyc} \subseteq \mathbb{R}^{\frac{m(m+1)}{2}}$ (respectively, $\mathcal{V}_{co} \subseteq \mathbb{R}^{\frac{m(m+1)}{2}}$) denote the linear span of cycles (respectively, co-cycles). It has been proved that $\mathcal{V}_{cyc}$ and $\mathcal{V}_{co}$ are orthogonal, $\dim(\mathcal{V}_{cyc}) = \binom{m-1}{2}$ (all 3-cycles containing a_1 in the increasing direction of subscripts constitute a non-orthogonal basis) and $\dim(\mathcal{V}_{co}) = m - 1$ (co-cycles centered at any fixed set of $m-1$ alternatives constitute a non-orthogonal basis). An *orthogonal decomposition* of a WMG G is a decomposition of G into its projections to $\mathcal{V}_{cyc}$ and $\mathcal{V}_{co}$, respectively.

We will sometimes use *fractional* profiles, i.e., the weights on the linear orders are allowed to be arbitrary non-negative numbers and can be larger than one. For example, $P = \frac{3}{2}@[a \succ b] + \frac{1}{2}@[b \succ a]$ is a fractional profile for $m = 2$, where the weight on $a \succ b$ is $\frac{3}{2}$ and the weight on $b \succ a$ is $\frac{1}{2}$. WMG, UMG, and KT distance can be naturally extended to fractional profiles, where the total weights are used to replace total number counts.

A special case is the fractional profile that corresponds to a distribution π, where the weight on each linear order R is its probability in the distribution, i.e., $\pi(R)$. The following example illustrates a fractional profile and its orthogonal decomposition.

Example 3. *Let $\theta = (a_1 \succ a_2 \succ a_3, \varphi)$ denote a parameter in $\mathcal{M}^{[\underline{\varphi},\overline{\varphi}]}_{Ma,3}$. WMG($\theta$), which is the WMG of the fractional profile represented by the distribution corresponding to θ, and its orthogonal decomposition are shown in Figure 1. Note that the weight on the co-cycle centered at a_3 is negative.*

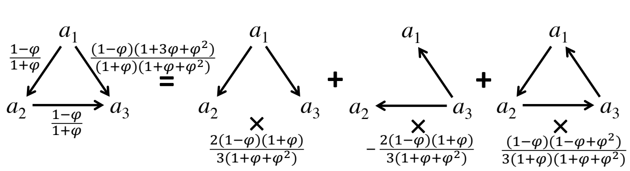

Figure 1: The WMG of $(a_1 \succ a_2 \succ a_3, \varphi)$ in $\mathcal{M}_{\text{Ma},3}$ and its orthogonal decomposition.

To present the theorem, we formally define some assumptions on the single-agent preference model $\vec{\mathcal{M}}$ as follows. The first assumption states that $\vec{\mathcal{M}}$ is easy to sample, the second assumption states that the model is neutral (see Definition 2), and the third assumption is introduced for technical reasons, which requires that $\vec{\mathcal{M}}$ is rich enough in the sense that there exists a distribution $\pi \in \Pi_m$ whose WMG has a non-negligible 3-cycle component in its orthogonal decomposition.

Assumption 1. $\vec{\mathcal{M}} = \{\mathcal{M}_m = (\Theta_m, \mathcal{L}(\mathcal{A}_m), \Pi_m) : m \ge 3\}$ *satisfies the following assumptions:*

(i) For any $m \ge 3$, $\mathcal{M}_m$ is P-samplable*, i.e., it admits a poly-time sampling algorithm.*

(ii) For any $m \ge 3$, $\mathcal{M}_m$ is neutral, i.e., for each distribution $\pi \in \Pi_m$ and any distribution σ over the alternatives, we have $\sigma(\pi) \in \Pi_m$.

(iii) There exist constants $k \ge 0$ and $A > 0$ such that for any $m \ge 3$, there exist $\pi_{3c} \in \Pi_m$ such that WMG(π_{3c}) has a 3-cycle component G_{3c} with WMG(π_{3c}) $\cdot\, G_{3c} > \frac{A}{m^k}$.

The first condition is the "most natural restriction" on general distributions (Bogdanov and Trevisan 2006, p. 17), which is less restrictive than the commonly-studied *P-computable* distributions (Bogdanov and Trevisan 2006, p. 18). The second and third conditions imply that $\vec{\mathcal{M}}$ is rich enough. As we will see in Example 4, a large class of models, including the Mallows series (Definition 4), satisfy Assumption 1.

Theorem 1 (Smoothed Hardness of KEMENY RANKING**).** *For any series of single-agent preference models $\vec{\mathcal{M}}$ that satisfies Assumption 1, if there exists a smoothed poly-time algorithm for* KEMENY RANKING *w.r.t. $\vec{\mathcal{M}}$, then* NP=RP.

Proof. The theorem is proved by contradiction. Suppose for the sake of contradiction that a smoothed poly-time algorithm Alg for KEMENY RANKING exists, we will prove that there exists an efficient randomized algorithm for the NP-hard problem EULERIAN FEEDBACK ARC SET (EFAS) (Perrot and Pham 2015). An instance of EFAS is denoted by (G, t), where $t \in \mathbb{N}$ and G is a directed unweighted Eulerian graph, which means that there exists a closed Eulerian walk that passes each edge exactly once. We are asked to decide whether G can be made acyclic by removing no more than t edges.

Given a single-agent preference model, a (fractional) *parameter profile* $P^\Theta \in \Theta_m^n$ is a collection of $n > 0$ parameters, where n may not be an integer. Note that P^Θ naturally leads to a fractional preference profile, where the weight on each ranking represents its total weighted "probability" under all parameters in P^Θ. Therefore, WMG and UMG can be naturally extended to parameter profiles.

The high-level idea of the proof is the following. For any EFAS instance (G, t), we will construct a fractional parameter profile P^Θ_G whose WMG is the same as the WMG equivalent of G, where the weight on each edge in G is 1 or -1, depending on its direction. Then, we sample a profile P' from P^Θ_G and try to run Alg to compute the Kemeny ranking R^*. If Alg successfully returns R^* in less than three times of its expected runtime (which is polynomial), then we proceed to check whether R^* leads to a YES answer to (G, t). Otherwise we give a No answer. More precisely, we give a NO answer if the ordering of vertices according to R^* has more than t backward edges in G, or Alg fails to terminate in time. Clearly this polynomial-time procedure always returns NO if (G, t) is a NO instance. We then prove that a YES instance will receive a YES answer with probability at least $1/2$, which means that EFAS is in RP and therefore proves the theorem.

Formally, the proof proceeds in three steps. In Step 1, we use permutations of π_{3c} guaranteed by Assumption 1 to construct a fractional parameter profile $P^{\Theta}_{G_{3c}}$ whose WMG is a 3-cycle, denoted by G_{3c}. In Step 2, we use $P^{\Theta}_{G_{3c}}$ to construct P^{Θ}_{G}. In Step 3, we show that Alg can be leveraged to Algorithm 1 to prove that EFAS is in RP as discussed above.

Step 1. Construct a parameter profile $P^{\Theta}_{G_{3c}}$ whose WMG is a 3-cycle. W.l.o.g. let $G_{3c} = a_1 \to a_2 \to a_3 \to a_1$ denote the target 3-cycle. Let σ_1 denote an arbitrary cyclic permutation among $\{a_1, a_2, a_3\}$ and let σ_2 denote an arbitrary cyclic permutation among $\mathcal{A} \setminus \{a_1, a_2, a_3\}$. We define the following set of $6(m-3)$ permutations.

$$\mathrm{O}_{G_{3c}} = \{\sigma_1^i \circ \sigma_2^t, \sigma_1^i \circ \sigma_2^{-t} : 1 \le i \le 3, 1 \le t \le m-3\},$$

where σ_1^i represents the application of σ_1 for i times. We note that $\mathrm{O}_{G_{3c}}$ can be naturally applied to linear orders, fractional profiles, distributions over $\mathcal{L}(\mathcal{A}_m)$, parameters in Θ_m, and weighted majority graphs. For example, for each linear order R, $\mathrm{O}_{G_{3c}}(R)$ is a set of $6(m-3)$ linear orders that are obtained from applying the $6(m-3)$ permutations in $\mathrm{O}_{G_{3c}}$ to R. For each (fractional) profile P, $\mathrm{O}_{G_{3c}}(P)$ is a set of profiles obtained from the union of the image of $\sigma^*(P)$ for all $\sigma^* \in \mathrm{O}_{G_{3c}}$. It follows that the total weight of linear orders in $\mathrm{O}_{G_{3c}}(P)$ is $6(m-3)n$, where n is the total weight of linear orders in P. When a permutation σ over $\mathcal{A}$ is applied to a parameter θ, the outcome is another parameter θ' such that $\pi_{\theta'} = \sigma(\pi_\theta)$, where we recall that $\pi_{\theta'}$ and π_θ are the distributions represented by θ' and θ, respectively. The existence of such θ' is guaranteed by the neutrality of the model (see Definition 2).

Let $\theta_{3c} \in \Theta_m$ denote the parameter corresponding to π_{3c} (the distribution guaranteed by Assumption 1 (iii)), i.e., $\pi_{\theta_{3c}} = \pi_{3c}$. Let $Q^{\Theta}_{G_{3c}} = \mathrm{O}_{G_{3c}}(\theta_{3c})$ denote the parameter profile obtained from θ_{3c} by applying permutations in $O_{G_{3c}}$. Some properties of $Q^{\Theta}_{G_{3c}}$ are described in the following claim, whose proof follows after the definition of $Q^{\Theta}_{G_{3c}}$. All missing proofs can be found in the full version of this paper on arXiv.

Claim 1. *$|Q^{\Theta}_{G_{3c}}| = O(m)$, $Q^{\Theta}_{G_{3c}}$ consists of $O(m)$ types of parameters. WMG($Q^{\Theta}_{G_{3c}}$) consists of the following two types of edges. (1) There are three edges $a_1 \to a_2$, $a_2 \to a_3$, $a_3 \to a_1$, each has weight $2(m-3)\alpha$, where $\alpha =$ WMG(π_{3c}) $\cdot\, G_{3c}$. (2) There are edges from $\{a_1, a_2, a_3\}$ to $\{a_4, \ldots, a_m\}$ whose weights are $\beta = 2\sum_{(d_1,d_2)\in\{a_1,a_2,a_3\}\times\{a_4,\ldots,a_m\}} w_{\pi_{3c}}(d_1, d_2)$.*

$|Q^{\Theta}_{G_{3c}}|$ and the types of parameters used in $Q^{\Theta}_{G_{3c}}$ will be crucial later in proving that Algorithm 1 runs in polynomial time. If $\beta = 0$, then we let $P^{\Theta}_{G_{3c}} = Q^{\Theta}_{G_{3c}}$. If $\beta > 0$, then we will provide a gadget soon to "cancel" edges between $\{a_1, a_2, a_3\}$ and $\mathcal{A}_m \setminus \{a_1, a_2, a_3\}$. For any alternative a, let σ_a denote the permutation such that $\sigma_a(a) = a$ and the remaining alternatives are permuted in the cyclic way in the increasing order of their subscripts. That is, $\sigma_a = a_{i_1} \to a_{i_2} \to \cdots \to a_{i_{m-1}} \to a_{i_1}$, where $i_1 < \cdots < i_{m-1}$ and $a_i \to a_j$ means that $\sigma_a(a_i) = a_j$. We let η denote the permutation that switches a_{i_s} and $a_{i_{m-s}}$ for all $1 \le s \le m-1$. For example, when $m = 5$, σ_{a_1} is the cyclic permutation $a_2 \to a_3 \to a_4 \to a_5 \to a_2$ and η_{a_1} switches two pairs of alternatives: (a_2, a_5) and (a_3, a_4).

Definition 6. *For any alternative a, let O_a denote the following set of $2(m-1)$ permutations over $\mathcal{A}_m$.*

$$O_a = \{\sigma_a^i, \sigma_a^i \circ \eta_a : 1 \le i \le m-1\}$$

Like $\mathrm{O}_{G_{3c}}$, O_a can be naturally applied to linear orders, fractional profiles, distributions over $\mathcal{L}(\mathcal{A}_m)$, parameters in Θ_m, and weighted majority graphs. Let $P^{\Theta}_a = \mathrm{O}_a(\theta_{3c})$. We have the following claim.

Claim 2. *$|P^{\Theta}_a| = 2(m-1)$ and P^{Θ}_a consists of $2(m-1)$ types of parameters. WMG(P^{Θ}_a) is a co-cycle whose center is a and the absolute weight on any non-zero edge is $2\sum_{(d_1,d_2)\in\{a\}\times\mathcal{A}\setminus\{a\}} w_{\pi_{3c}}(d_1, d_2)$.*

Let $P^{\Theta}_1 = \mathrm{O}_{a_1}(Q^{\Theta}_{G_{3c}})$. We have the following claim..

Claim 3. *$|P^{\Theta}_1| = O(m^2)$ and P^{Θ}_1 consists of $O(m^2)$ types of parameters. WMG(P^{Θ}_1) is a co-cycle whose center is a_1 and the absolute weight on any non-zero edge is $2(m-3)\beta$.*

For any pair of alternative a, d, let $\sigma_{a\leftrightarrow d}$ denote the permutation that exchanges a and d. Because the WMG of P^{Θ}_1 is a co-cycle centered at a_1, $\sigma_{a_1\leftrightarrow d}(P^{\Theta}_1)$ is a parameter profile whose WMG is a co-cycle centered at d. We are now ready to define the fractional parameter profile $P^{\Theta}_{G_{3c}}$ whose WMG resembles G_{3c}. Recall that we defined $\alpha =$ WMG(π_{3c}) $\cdot\, G_{3c}$ in Claim 1.

Definition 7. *Let $P^{\Theta}_{G_{3c}}$ denote the fractional parameter profile that consists of (1) $\frac{1}{2(m-3)\alpha}$ copies of $Q^{\Theta}_{G_{3c}}$, and (2) for each $d \in \{a_4, \ldots, a_m\}$, $\frac{1}{4(m-3)^2\alpha}$ copies of $\sigma_{a_1\leftrightarrow d}(P^{\Theta}_1)$.*

We have the following claim about $P^{\Theta}_{G_{3c}}$.

Claim 4. *$|P^{\Theta}_{G_{3c}}| = O(m^k)$ and $P^{\Theta}_{G_{3c}}$ consists of $O(m^3)$ different types of parameters. WMG($P^{\Theta}_{G_{3c}}$) $= G_{3c}$.*

Step 2. Construct a parameter profile P^{Θ}_G for EFAS. Because $\mathcal{M}_m$ is neutral, for any 3-cycle $G' = a_{i_1} \to a_{i_2} \to a_{i_3} \to a_{i_1}$ we can apply a permutation $\sigma_{G'}$ that maps a_s to a_{i_s} ($s = 1, 2, 3$) on $P^{\Theta}_{G_{3c}}$, which means that WMG($\sigma_{G'}(P^{\Theta}_{G_{3c}})$) resembles G'.

It is not hard to verify that any cycle of length T can be obtained from the union of $(T-2)$ individual 3-cycles, which can be computed in $O(m^2)$ time. Therefore, given the EFAS instance (G, t), we first compute $G = \bigcup_{s=1}^{S} G'_s$, where each G'_s is a 3-cycle, and $S \le \binom{m}{2}$. Then, we

ALGORITHM 1: Algorithm for EFAS.

Input: An EFAS instance (G, t), Alg for KEMENY RANKING whose smoothed runtime is T.
Compute a parameter profile $P_G^{\Theta *}$ according to (2).
Sample a profile P' from $\vec{\mathcal{M}}_m$ given $P_G^{\Theta *}$ and run Alg on P'.
if *Alg returns* R^* *within* $3T$ *time and* R^* *is a solution to* (G, t) **then**
  return YES
else
  return NO
end

let $P_G^{\Theta} = \bigcup_{s=1}^{S} \sigma_{G'_s}(P_{G_{3c}}^{\Theta})$. It follows from Claim 4 that $|P_G^{\Theta}| = O(m^{k+2})$, P_G^{Θ} consists of $O(m^5)$ types of parameters, and WMG$(P_G^{\Theta}) = G$.

Step 3. Use Alg to solve EFAS. Let $K = 11 + 2k$, which means that $K > 9$. We first define a parameter profile $P_G^{\Theta *}$ of $n = \Theta(m^K)$ parameters that is approximately $\frac{m^K}{|P_G^{\Theta}|}$ copies of P_G^{Θ} up to $O(m^5)$ in L_∞ error. Formally, let

$$P_G^{\Theta *} = \lfloor P_G^{\Theta} \cdot \frac{m^K}{|P_G^{\Theta}|} \rfloor \tag{2}$$

Let $n = |P_G^{\Theta *}|$. Because the number of different types of parameters in $P_G^{\Theta *}$ is $O(m^5)$, we have $n = m^K - O(m^5)$, $\|\text{WMG}(P_G^{\Theta *}) - \text{WMG}(P_G^{\Theta} \cdot \frac{m^K}{|P_G^{\Theta}|})\|_\infty = O(m^5)$, and $\|\text{WMG}(P_G^{\Theta *}) - G \cdot \frac{m^K}{|P_G^{\Theta}|})\|_\infty = O(m^5)$.

We now prove that Alg can be leveraged to provide an RP algorithm (Algorithm 1) for EFAS.

Notice that sampling P' from $P_G^{\Theta *}$ takes polynomial time because $\vec{\mathcal{M}}$ is P-samplable (Assumption 1 (i)). It follows that Algorithm 1 is a polynomial-time randomized algorithm. Clearly, if (G, t) is a NO instance, then Algorithm 1 returns NO. Therefore, to prove that Algorithm 1 is an RP algorithm it suffices to prove that if (G, t) is a YES instance, then Algorithm 1 returns YES with $> \frac{1}{2}$ probability.

Let $G_n = G \cdot \frac{m^K}{|P_G^{\Theta}|}$. We first prove in the following claim that with exponentially small probability WMG(P') is different from G_n by more than $\Omega(m^{\frac{K+1}{2}})$.

Claim 5. $\Pr(\|WMG(P') - G_n\|_\infty > \Omega(m^{\frac{K+1}{2}})) < \exp^{-\Omega(m)}$.

Proof. We first show that for each pairwise comparison b vs. c, $\Pr(|w_{P'}(b,c) - w_G(b,c) \cdot \frac{m^K}{|P_G^{\Theta}|}| = \Omega(m^{\frac{K+1}{2}})) < \exp^{-\Omega(m)}$, then apply union bound to all pairwise comparisons. Notice that $w_{P'}(b,c)$ can be viewed as the sum of n independent (not necessarily identical) bounded random variables, each of which corresponds to the pairwise comparison between b and c in a ranking—if $b \succ c$ in the ranking, then the random variable takes 1, otherwise the random variable takes -1. By Hoeffding's inequality for bounded random variables, we have:

$$\Pr(|w_{P'}(b,c) - \mathbb{E}(w_{P'}(b,c))| > \Omega(m^{\frac{K+1}{2}}))$$
$$< \exp\{-\frac{\Omega(m^{\frac{K+1}{2}}/n)^2 n^2}{4n}\} = \exp\{-\Omega(m)\}$$

Also notice that $\mathbb{E}(w_{P'}(b,c)) = w_{P_G^{\Theta *}}(b,c)$ and $\|\text{WMG}(P_G^{\Theta *}) - G_n\| = O(m^5) = O(m^{\frac{K+1}{2}})$. □

Suppose (G, t) is a YES instance. That is, there exists a ranking R' whose KT distance to G is no more than t. Due to Markov's inequality, Alg returns a Kemeny ranking R^* with probability $\geq \frac{2}{3}$. We now prove that R^* is a solution to (G, t) with probability $1 - \exp^{-\Omega(m)}$.

Note that for any ranking R, $|\text{KT}(R, P') - \text{KT}(R, G_n)| = O(m^{\frac{K+5}{2}})$ holds with probability $1 - \exp^{-\Omega(m)}$, which follows after applying Claim 5 and the union bound to all $\Theta(m^2)$ pairwise comparisons. In this case R^* is a solution to (G, t), because if it is not, then $\text{KT}(R^*, G_n) - \text{KT}(R', G_n) > \frac{m^K}{|P_G^{\Theta}|} = \Omega(m^{K-k-2})$. We note that $|\text{KT}(R', P') - \text{KT}(R', G_n)| = O(m^{\frac{K+1}{2}}) \times O(m^2) = O(m^{\frac{K+5}{2}})$ and $|\text{KT}(R^*, P') - \text{KT}(R^*, G_n)| = O(m^{\frac{K+5}{2}})$, which means that $\text{KT}(R^*, P') - \text{KT}(R', P') = \Theta(m^{K-k-2}) - 2O(m^{\frac{K+5}{2}}) > 0$ (because $K = 11 + 2k$), which contradicts the optimality of R^*.

Therefore, Algorithm 1 returns YES with probability at least $\frac{2}{3} - \exp^{-\Omega(m)}$, which proves that EFAS is in RP. Since EFAS is NP-hard (Perrot and Pham 2015) and RP⊆NP. It follows that RP=NP. □

We now prove a similar theorem for SLATER RANKING with an additional condition. The theorem requires one more richness assumption in addition to Assumption 1, which can be viewed as the co-cycle counterpart to (iii)—the new assumption requires that Π_m contains a distribution whose WMG has a non-negligible co-cycle component. While this assumption is quite technical, we expect it to hold for many models in social choice, as Example 4 shows.

(iv) There exist constants $k^* \geq 0$ and $B > 0$ such that for any $m \geq 3$, there exist $\pi_{co} \in \Pi_m$ such that WMG(π_{co}) has a co-cycle component G_{co} with WMG$(\pi_{co}) \cdot G_{co} > \frac{B}{m^{k^*}}$.

Theorem 2 (**Smoothed Hardness of** SLATER RANKING). *For any series of single-agent preference models $\vec{\mathcal{M}}$ that satisfies Assumption 1 and (iv), if there exists a smoothed poly-time algorithm for* SLATER RANKING *w.r.t. $\vec{\mathcal{M}}$, then* NP=RP.

Proof sketch. The high-level idea of the proof is similar to the proof of Theorem 1. The difference is that in this proof we use the TOURNAMENT FEEDBACK ARC SET (TFAS) problem, which is NP-hard (Alon 2006; Conitzer 2006). Since it is unknown whether EULERIAN TOURNAMENT FEEDBACK ARC SET is NP-hard, in a TFAS instance (G, t) it is possible that G is not Eulerian. Therefore, we need a co-cycle component to construct a parameter profile whose

WMG is G. Moreover, the weight on the co-cycle component cannot be too small, otherwise the construction will not be polynomial. Condition (iv) is used to guarantee the existence of a desirable co-cycle component as described.

Slightly more formally, the proof proceeds in four steps. Step 1 is the same as the Step 1 in the proof of Theorem 1. In Step 2, we use permutations of π_{co}, which is guaranteed by (iv), to construct a parameter profile whose WMG is a co-cycle. The reduction from TFAS will be presented in Step 3. In Step 4 we show that a smoothed poly-time algorithm for SLATER RANKING can be used to prove that TFAS is in RP. □

The following example shows that Assumption 1 and (iv) hold for a large class of Mallows-based models.

Example 4. *We show that for any $\underline{\varphi} \neq 1$, $\vec{\mathcal{M}}_{Ma}^{[\underline{\varphi},\overline{\varphi}]}$ satisfies Assumption 1 and (iv). (i) and (ii) have been discussed in Example 3. For any $\varphi \in [\underline{\varphi}, \overline{\varphi}] \cap (0,1)$, we show that (iii) and (iv) hold for $\pi = (a_1 \succ \cdots \succ a_m, \varphi)$.*

For (iii), let $G_{3c} = a_1 \to a_2 \to a_3 \to a_1$. Mallows (1957) proved that under Mallows' model, the probability for $a \succ b$ only depends on m, φ, and the difference in the ranks of a and b in the central ranking. Therefore, for any m, WMG$(\pi) \cdot G_{3c} = 2 \cdot \frac{1}{1+\varphi} - \frac{1+2\varphi}{(1+\varphi)(1+\varphi+\varphi^2)} = \frac{1+2\varphi^2}{(1+\varphi)(1+\varphi+\varphi^2)} = \Theta(1)$, see Figure 1. This means that $k = 0$ and $A = \Theta(1)$.

For (iv), it is not hard to verify that $k^* = 0$ and $B = \Theta(1)$ for the co-cycle centered at a_1.

4 Parameterized Typical-Case Smoothed Complexity of KEMENY RANKING

Following the idea in *probably polynomial smoothed complexity* of perceptron (Blum and Dunagan 2002), in this section we prove a similar result on the parameterized typical-case smoothed complexity of Betzler et al. (2009)'s dynamic programming algorithm for KEMENY RANKING, denoted by Alg_{KS}, under the Mallows series $\vec{\mathcal{M}}_{\text{Ma}}$ (Definition 4). Under $\vec{\mathcal{M}}_{\text{Ma}}$, for convenience, sometimes we rewrite a parameter $\vec{\theta} \in \Theta_m^n$ as $(P, \vec{\varphi})$, where $P \in \mathcal{L}(\mathcal{A}_m)^n$ is called the *central profile* and $\vec{\varphi} \in (0,1]^n$ is called the *dispersion vector*. For any n-profile P, *the average KT distance* is defined as

$$\overline{\text{KT}}(P) = \frac{1}{n(n-1)} \sum_{R_1, R_2 \in P} \text{KT}(R_1, R_2)$$

The high-level idea behind Theorem 3 below is as follows. Betzler et al. (2009) proposed Alg_{KS} and proved that for any profile P, its runtime is $O(exp(\overline{\text{KT}}(P))poly(mn))$. Therefore, to show that with high probability Alg_{KS} runs in polynomial time, it suffices to show that with high probability, $\overline{\text{KT}}(P')$ is upper bounded by a constant, where P' is a randomly generated profile according to $\vec{\mathcal{M}}_{\text{Ma}}$ given some central profile P and dispersion vector $\vec{\varphi}$. Intuitively, when P is sufficiently centralized, i.e., $\overline{\text{KT}}(P)$ is upper bounded by a constant, and $\vec{\varphi}$ is not too large on average, P' should also be sufficiently centralized. The formal relationship between centralization of P, $\vec{\varphi}$, probability for Alg_{KS} to be in P, and the runtime is stated in the following theorem.

Theorem 3 (Parameterized Typical-Case Smoothed Complexity of KEMENY RANKING). *For any $m \geq 3$, $n \geq 1$, any central profile $P \in \mathcal{L}(\mathcal{A}_m)^n$, any dispersion vector $\vec{\varphi} \in (0,1]^n$, and any $t > 0$, let $\varphi^* = \frac{1}{n}\sum_{j \leq n} \min(m^2\varphi_j, \frac{m\varphi_j}{(1-\varphi_j)^2(1-\varphi_j^2)})$ and $d = \lceil \overline{KT}(P) + 2\varphi^* + t \rceil$. We have*

$$\Pr_{P' \sim (P,\vec{\varphi})}(\textit{Time}_{Alg_{KS}}(P') > O(16^d(d^2n^2m^2\log m)) < \exp\left(-\frac{2nt^2}{m^2(m-1)^2}\right)$$

Theorem 3 contains a parameter t to control the tradeoff between the probability and the runtime guarantee. A larger t corresponds to a larger d, which corresponds to a higher upper bound on the runtime. Meanwhile, the probability for the runtime of Alg_{KS} to exceed the upper bound in the theorem decreases exponentially in t.

Proof. Betzler et al. (2009) proved that the runtime of Alg_{KS} is $O(16^{\bar{d}}(\bar{d}^2n^2m^2\log m))$, where $\bar{d} = \lceil \overline{\text{KT}}(P') \rceil$. Therefore, it suffices to prove that with high probability, the average KT distance in P' is at most d.

We first prove a claim on the expected KT distance between the central ranking and randomly generated ranking.

Claim 6. *For any single-agent Mallows model and any $\theta = (R, \varphi)$, we have*

$$\mathbb{E}_{W\sim\theta}\textit{KT}(R, W) \leq \min(m^2\varphi, \frac{m\varphi}{(1-\varphi)^2(1-\varphi^2)})$$

The proof is done by directly calculating $\mathbb{E}_{W\sim\theta}\text{KT}(R, W)$ using Mallows (1957)'s closed-form formulas for probability of pairwise comparisons.

Claim 7. *We have:*

$$\Pr(\overline{KT}(P') > \overline{KT}(P) + 2\varphi^* + t) \leq \exp\left(-\frac{2nt^2}{m^2(m-1)^2}\right)$$

Proof. For any pair of agents $1 \leq j_1 < j_2 \leq n$, let R'_{j_1} and R'_{j_2} denote their rankings in P', respectively. We have $\text{KT}(R'_{j_1}, R'_{j_2}) \leq \text{KT}(R_{j_1}, R_{j_2}) + \text{KT}(R'_{j_1}, R_{j_1}) + \text{KT}(R'_{j_2}, R_{j_2})$. Therefore, $\mathbb{E}(\overline{\text{KT}}(P')) = \overline{\text{KT}}(P) + \frac{(n-1)\sum_{j=1}^n \text{KT}(R'_j, R_j)}{n(n-1)/2} = \overline{\text{KT}}(P) + \frac{2\sum_{j=1}^n \text{KT}(R'_j, R_j)}{n}$. Notice that for each $j \leq n$, $\text{KT}(R'_j, R_j)$ is a random variable in $[0, m(m-1)/2]$ whose mean is no more than $\min(m^2\varphi_j, \frac{m\varphi_j}{(1-\varphi_j)^2(1-\varphi_j^2)})$. Let $S_n = 2\sum_{j=1}^n \text{KT}(R'_j, R_j)$. We have that $\mathbb{E}(S_n) \leq 2\varphi^* n$. Therefore, for any $t > 0$, by Hoeffding's inequality, we have:

$$\Pr(\frac{S_n}{n} > 2\varphi^* + t) \leq \Pr(\frac{S_n}{n} > \mathbb{E}(\frac{S_n}{n}) + t)$$
$$\leq \Pr(S_n > \mathbb{E}(S_n) + nt) \leq \exp\left(-\frac{2(nt)^2}{n(m(m-1))^2}\right)$$
$$= \exp\left(-\frac{2nt^2}{m^2(m-1)^2}\right)$$

□

The theorem follows after Claim 7. □

We note that $\varphi^* \leq \frac{m^2}{n}\sum_{j=1}^n \varphi_j$ and $\varphi^* \leq \frac{m}{n}\sum_{j=1}^n \frac{\varphi}{(1-\varphi)^2(1-\varphi^2)}$, and the former upper bound is the average dispersion of the agents multiplied by m^2. Neither upper bound implies the other. The former is stronger when some φ_j is close to 1. The latter is stronger when all $\varphi_j = O(\frac{1}{m})$. We immediately obtain the following corollary by combining the former upper bound with Theorem 3.

Corollary 1. *Given $\vec{\mathcal{M}}_{Ma,m}$, any $n = \Omega(m^4)$, and any parameter $(P, \vec{\varphi})$ such that (1) $\overline{KT}(P) = O(\log m + \log n)$, and (2) $\frac{\vec{\varphi}\cdot\vec{1}}{n} = O(\frac{\log m + \log n}{m^2})$, we have:*

$$\Pr_{P' \sim (P,\vec{\varphi})}(\textit{Time}_{\textit{Alg}_{KR}}(P') = \omega(\textit{poly}(mn)^t)) = \exp(-\Omega(t^2))$$

Corollary 1 states that when n is sufficiently large, and the average KT distance in P and the average dispersion are not too large, with high probability Alg_{KR} solves KEMENY RANKING in polynomial time.

5 An Attempt to Apply Bläser and Manthey's (2015) Framework

In Bläser and Manthey's (2015) framework, the smoothed runtime of an algorithm is analyzed w.r.t. a *perturbation model* over the input, denoted by $\mathcal{D} = \{D_{\ell,x,\phi}\}$, where each $D_{\ell,x,\phi}$ is a distribution over data, x represents the "ground truth", ϕ is the maximum probability of any data point under $D_{\ell,x,\phi}$, and ℓ is the size of x. Let $N_{\ell,x}$ denote the size of support of $D_{\ell,x,\phi}$, i.e. the number of data points that receive non-zero probability under $D_{\ell,x,\phi}$. Bläser and Manthey (2015) defined the following notion of smoothed poly-time algorithms (some technical assumptions are omitted for better presentation).

Definition 8 (BM-Smoothed poly-time (Bläser and Manthey 2015)). *Given a perturbation model $\mathcal{D} = \{D_{\ell,x,\phi}\}$, an algorithm Alg is* BM-smoothed poly-time *if there exists $\epsilon > 0$ such that for all $D_{\ell,x,\phi} \in \mathcal{D}$,*

$$\mathbb{E}_{y\sim D_{\ell,x,\phi}}(\textit{Time}_{\textit{Alg}}(y)^{\epsilon}) = O(\ell \cdot \phi \cdot N_{\ell,x}) \qquad (3)$$

In words, Alg is BM-smoothed poly-time if for every distribution $D_{\ell,x,\phi}$, the expected runtime$^{\epsilon}$ for some constant $\epsilon > 0$ is linear in n and $\phi \cdot N_{\ell,x}$. As commented by Bläser and Manthey (2015), this does not mean that the expected runtime of the algorithm is polynomial in the input size. While this formulation is adopted to build a sound theory for smoothed complexity analysis, we feel that its relevance in social choice is not clear.

Brute force search is BM-smoothed poly-time. Let BF denote the brute force search algorithm that first computes the Kemeny scores (respectively, Slater scores) for all $m!$ rankings, then chooses a ranking with the minimum score. Note that the runtime of BF is $\Omega(m!)$ for any input. The following simple observation states that BF is BM-smoothed poly-time for a large class of models.

Proposition 1. *For any fixed $0 < \underline{\varphi} \leq \overline{\varphi} < 1$, when $m \geq 2^{(3-\overline{\varphi})/(1-\overline{\varphi})}$, BF is BM-smoothed poly-time for* KEMENY RANKING *and* SLATER RANKING *w.r.t. $\vec{\mathcal{M}}_{Ma}^{[\underline{\varphi},\overline{\varphi}]}$, where the dispersion parameter is discretized.*

Proof. We first translate $\vec{\mathcal{M}}_{\text{Ma}}^{[\underline{\varphi},\overline{\varphi}]}$ to the $D_{\ell,x,\phi}$ notation. For any m, n, and any $\vec{\pi} \in \Pi_m^n$, let $x = \vec{\pi}$ represent the central rankings and dispersion parameters for the n agents. Therefore, $\ell = \Theta(nm\log m)$, $N_{\ell,x} = (m!)^n$, and $\phi \geq (\frac{1}{Z_{\underline{\varphi}}})^n \geq (1-\overline{\varphi})^{n(m-1)}$. Let $\epsilon = \frac{1}{2}$. For any n-profile $P' \in \mathcal{L}(\mathcal{A}_m)^n$, $\text{Time}_{\text{BF}}(P') = O(m!nm^2)$ for both KEMENY RANKING and SLATER RANKING. Therefore, $\mathbb{E}_{P'\sim D_{\ell,x,\phi}}(\text{Time}_{\text{BF}}(P')^{\epsilon}) = O((m!nm^2)^{\epsilon})$. To prove the proposition, it suffices to prove that for any $m > 2^{(3-\overline{\varphi})/(1-\overline{\varphi})}$ and $n \geq 1$,

$$nm\log m(m!(1-\overline{\varphi})^{m-1})^n > (m!nm^2)^{1/2}$$

This is done in the following series of inequalities.

$$\begin{aligned}
&m > 2^{(3-\overline{\varphi})/(1-\overline{\varphi})} \Leftrightarrow m\log m > m(\frac{2}{1-\overline{\varphi}}+1)\\
\Rightarrow &(m+\frac{1}{2})\log m - m > (m-1)\frac{2}{1-\overline{\varphi}}\\
\Rightarrow &\log m! > (m-1)\frac{2}{1-\overline{\varphi}} \quad \text{(Stirling's formula)}\\
\Rightarrow &(m!)^{1-\frac{1}{2n}}(1-\overline{\varphi})^{m-1} > 1\\
\Rightarrow &nm\log m(m!(1-\overline{\varphi})^{m-1})^n > (m!nm^2)^{1/2}
\end{aligned}$$

□

Proposition 1 may appear paradoxical because it calls the always-exponential-time BF (BM-smoothed) poly-time. Technically, this is allowed in Bläser and Manthey's (2015) framework. As Bläser and Manthey (2015) commented, the expected runtime$^{\epsilon}$ in (3) is allowed to be exponentially large when $N_{\ell,x}\phi$ is exponentially large, i.e., the perturbation (which corresponds to $1/\phi$) is relatively small compared to the size of support of the distributions (i.e., $N_{\ell,x}$). Therefore, we believe that the real power of Bläser and Manthey's (2015) framework is in the analysis of scenarios where the perturbation is large compared to the size of the support of the distributions, i.e., $N_{\ell,x}\phi$ is not too large. Unfortunately, this is not the case for the Mallows-series model $\vec{\mathcal{M}}_{\text{Ma}}^{[\underline{\varphi},\overline{\varphi}]}$, which we believe to be a reasonable model for social choice. It is possible that when $\underline{\varphi}$ and $\overline{\varphi}$ are allowed to change as m increases, BF is no longer a smoothed-poly-time algorithm in Bläser and Manthey's (2015) framework.

6 Summary and Future Work

We prove the smoothed hardness of Kemeny and Slater, and a parameterized typical-case smoothed easiness result for Kemeny. An immediate open question is the smoothed complexity of the Dodgson rule and the Young rule. Smoothed complexity of other problems in computational social choice is also an obvious open question as Baumeister, Hogrebe, and Rothe (2020) pointed out.

Acknowledgments

LX acknowledges support from NSF #1453542 and #1716333, ONR #N00014-171-2621, and a gift fund from Google. We thank the anonymous reviewers for their very helpful comments and suggestions.

Broader Impact

This paper aims to understand the smoothed complexity of computing commonly-studied voting rules, which are important tools for collective decision making. The results will be important to multi-agent systems, where voting is used to achieve consensus. Success of the research will benefit general public beyond the CS research community because voting is a key component of democracy.

References

Ailon, N.; Charikar, M.; and Newman, A. 2008. Aggregating inconsistent information: Ranking and clustering. *Journal of the ACM* 55(5): Article No. 23.

Ali, A.; and Meila, M. 2012. Experiments with Kemeny ranking: What works when? *Mathematical Social Sciences* 64(1): 28–40.

Alon, N. 2006. Ranking tournaments. *SIAM Journal of Discrete Mathematics* 20: 137–142.

Bartholdi, III, J.; Tovey, C.; and Trick, M. 1989. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare* 6: 157–165.

Baumeister, D.; Hogrebe, T.; and Rothe, J. 2020. Towards Reality: Smoothed Analysis in Computational Social Choice. In *Proceedings of AAMAS*, 1691–1695.

Betzler, N.; Fellows, M. R.; Guo, J.; Niedermeier, R.; and Rosamond, F. A. 2009. Fixed-parameter algorithms for Kemeny rankings. *Theoretical Computer Science* 410: 4554–4570.

Bläser, M.; and Manthey, B. 2015. Smoothed Complexity Theory. *ACM Transactions on Computation Theory* Article 6.

Blum, A.; and Dunagan, J. D. 2002. Smoothed Analysis of the Perceptron Algorithm for Linear Programming. In *Proceedings of SODA*, 905–914.

Bogdanov, A.; and Trevisan, L. 2006. Average-Case Complexity. *Foundations and Trends in Theoretical Computer Science* 2(1): 1–106.

Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.

Conitzer, V. 2006. Computing Slater Rankings Using Similarities among Candidates. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 613–619. Boston, MA, USA. Early version appeared as IBM RC 23748, 2005.

Conitzer, V.; Davenport, A.; and Kalagnanam, J. 2006. Improved Bounds for Computing Kemeny Rankings. In *Proceedings of the National Conference on Artificial Intelligence*, 620–626.

Cornaz, D.; Galand, L.; and Spanjaard, O. 2013. Kemeny elections with bounded single-peaked or single-crossing width. In *Proceedings of IJCAI*, 76–82.

Davenport, A.; and Kalagnanam, J. 2004. A Computational Study of the Kemeny Rule for Preference Aggregation. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 697–702. San Jose, CA, USA.

Doignon, J.-P.; Pekeč, A.; and Regenwetter, M. 2004. The repeated insertion model for rankings: Missing link between two subset choice models. *Psychometrika* 69(1): 33–54.

Hemaspaandra, E.; Spakowski, H.; and Vogel, J. 2005. The complexity of Kemeny elections. *Theoretical Computer Science* 349(3): 382–391.

Huang, L.-S.; and Teng, S.-H. 2007. On the Approximation and Smoothed Complexity of Leontief Market Equilibria. In *Proceedings of FAW*, 96–107.

Karpinski, M.; and Schudy, W. 2010. Faster Algorithms for Feedback Arc Set Tournament, Kemeny Rank Aggregation and Betweenness Tournament. In *Proceedings of ISAAC*, 3–14.

Kenyon-Mathieu, C.; and Schudy, W. 2007. How to Rank with Few Errors: A PTAS for Weighted Feedback Arc Set on Tournaments. In *Proceedings of the Thirty-ninth Annual ACM Symposium on Theory of Computing*, 95–103. San Diego, California, USA.

Levin, L. A. 1986. Average case complete problems. *SIAM Journal on Computing* 15(1): 285—286.

Mallows, C. L. 1957. Non-null ranking model. *Biometrika* 44(1/2): 114–130.

Perrot, K.; and Pham, T. V. 2015. Feedback Arc Set Problem and NP-Hardness of Minimum Recurrent Configuration Problem of Chip-Firing Game on Directed Graphs 19: 373–396.

Rothe, J.; Spakowski, H.; and Vogel, J. 2003. Exact Complexity of the Winner Problem for Young Elections. In *Theory of Computing Systems*, volume 36(4), 375–386. Springer-Verlag.

Spielman, D. A.; and Teng, S.-H. 2004. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM* 51(3).

Spielman, D. A.; and Teng, S.-H. 2009. Smoothed Analysis: An Attempt to Explain the Behavior of Algorithms in Practice. *Communications of the ACM* 52(10): 76–84.

van Zuylen, A.; and Williamson, D. P. 2007. Deterministic Algorithms for Rank Aggregation and Other Ranking and Clustering Problems. In *Proceedings of WOWA*, 260–273.

Xia, L. 2020. The Smoothed Possibility of Social Choice. In *Proceedings of NeurIPS*.

Young, H. P. 1974. An axiomatization of Borda's rule. *Journal of Economic Theory* 9(1): 43–52.

Zwicker, W. S. 2018. Cycles and Intractability in a Large Class of Aggregation Rules. *Journal of Artificial Intelligence Research* 61(1).