
Bayesian Estimators As Voting Rules

Lirong Xia

Rensselaer Polytechnic Institute

Troy, NY, USA

xial@cs.rpi.edu

Abstract

We investigate the fairness of Bayesian estimators (BEs) by viewing them as (irresolute) voting rules and evaluating them by satisfaction of desirable social choice axioms. We characterize the class of BEs that satisfy *neutrality* by the class of BEs with neutral structures. We prove that a BE with a neutral structure is a minimax rule if it further satisfies *parameter connectivity*. We prove that no BE satisfies *strict Condorcet criterion*. We also propose three new BEs of natural frameworks and investigate their computational complexity and satisfaction of *monotonicity* and *Condorcet criterion*.

1 INTRODUCTION

Bayesian estimators have been widely applied in rank aggregation. For example, IMDb uses Bayesian estimators to aggregate users' votes to create the top-250 movie list [2]. However, users have complaint that such mechanisms are "unfair" because the rank of a seemingly good film is not high [1]. While this particular complaint may not be hard to address, ideally we would like to use a fair rank aggregation method with high statistical efficiency. This raises the following important questions.

Q1. How can we measure fairness in rank aggregation?

Q2. How can we design fair Bayesian estimators?

Same questions arise in many other rank aggregation situations, especially those where the voting agents are human beings. For example, in political domains, important public decisions are made by aggregating citizens' votes; in low-stakes voting scenarios, friends vote to decide the place for dinner; in crowdsourcing, online workers' noisy answers are aggregated to estimate the correct answer [22].

Q1 has been partially answered by social choice theory. Following Arrow's celebrated impossibility theorem [4],

various kinds of measures on fairness, called *axioms*, have been formulated and used to evaluate voting rules in political elections. For example, the *anonymity* axiom states that the voting rule is insensitive to permutations over agents' votes, which can be seen as fairness for voters; *neutrality* is a fairness condition for the alternatives; and *Condorcet criterion* (informally) states that an obviously socially strong alternative should win, which is similar in spirit to the complaint by the IMDb user. The axiomatic approach has gone beyond political elections to e.g. ranking systems [3], recommender systems [25], and community detection [9].

While there has been a growing literature on statistical properties of commonly studied voting rules, there is little work in the reverse direction, i.e. studying the satisfaction of social choice axioms for commonly studied statistical estimators, especially Bayesian estimators. Recently Azari Soufiani et al. [7] proposed a statistical decision-theoretic framework (framework for short) to obtain new voting rules as Bayesian estimators, and investigated the satisfaction of some axioms for two Bayesian estimators. To the best of our knowledge, there is no general characterizations of social choice axioms for Bayesian estimators.

Our Contributions. We study the satisfaction of axioms for Bayesian estimators (BEs) under the framework proposed by Azari Soufiani et al. [7]. We answer Q2 for two well-studied axioms: *neutrality* and *strict Condorcet criterion*. We characterize BEs that satisfy neutrality by the BEs of *neutral* frameworks. Therefore, to design neutral BEs we only need to focus on neutral frameworks. We also prove that no BE satisfies strict Condorcet criterion.

In addition, we prove that if a neutral framework satisfies *parameter connectivity*, then its BE is a minimax rule, which means that the BE is optimal w.r.t. the worst-case frequentist expected loss. We believe that this result is of independent interest.

We also analyze the satisfaction of *Condorcet criterion*, *monotonicity*, and computational complexity for four classes of BEs. Each BE in each class is identified by a

	Anonymity	Strict Condorcet	Neutrality	Minimax	Condorcet	Monotonicity	Comp.
BEs			Y/N	Y/N	Y/N	Y/N	P/NP-hard
$f_{Ma,\varphi}^{\text{Top}}$	Y (trivial)	N (Thm. 4)	Y (Thm. 2)	Y (Thm. 1)	Y iff $\frac{\varphi(1-\varphi^{m-1})}{1-\varphi} \leq 1$ (Thm. 5)	Y [5]	NP-hard [25]
$f_{Co,\varphi}^{\text{Borda}}$					Y iff $\varphi \leq \frac{1}{m-1}$ (Thm. 7)	Y (Prop. 1)	P (Thm. 6)
$f_{\text{Pair},\varphi}^1$					Y iff $\varphi \leq \frac{1}{m-1}$ (Thm. 10)	Y (Prop. 2)	P (Thm. 9)
$f_{\text{Pair},\varphi}^2$					N (Thm. 10)		

Table 1: Main results. m is the number of alternatives.

dispersion value $0 < \varphi < 1$. The first class has been studied [30, 27, 7] while the remaining three classes are new. The four classes are (1) $f_{Ma,\varphi}^{\text{Top}}$ is the BE of Mallows’ model with the top loss function. Condorcet criterion has been studied for $f_{Ma,\varphi}^{\text{Top}}$ for $\varphi > \frac{1}{\sqrt{2}}$ but the remaining cases are open [7]¹. (2) $f_{Co,\varphi}^{\text{Borda}}$ is the BE of Condorcet’s model with the Borda loss function. (3) $f_{\text{Pair},\varphi}^1$ and (4) $f_{\text{Pair},\varphi}^2$ are the BEs of a new model with different loss functions, where a parameter can be interpreted as the “strongest pairwise comparison”. Our results are summarized in Table 1.

The second row in Table 1 are results for general BEs. A “Y/N” means that some Bayesian estimators satisfy the axiom and some do not. For $f_{Ma,\varphi}^{\text{Top}}$, we prove a dichotomy theorem on its satisfaction of Condorcet criterion: it satisfies the Condorcet criterion if and only if $\frac{\varphi(1-\varphi^{m-1})}{1-\varphi} \leq 1$, where m is the number of alternatives (Theorem 5). We also proved similar dichotomy theorems for $f_{Co,\varphi}^{\text{Borda}}$ and $f_{\text{Pair},\varphi}^1$, where the threshold is $\frac{1}{m-1}$ (Theorem 7 and 10). We would like to highlight two new classes of BEs: $f_{Co,\varphi}^{\text{Borda}}$ and $f_{\text{Pair},\varphi}^1$, because they can satisfy all axioms studied in this paper (except strict Condorcet criterion, which is not satisfied by any BE) and can be computed in polynomial time.

In addition to satisfaction of axioms, we also study the limiting cases of the three new BEs as $\varphi \rightarrow 0$ and $\varphi \rightarrow 1$. While all classes converge to refinements of the Borda rule as $\varphi \rightarrow 1$, they converge to refinements of different rules as $\varphi \rightarrow 0$. Interestingly, $f_{Co,\varphi}^{\text{Borda}}$ converges to a refinement of Copeland_{0.5} (Theorem 8) and for any $\varphi \leq \frac{1}{m-1}$, $f_{\text{Pair},\varphi}^1$ is a refinement of maximin (Theorem 11).

Related Work and Discussions. As discussed above our theorems on neutrality and strict Condorcet criterion answer Q2 for the two axioms. We are not aware of other general results on satisfaction of axioms for Bayesian estimators. In particular, Azari Soufiani et al. [7] studied the satisfaction of some axioms for two classes of BEs but did not obtain general results for BEs.

Most previous work at the intersection of social choice and statistics focused on computational aspects of the *maximum likelihood estimators (MLEs)* of various ranking models [14, 10, 15, 17, 29, 20, 24, 27, 5, 6, 16, 19]. The focuses of our work are different. We focus on Bayesian estimators, which are more general than MLEs, and we focus on the satisfaction of axioms rather than computation.

Minimax rules for various statistical models with continuous parameter spaces have been characterized by Berger [8]. Choirat and Seri [12] provided a sufficient condition on discrete-parameter models for MLEs to be minimax. In the social choice context, Caragiannis et al. [11] proved that the uniformly randomized MLE has the least sample complexity w.r.t. Mallows’ model, which is equivalent to minimaxity. Our minimaxity proof can be seen as an application of techniques by Berger [8] to social choice frameworks. As we will see in Corollary 1, our results can be easily applied to Mallows’ model and other models.

Our work is also related to statistical justification of commonly studied voting rules. Conitzer and Sandholm [14] studied whether some commonly studied voting rules can be rationalized as MLEs of some statistical models. They showed that if a voting rule does not satisfy *consistency*, then it cannot be an MLE. Pivato [26] further investigated voting rules that can be viewed as MLEs, maximum a posteriori estimators, and Bayesian estimators. Our impossibility theorem on strict Condorcet criterion can be used to prove that a voting rule *cannot* be justified a Bayesian estimator. In Corollary 2, we show that a number of voting rules including Copeland₁ and maximin are not Bayesian estimators. On the other hand, we prove that as $\varphi \rightarrow 0$, $f_{Co,\varphi}^{\text{Borda}}$ converges to a refinement of Copeland_{0.5} (Theorem 8), and for all for any $\varphi < \frac{1}{m-1}$, $f_{\text{Pair},\varphi}^1$ is a refinement of maximin (Theorem 11). Therefore, $f_{\text{Pair},\varphi}^1$ for $\varphi < \frac{1}{m-1}$ are desirable refinements of maximin because they can be justified as BEs. Previously it was only known that a refinement of Kemeny is a BE [27] and a refinement of Tideman’s rule is a BE [16].

¹The original paper has a typo on the direction of the inequality.

2 PRELIMINARIES

Let $\mathcal{A} = \{a_1, \dots, a_m\}$ denote a set of m alternatives and let $\mathcal{L}(\mathcal{A})$ denote the set of all linear orders over \mathcal{A} . Let n denote the number of agents. Each agent's vote is a linear order in $\mathcal{L}(\mathcal{A})$. The collection P of all agents' votes is called a *profile*. An *irresolute voting rule* r maps each profile to a non-empty set of winning alternatives. That is, $r : \bigcup_{n=1}^{\infty} \mathcal{L}(\mathcal{A})^n \rightarrow (2^{\mathcal{A}} \setminus \{\emptyset\})$.

For example, an irresolute *positional scoring rule* is characterized by a scoring vector $\vec{s} = (s_1, \dots, s_m)$ with $s_1 \geq s_2 \geq \dots \geq s_m$. For any alternative a and any linear order V , we let $\vec{s}(V, a) = s_j$, where j is the rank of a in V . Given a profile P , an irresolute positional scoring rule chooses all alternatives a with maximum $\sum_{V \in P} \vec{s}(V, a)$, where P is viewed as a multi-set of votes. The Borda rule is a positional scoring rule with $\vec{s} = (m-1, m-2, \dots, 1)$.

For any profile P and any pair of alternatives a, b , we let $P(a \succ b)$ denote the number of votes in P where a is preferred to b . The *weighted majority graph* of P , denoted by $\text{WMG}(P)$ is a directed weighted graph where the weight $w_P(a, b)$ on any edge $a \rightarrow b$ is $w_P(a, b) = P(a \succ b) - P(b \succ a)$. Clearly $w_P(a, b) = -w_P(b, a)$.

Given $0 \leq \alpha \leq 1$, the Copeland $_{\alpha}$ score of an alternative a in a profile P is the number of alternatives beaten by a in head-to-head competitions plus α times the alternatives tied with a . Copeland $_{\alpha}$ chooses all alternatives with the maximum Copeland $_{\alpha}$ score as the winners. The *maximin* rule chooses all alternatives a with the maximum *min-score*. The min-score of a is $\min_b w_P(a, b)$.

We will focus on the following axioms in this paper. An irresolute r satisfies

- *anonymity*, if r is insensitive to permutations over agents;
- *neutrality*, if r is insensitive to permutations over alternatives;
- *monotonicity*, if for any P , any $a \in r(P)$, and any P' that is obtained from P by only raising the positions of a , we have $a \in r(P')$;
- *Condorcet criterion*, if for any profile P , whenever a *Condorcet winner* a exists, it must be the unique winner. That is, $r(P) = \{a\}$. A Condorcet winner is an alternative that beats all other alternatives in their head-to-head competitions;
- *strict Condorcet criterion* [18], if for any profile P , whenever the set of *weak Condorcet winners* is non-empty, it must be the output of r . A weak Condorcet winner is an alternative that never loses to any other alternative in their head-to-head competition.

Azari Soufiani et al. [7] defined a *statistical decision-theoretic framework for social choice* (framework for short) to be a tuple $\mathcal{F} = (\mathcal{M}_{\mathcal{A}}, \mathcal{D}, L)$, where \mathcal{A} is the set of alternatives, $\mathcal{M}_{\mathcal{A}} = (\Theta, \vec{\pi})$ is a parametric ranking model, \mathcal{D} is the decision space, and $L : \Theta \times \mathcal{D} \rightarrow \mathbb{R}$ is a loss function.

$\mathcal{M}_{\mathcal{A}} = (\Theta, \vec{\pi})$ has two parts: a *parameter space* Θ and a set of probability distributions $\vec{\pi} = \{\pi_{\theta} : \theta \in \Theta\}$ over $\mathcal{L}(\mathcal{A})$. Agents' votes are generated i.i.d. according to $\mathcal{M}_{\mathcal{A}}$, which means that the sample space is $\mathcal{L}(\mathcal{A})^n$ and is omitted for simplicity. In this paper we focus on frameworks with finite parameter spaces and finite decision spaces.

We now recall two popular parametric ranking models. For any pair of linear orders V, W in $\mathcal{L}(\mathcal{A})$, let $\text{Kd}(V, W)$ denote the *Kendall-tau distance* between V and W , which is the total number of pairwise disagreements between V and W .

Definition 1 (Mallows' model with fixed dispersion [21]). *Given $0 < \varphi < 1$, the Mallows model with fixed dispersion φ is denoted by $\mathcal{M}_{\text{Ma}, \varphi} = (\mathcal{L}(\mathcal{A}), \vec{\pi})$, where the parameter space is $\mathcal{L}(\mathcal{A})$ and for any $V, W \in \mathcal{L}(\mathcal{A})$, $\pi_W(V) = \frac{1}{Z} \varphi^{\text{Kd}(V, W)}$, where Z is the normalization factor with $Z = \sum_{V \in \mathcal{L}(\mathcal{A})} \varphi^{\text{Kd}(V, W)}$.*

Let $\mathcal{B}(\mathcal{A})$ denote the set of all irreflexive, antisymmetric, and total binary relations over \mathcal{A} . We have $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{B}(\mathcal{A})$ and the Kendall-tau distance can be easily extended to $\mathcal{B}(\mathcal{A})$ by counting the number of pairwise disagreements.

Definition 2 (Condorcet's model [13, 30]). *Given $0 < \varphi < 1$, the Condorcet model is denoted by $\mathcal{M}_{\text{Co}, \varphi} = (\mathcal{B}(\mathcal{A}), \vec{\pi})$, where the parameter space is $\mathcal{B}(\mathcal{A})$ and for any $W \in \mathcal{B}(\mathcal{A})$ and $V \in \mathcal{L}(\mathcal{A})$, $\pi_W(V) = \frac{1}{Z} \varphi^{\text{Kd}(V, W)}$, where Z is the normalization factor.*

Next, we give three examples of loss functions. When $\Theta = \mathcal{D}$, the 0-1 *loss function*, denoted by $L_{0,1}(\theta, d)$, outputs 0 if $\theta = d$, otherwise it outputs 1. When $\mathcal{D} = \mathcal{A}$ and Θ is $\mathcal{L}(\mathcal{A})$ or $\mathcal{B}(\mathcal{A})$, the *top loss function*, denoted by $L_{\text{top}}(\theta, d)$, outputs 0 if for all other alternatives $c \in \mathcal{A}$, $d \succ c$ in θ , otherwise it outputs 1. The *Borda loss function*, denoted by $L_{\text{Borda}}(\theta, d)$, outputs the number of alternatives that are preferred to d in θ , that is, $L_{\text{Borda}}(\theta, d) = \#\{c \in \mathcal{A} : c \succ_{\theta} d\}$. All loss functions can be naturally generalized to evaluate a subset D of \mathcal{D} by computing the average loss of the decisions in D . More precisely, for any $D \subseteq \mathcal{D}$ and any $\theta \in \Theta$, we let $L(\theta, D) = \sum_{d \in D} L(\theta, d) / |D|$.

Given a framework \mathcal{F} , the *Bayesian expected loss* of $d \in \mathcal{D}$ given a profile P is $\text{EL}_{\mathcal{F}}(d|P) = \sum_{\theta \in \Theta} \text{Pr}(\theta|P) L(\theta, d)$. The subscript \mathcal{F} is often omitted without introducing confusions. In this paper we focus on the uniform prior. The *Bayesian estimator* of \mathcal{F} , denoted by $\text{BE}_{\mathcal{F}}$, takes a profile P as input and outputs all decisions with minimum expected Bayesian loss. That is, $\text{BE}_{\mathcal{F}}(P) = \arg \min_{d \in \mathcal{D}} \text{EL}(d|P)$.

Let $f_{\text{Ma}, \varphi}^{\text{Top}}$ denote the Bayesian estimator of the framework $(\mathcal{M}_{\text{Ma}, \varphi}, L_{\text{top}})$. It was proved by Azari Soufiani et al. [7] that $f_{\text{Ma}, \varphi}^{\text{Top}}$ satisfy anonymity, neutrality, monotonicity, but fails to satisfy the Condorcet criterion for some φ . Let $f_{\text{Co}, \varphi}^{\text{Borda}}$ denote the Bayesian estimator of the framework $(\mathcal{M}_{\text{Co}, \varphi}, L_{\text{Borda}})$. We will study the satisfaction of axioms

for $f_{\text{Co},\varphi}^{\text{Borda}}$.

Given a framework \mathcal{F} , a parameter $\theta \in \Theta$, $n \in \mathbb{N}$, and a voting rule r , the *frequentist loss* $\text{FL}_n(\theta, r)$ is the expected loss of the output of r against θ for randomly generated profiles of n votes. More precisely,

$$\text{FL}_n(\theta, r) = \sum_{P_n \in \mathcal{L}(\mathcal{A})^n} \pi_\theta(P_n) L(\theta, r(P_n))$$

Definition 3 ([8]). *Given a framework $\mathcal{F} = (\mathcal{M}_{\mathcal{A}}, \mathcal{D}, L)$, a voting rule r is minimax, if $r \in \arg \min_{r^*} \max_{\theta \in \Theta} \text{FL}_n(\theta, r^*)$.*

That is, a minimax rule minimizes the worst-case frequentist loss among all deterministic or randomized rules. Minimality is an important statistical criteria for decision functions—a minimax rule is the most robust rule against the adversarial nature who controls the true state of the world (the parameter). A minimax rule can be seen as having the minimum *sample complexity* [11].

3 NEUTRAL FRAMEWORKS AND MINIMAXITY

We first define the neutrality of a framework for general decision spaces. Intuitively, a framework $\mathcal{F} = (\mathcal{M}_{\mathcal{A}}, \mathcal{D}, L)$ is neutral if and only all of its three components are neutral w.r.t. permutations σ over \mathcal{A} . Because σ may not be well-defined for the parameter space and the decision space, we require the existence of homomorphisms from the permutation group over \mathcal{A} to the permutation groups over Θ and \mathcal{D} , respectively. Formally, we have the following definition.

Definition 4. *A framework $\mathcal{F} = (\mathcal{M}_{\mathcal{A}}, \mathcal{D}, L)$ where $\mathcal{M}_{\mathcal{A}} = (\Theta, \bar{\pi})$ is neutral, if each permutation σ over \mathcal{A} is mapped to a permutation σ_Θ over Θ and a permutation $\sigma_{\mathcal{D}}$ over \mathcal{D} that satisfy the following conditions.*

(i) Homomorphism. *For any pair of permutations γ and β over \mathcal{A} , $(\gamma \circ \beta)_\Theta = \gamma_\Theta \circ \beta_\Theta$ and $(\gamma \circ \beta)_{\mathcal{D}} = \gamma_{\mathcal{D}} \circ \beta_{\mathcal{D}}$.*

(ii) Model neutrality. *For any $\theta \in \Theta$, any $V \in \mathcal{L}(\mathcal{A})$, and any permutation σ over \mathcal{A} , we have $\pi_\theta(V) = \pi_{\sigma_\Theta(\theta)}(\sigma(V))$.*

(iii) Loss function neutrality. *For any $\theta \in \Theta$, any $d \in \mathcal{D}$, and any permutation σ over \mathcal{A} , we have $L(\theta, d) = L(\sigma_\Theta(\theta), \sigma_{\mathcal{D}}(d))$.*

Example 1. *For any $0 < \varphi < 1$, $(\mathcal{M}_{\text{Ma},\varphi}, \mathcal{A}, L_{\text{top}})$, $(\mathcal{M}_{\text{Ma},\varphi}, \mathcal{A}, L_{\text{Borda}})$, $(\mathcal{M}_{\text{Co},\varphi}, \mathcal{A}, L_{\text{top}})$, $(\mathcal{M}_{\text{Co},\varphi}, \mathcal{A}, L_{\text{Borda}})$ are neutral, where $\sigma_\Theta = \sigma_{\mathcal{D}} = \sigma$.*

The main theorem of this section states that if a neutral framework further satisfies the following connectivity condition, then its Bayesian estimator is a minimax rule.

(iv) Parameter connectivity. *For any pair $\theta_1, \theta_2 \in \Theta$, there exists a permutation σ over \mathcal{A} such that $\sigma_\Theta(\theta_1) = \theta_2$.*

Theorem 1. *For any neutral framework \mathcal{F} that satisfies*

parameter connectivity and any $n \in \mathbb{N}$, $\text{BE}_{\mathcal{F}}$ is a minimax rule.²

Proof: Any deterministic Bayesian estimator $\text{BE}_{\mathcal{F}}$ can be seen as a randomized rule that chooses a single decision uniformly at random from the output of $\text{BE}_{\mathcal{F}}$. Our proof is based on the following lemma.

Lemma 1 (Section 5.3.2 III in [8]). *Given a framework \mathcal{F} . Let r_{π^*} denote a Bayesian estimator for prior π^* . If $\text{FL}_n(\theta, r_{\pi^*})$ are equal for all $\theta \in \Theta$, then r_{π^*} is minimax.*

Let r_U denote the randomized decision rule that outputs all decisions with the minimum Bayesian expected loss uniformly at random. By Lemma 1, it suffices to show that for all $\theta \in \Theta$, $\text{FL}_n(\theta, r_U)$ are equal. For any pair of parameters $\theta_1, \theta_2 \in \Theta$, we let $\sigma^{(\theta_1, \theta_2)}$ denote a permutation over \mathcal{A} such $\sigma_\Theta^{(\theta_1, \theta_2)}(\theta_1) = \theta_2$, which is guaranteed by Condition (iv).

Claim 1. *For any profile P of n votes and any $d^* \in \mathcal{D}$, $r_U(P)(d^*) > 0$ if and only if $r_U(\sigma^{(\theta_1, \theta_2)}(P))(\sigma_{\mathcal{D}}^{(\theta_1, \theta_2)}(d^*)) > 0$.*

Proof: To simplify the notation, in this proof we let σ^* denote $\sigma^{(\theta_1, \theta_2)}$. $r_U(P)$ has positive probability on d^* if and only if d^* minimizes the Bayesian loss at P under uniform prior, which is equivalent to requiring that for all $d' \in \mathcal{D}$, $\sum_\theta L(\theta, d^*) \Pr(\theta|P) \leq \sum_\theta L(\theta, d') \Pr(\theta|P)$. We have the following calculation, where $\Pr(P|\theta) = \pi_\theta(P)$.

$$\begin{aligned} & \sum_\theta L(\theta, d^*) \Pr(\theta|P) \leq \sum_\theta L(\theta, d') \Pr(\theta|P) \\ \Leftrightarrow & \sum_\theta L(\theta, d^*) \Pr(P|\theta) \leq \sum_\theta L(\theta, d') \Pr(P|\theta) \quad (1) \end{aligned}$$

$$\begin{aligned} \Leftrightarrow & \sum_\theta L(\sigma_\Theta^*(\theta), \sigma_{\mathcal{D}}^*(d^*)) \Pr(\sigma^*(P)|\sigma_\Theta^*(\theta)) \\ & \leq \sum_\theta L(\sigma_\Theta^*(\theta), \sigma_{\mathcal{D}}^*(d')) \Pr(\sigma^*(P)|\sigma_\Theta^*(\theta)) \quad (2) \end{aligned}$$

$$\begin{aligned} \Leftrightarrow & \sum_\theta L(\theta, \sigma_{\mathcal{D}}^*(d^*)) \Pr(\sigma^*(P)|\theta) \\ & \leq \sum_\theta L(\theta, \sigma_{\mathcal{D}}^*(d')) \Pr(\sigma^*(P)|\theta) \quad (3) \end{aligned}$$

$$\begin{aligned} \Leftrightarrow & \sum_\theta L(\theta, \sigma_{\mathcal{D}}^*(d^*)) \Pr(\theta|\sigma^*(P)) \\ & \leq \sum_\theta L(\theta, \sigma_{\mathcal{D}}^*(d')) \Pr(\theta|\sigma^*(P)) \quad (4) \end{aligned}$$

Therefore, d^* minimizes the Bayesian loss for P if and only if $\sigma_{\mathcal{D}}^*(d^*)$ minimizes the Bayesian loss for $\sigma^*(P)$, which means that $r_U(\sigma^*(P))$ has positive probability on $(\sigma_{\mathcal{D}}^*(d^*))$. (1) and (4) are due to Bayes' rule and the uniform prior assumption. (2) is obtained from applying σ^* , σ_Θ^* , and $\sigma_{\mathcal{D}}^*$ on both sides of the inequality and then considering the neutrality of the framework. (3) is a change of

²A similar result was presented at COMSOC-14 workshop [28].

variable names, which is possible because for any pair of parameters $\theta \neq \theta'$ and any permutation σ over \mathcal{A} , we must have $\sigma_{\Theta}(\theta) \neq \sigma_{\Theta}(\theta')$, because σ_{Θ} is a permutation over Θ . \square

We note that r_U chooses a decisions d with $r_U(P)(d) > 0$ uniformly at random. Therefore, by Claim 1, for any $\theta_1 \neq \theta_2$ and any profile P , we have $r_U(\sigma^{(\theta_1, \theta_2)}(P)) = \sigma_{\mathcal{D}}^{(\theta_1, \theta_2)}(r_U(P))$. By neutrality of \mathcal{F} , we have $L(\theta_1, r_U(P)) = L(\sigma_{\Theta}^{(\theta_1, \theta_2)}(\theta_1), \sigma_{\mathcal{D}}^{(\theta_1, \theta_2)}(r_U(P))) = L(\theta_2, r_U(\sigma^{(\theta_1, \theta_2)}(P)))$.

Finally, we have:

$$\begin{aligned} \text{FL}_n(\theta_1, r_U) &= \sum_P L(\theta_1, r_U(P)) \Pr(P|\theta_1) \\ &= \sum_P L(\theta_2, r_U(\sigma^{(\theta_1, \theta_2)}(P))) \Pr(\sigma^{(\theta_1, \theta_2)}(P)|\theta_2) \\ &= \text{FL}_n(\theta_2, r_U) \end{aligned}$$

By Lemma 1, r_U is a minimax rule. We note that for any θ and any profile P , $L(\theta, r_U(P)) = L(\theta, \text{BE}_{\mathcal{F}}(P))$. This means that $\text{BE}_{\mathcal{F}}$ is a minimax rule. \square

It is not hard to verify that all models mentioned in Example 1 satisfy parameter connectivity. Therefore, their Bayesian estimators are minimax rules. In particular, $f_{\text{Ma}, \varphi}^{\text{Top}}$ is a minimax rule for $(\mathcal{M}_{\text{Ma}, \varphi}, L_{\text{Top}})$. When the 0-1 loss function is used, the Bayesian estimator becomes *maximum likelihood estimator (MLE)*. Therefore, Theorem 1 immediately implies that MLE is minimax.

Corollary 1. *For any neutral framework $\mathcal{F} = (\mathcal{M}_{\mathcal{A}}, \mathcal{A}, L_{0-1})$, its MLE (that outputs all alternatives with the maximum likelihood) is a minimax rule.*

The special case of Corollary 1 for $\mathcal{M}_{\text{Ma}, \varphi}$ was proved by Caragiannis et al. [11]. We note that in Appendix A of [11], an example was shown to illustrate that MLE is not minimax w.r.t. L_{0-1} . This does not contradict Theorem 1 and Corollary 1 because the framework in the example is not neutral.

As shown in the following example, not all Bayesian estimators of neutral frameworks satisfy minimaxity.

Example 2. *Let $\mathcal{A} = \{a, b\}$. Consider a framework (\mathcal{M}, L) for two alternatives where $\mathcal{M} = (\Theta, \bar{\pi})$ combines two Mallows' models with dispersion parameter 0.6 and 0.7 respectively. Formally, let $\Theta = \{0.6, 0.7\} \times \{a \succ b, b \succ a\}$. For each $(\varphi, W) \in \Theta$, $\pi_{(\varphi, W)}$ is the same as π_W in Mallows' model with dispersion φ . For any $W \in \mathcal{L}(\mathcal{A})$ and $c \in \mathcal{A}$, we let $L((0.6, W), c) = L_{\text{top}}(W, c)$ and $L((0.7, W), c) = 1 - L_{\text{top}}(W, c)$.*

It can be verified that \mathcal{F} is neutral by letting γ_{Θ} be a permutation that only applies to the second component of the parameter (the ranking). Let $n = 1$. When the vote is $a \succ b$, the posterior distribution is the following.

Parameter	(0.6, $a \succ b$)	(0.6, $b \succ a$)	(0.7, $a \succ b$)	(0.7, $b \succ a$)
Post. Prob.	$\frac{1}{3.2}$	$\frac{0.6}{3.2}$	$\frac{1}{3.4}$	$\frac{0.7}{3.4}$
Loss for a	0	1	1	0

Therefore, $EL(a|\{a \succ b\}) = \frac{0.6}{3.2} < \frac{0.7}{3.4} = EL(b|\{a \succ b\})$. Therefore, $BE_{\mathcal{F}}(a \succ b) = a$. Similarly $BE_{\mathcal{F}}(b \succ a) = b$.

When the ground truth parameter is $(0.7, a \succ b)$, the frequentist expected loss of $BE_{\mathcal{F}}$ is $\frac{1}{1.7} > \frac{1}{2}$. We note that the worst-case frequentist loss of the voting rule that always output \mathcal{A} is $\frac{1}{2}$, which means that $BE_{\mathcal{F}}$ is not a minimax rule. \square

4 GENERAL RESULTS ON SATISFACTION OF AXIOMS

To analyze the satisfaction of axioms of Bayesian estimators, in the rest of this paper we focus on a special class of frameworks where the decision space is \mathcal{A} . We let $\mathcal{F} = (\mathcal{M}_{\mathcal{A}}, L)$ denote such a framework where the decision space is omitted. For neutral frameworks, we further require that $\sigma_{\mathcal{A}} = \sigma$, where $\sigma_{\mathcal{A}}$ is the corresponding permutation over the decision space, which is \mathcal{A} in this section.

Theorem 2. *The Bayesian estimator of any neutral framework satisfies neutrality.*

Proof: Let $S(\mathcal{A})$ denote the set of all permutations over \mathcal{A} . It suffices to prove that the expected loss function is insensitive to permutations. For any neutral framework \mathcal{F} , any profile P , any alternative a , and any $\gamma \in S(\mathcal{A})$, we have

$$\begin{aligned} EL(a|P) &= \sum_{\theta \in \Theta} \Pr(\theta|P) L(\theta, a) \propto \sum_{\theta \in \Theta} \Pr(P|\theta) L(\theta, a) \\ &= \sum_{\theta \in \Theta} \Pr(\gamma(P)|\gamma_{\Theta}(\theta)) L(\gamma_{\Theta}(\theta), \gamma(a)) \\ &\propto \sum_{\theta \in \Theta} \Pr(\gamma_{\Theta}(\theta)|\gamma(P)) L(\gamma_{\Theta}(\theta), \gamma(a)) = EL(\gamma(a)|\gamma(P)) \end{aligned}$$

\square

Theorem 3. *If the Bayesian estimator $BE_{\mathcal{F}}$ of a framework \mathcal{F} satisfies neutrality then there exists a neutral framework \mathcal{F}^* such that $BE_{\mathcal{F}^*} = BE_{\mathcal{F}}$.*

All missing proofs can be found in the full version on arXiv.

Theorem 4. *No Bayesian estimator satisfies strict Condorcet criterion.*

Proof: For the sake of contradiction suppose a Bayesian estimator r of $\mathcal{F} = (\mathcal{M}_{\mathcal{A}}, L)$ satisfies strict Condorcet criterion where $\mathcal{M}_{\mathcal{A}} = (\Theta, \bar{\pi})$.

We first prove that for any profile P , if alternatives a and b are tied in their head-to-head competition, then expected loss for a must be the same as the expected loss for b .

Lemma 2. *Suppose $r = BE_{\mathcal{F}}$ satisfies strict Condorcet criterion. For any profile P and any pair of alternatives (a, b) , if $w_P(a, b) = 0$ then $EL(a|P) = EL(b|P)$.*

Proof: For any distribution π over Θ , let $\mathcal{S}_\pi = \{S_1, \dots, S_p\}$ denote the partition of Θ into equivalent classes according to π , where p is the number of equivalent classes. That is, for any $S \in \mathcal{S}_\pi$ and any $\theta_1, \theta_2 \in S$, we have $\pi(\theta_1) = \pi(\theta_2)$. Let \mathcal{T}_π denote the total order over \mathcal{S}_π such that for any pair $S, S' \in \mathcal{S}_\pi$, we have $S \succ_{\mathcal{T}_\pi} S'$ if and only if the π value of parameters in S is strictly larger than the π value of parameters in S' .

For any profile P , let \mathcal{S}_P denote $\mathcal{S}_{Pr(\cdot|P)}$. That is, \mathcal{S}_P is the partition of Θ according to the posterior distribution over Θ given P . \mathcal{T}_P is defined similarly. The next lemma states that for any profile P and any pair of co-winners (a, b) , the total loss of a and b within each equivalent class in \mathcal{S}_P must be the same. For any $S \subseteq \Theta$, we let $L(S, a) = \sum_{\theta \in S} L(\theta, a)$.

Lemma 3. *Suppose $r = BE_{\mathcal{F}}$ satisfies strict Condorcet criterion. For any profile P and any $S \in \mathcal{S}_P$, if there are at least two weak Condorcet winners $\{a, b\}$ in P , then $L(S, a) = L(S, b)$.*

Proof: For the sake of contradiction suppose the lemma does not hold for a profile P where $\{a, b\}$ are two weak Condorcet winners. Let $\mathcal{T}_P = S_1 \succ S_2 \succ \dots \succ S_p$. Let S_i denote the highest-ranked equivalent class in \mathcal{T}_P such that the total loss of a and the total loss of b on S_i are different. W.l.o.g. suppose $L(S_i, a) > L(S_i, b)$. For any natural number k , it follows that a and b are also weak Condorcet winners in kP , whose weighted majority graph is exactly $WMG(P)$ times k . We next show that when k is sufficiently large, $EL(a|kP) > EL(b|kP)$. For any $i \leq p$, let $\theta_i \in S_i$ be an arbitrary parameter in S_i .

$$\begin{aligned} EL(a|kP) &= \sum_{\theta \in \Theta} \Pr(\theta|kP) L(\theta, a) \\ &\propto \sum_{\theta \in \Theta} \Pr(\theta|P)^k L(\theta, a) = \sum_{i=1}^p \sum_{\theta \in S_i} \Pr(\theta|P)^k L(\theta, a) \\ &= \sum_{i=1}^p \Pr(\theta_i|P)^k L(S_i, a) \end{aligned}$$

Because for any $i' > i$ we have $\Pr(\theta_i|P) > \Pr(\theta_{i'}|P)$, there exists $k \in \mathbb{N}$ such that $(\frac{\Pr(\theta_i|P)}{\Pr(\theta_{i+1}|P)})^k > \frac{\sum_{i=i+1}^p (L(S_i, a) - L(S_i, b))}{L(S_i, a) - L(S_i, b)}$. Therefore, for such k we have $\sum_{i=1}^p \Pr(\theta_i|P)^k L(S_i, a) > \sum_{i=1}^p \Pr(\theta_i|P)^k L(S_i, b) \propto EL(b|kP)$. This means that b cannot be a co-winner in $r(kP)$, which contradicts the assumption that r satisfies strict Condorcet criterion. \square

For any pair of partitions \mathcal{S}_1 and \mathcal{S}_2 of Θ , we let $\mathcal{S}_1 \oplus \mathcal{S}_2$ denote the coarsest partition of Θ that refines both \mathcal{S}_1 and

\mathcal{S}_2 . That is,

$$\mathcal{S}_1 \oplus \mathcal{S}_2 = \{S_1 \cap S_2 : S_1 \in \mathcal{S}_1, S_2 \in \mathcal{S}_2\} \setminus \{\emptyset\}$$

Lemma 4. *For any statistical model and any pair of profiles P_1, P_2 , there exists $k \in \mathbb{N}$ such that $\mathcal{S}_{kP_1 \cup P_2} = \mathcal{S}_{P_1} \oplus \mathcal{S}_{P_2}$.*

Proof: We let $P^* = kP_1 \cup P_2$ for a sufficiently large k such that the ‘‘gap’’ between two equivalent classes in kP is large enough that the only effect of P_2 is to refine the equivalent classes in kP . More formally, we choose $k \in \mathbb{N}$ such that for any $\theta_1, \theta_2 \in \Theta$, $\Pr(\theta_1|P_1)^k \Pr(\theta_1|P_2) > \Pr(\theta_2|P_1)^k \Pr(\theta_2|P_2)$ if and only if one of the following two conditions hold: (1) $\Pr(\theta_1|P_1) > \Pr(\theta_2|P_1)$, or (2) $\Pr(\theta_1|P_1) = \Pr(\theta_1|P_2)$ and $\Pr(\theta_1|P_2) > \Pr(\theta_2|P_2)$. \square

For any $a, b \in \mathcal{A}$, let \mathcal{L}_{ab} denote the set of all rankings where $a \succ b$. Let \mathcal{P}_{ab} denote the set of all two-agent profiles where one vote comes from \mathcal{L}_{ab} and the other vote comes from \mathcal{L}_{ba} . That is,

$$\mathcal{P}_{ab} = \{\{V_1, V_2\} : V_1 \in \mathcal{L}_{ab}, V_2 \in \mathcal{L}_{ba}\}$$

Let \mathcal{S}_{ab} denote the finest partition of Θ that refines all partitions induced by profiles in \mathcal{P}_{ab} . That is, $\mathcal{S}_{ab} = \oplus \mathcal{P}_{ab}$. By Lemma 4, there exists a profile P_{ab} such that $\mathcal{S}_{P_{ab}} = \mathcal{S}_{ab}$.

Lemma 5. *Suppose $r = BE_{\mathcal{F}}$ satisfies strict Condorcet criterion. For any $a, b \in \mathcal{A}$ and any $S \in \mathcal{S}_{ab}$, we have $L(S, a) = L(S, b)$.*

Proof: Let P^* be an arbitrary profile with the following conditions. (1) $w_{P^*}(a, b) = w_{P^*}(b, a) = 0$. (2) For any $c \notin \{a, b\}$, we have $w_{P^*}(a, c) > 0$ and $w_{P^*}(b, c) > 0$. By Lemma 4, there exists a sufficiently large $k \in \mathbb{N}$ such that both conditions still hold for $kP^* \cup P_{ab}$, and $\mathcal{S}_{kP^* \cup P_{ab}} = \mathcal{S}_{ab}$. The latter is because P^* can be seen as the union of $|P^*|/2$ profiles in \mathcal{P}_{ab} , which means that \mathcal{S}_{ab} is a refinement of \mathcal{S}_{kP^*} . The lemma follows after Lemma 3. \square

We note that for any $a, b \in \mathcal{A}$, any profile P where $w_P(a, b) = 0$ can be seen as the union of $|P|/2$ profiles in \mathcal{P}_{ab} . This means that \mathcal{S}_{ab} is a refinement of \mathcal{S}_P . Therefore, any $S \in \mathcal{S}_P$ must be the union of some equivalent classes in \mathcal{S}_{ab} . By Lemma 5 we have that $L(S, a) = L(S, b)$. We have $EL(a|P) = \sum_{S \in \mathcal{S}_P} \Pr(\theta_S|P) L(S, a) = \sum_{S \in \mathcal{S}_P} \Pr(\theta_S|P) L(S, b) = EL(b|P)$, where θ_S denote an arbitrary element in S . This proves Lemma 2. \square

Consider any profile P where $w_P(a, b) = w_P(b, c) = 0$, $w_P(a, c) = 2$, a and b are the only two weak Condorcet winners, and c loses to all other alternatives in head-to-head competitions. Such a profile exists due to McGarvey’s theorem [23]. By Lemma 2, $EL(a|P) = EL(b|P) = EL(c|P)$. However, because r satisfies strict Condorcet criterion, $c \notin r(P)$, which is a contradiction. \square

A direct corollary is that any voting rule that satisfies strict Condorcet criterion cannot be the BE of any framework.

Corollary 2. *Copeland₁, maximin, Black's function,³ Dodgson's function, Young's function, Condorcet's function, and Fishburn's function cannot be the Bayesian estimator of any framework.*

5 NEW BAYESIAN ESTIMATORS AS VOTING RULES

The following theorem solves the open question about the satisfaction of Condorcet criterion for $f_{\text{Ma},\varphi}^{\text{Top}}$ [7].

Theorem 5. $f_{\text{Ma},\varphi}^{\text{Top}}$ satisfies the Condorcet criterion if and only if $\frac{\varphi(1-\varphi^{m-1})}{1-\varphi} \leq 1$.

Proof: The “if part”: suppose $\frac{\varphi(1-\varphi^{m-1})}{1-\varphi} \leq 1$. Let P be a profile where a is the Condorcet winner. For any $c, d \in \mathcal{A}$, we let $\mathcal{A}_{-c} = \mathcal{A} \setminus \{c\}$ and $\mathcal{A}_{-cd} = \mathcal{A} \setminus \{c, d\}$. For any $c \in \mathcal{A}$, let \mathcal{L}_c denote the set of all rankings where c is ranked at the top. For any profile P , let $P|_{-a}$ denote its restriction on \mathcal{A}_{-a} .

$$\begin{aligned} 1 - \text{EL}(c|P) &= \sum_{V \in \mathcal{L}_c} \text{Pr}(V|P) \\ &\propto \sum_{V \in \mathcal{L}_c} \text{Pr}(P|V) \propto \sum_{V \in \mathcal{L}_c} \varphi^{\text{Kd}(P,V)} \end{aligned}$$

Fix $b \neq a$. For any ranking $V_{-ab} \in \mathcal{L}(\mathcal{A}_{-ab})$, we let $Q(V_{-ab})$ denote the set of $m-1$ rankings over $\mathcal{A} \setminus \{a\}$ obtained by inserting b to V_{-ab} without changing the relative positions of other alternatives. Let $J(V_{-ab}) \in \mathcal{L}(\mathcal{A}_{-a})$ be the ranking in $Q(V_{-ab})$ with the minimum Kentall-tau distance from $P|_{-a}$. If there are multiple such rankings, let $J(V_{-ab})$ be the one where b is ranked at the highest position.

Let $H : \mathcal{L}_b \rightarrow \mathcal{L}_a$ denote the following mapping. For any $b \succ V_{-b} \in \mathcal{L}_b$ we first look at V_{-ab} and decide the best position to insert b , then put a at the top, where V_{-ab} is obtained from V_{-b} by removing a . Formally, $H(b \succ A_{-b}) = a \succ J(V_{-ab})$.

It follows that for any pair of rankings $V, W \in \mathcal{L}_b$, where the only difference is the position of a , we have $H(V) = H(W)$. Therefore, for any $V \in H(\mathcal{L}_b)$, $H^{-1}(V)$ contains exactly $m-1$ rankings in \mathcal{L}_b that correspond to the $m-1$ positions of a (from the second position to the m -th position—the first position is occupied by b). For each $2 \leq i \leq m$, let $W_i \in H^{-1}(V)$ denote the ranking where a

is ranked at the i -th position.

$$\begin{aligned} &\text{Kd}(P, W_i) \\ &= \text{Kd}(P|_{-a}, (W_i)_{-a}) + \sum_{d \succ_{W_i} a} P(a \succ d) \\ &\quad + \sum_{a \succ_{W_i} d} P(d \succ a) \\ &\geq \text{Kd}(P|_{-a}, J((W_i)_{-ab})) + \sum_{d \neq a} P(d \succ a) + i - 1 \quad (5) \\ &= \text{Kd}(P, V) + i - 1 \end{aligned}$$

Inequality (5) is because a is the Condorcet winner, which means that for any $d \neq a$ we have $\#_P(a \succ d) \geq \#_P(d \succ a) + 1$. Therefore, for each $V \in H(\mathcal{L}_b)$ we have

$$\begin{aligned} \sum_{W \in H^{-1}(V)} \varphi^{\text{Kd}(P,W)} &\leq (\varphi + \dots + \varphi^{m-1}) \varphi^{\text{Kd}(P,V)} \\ &= \frac{\varphi(1-\varphi^{m-1})}{1-\varphi} \varphi^{\text{Kd}(P,V)} \end{aligned}$$

Therefore,

$$\begin{aligned} \sum_{W \in \mathcal{L}_b} \varphi^{\text{Kd}(P,W)} &\leq \frac{\varphi(1-\varphi^{m-1})}{1-\varphi} \sum_{V \in H(\mathcal{L}_b)} \varphi^{\text{Kd}(P,V)} \\ &< \frac{\varphi(1-\varphi^{m-1})}{1-\varphi} \sum_{V \in \mathcal{L}_a} \varphi^{\text{Kd}(P,V)} \leq \sum_{V \in \mathcal{L}_a} \varphi^{\text{Kd}(P,V)} \end{aligned}$$

Therefore, we have $1 - \text{EL}(a|P) > 1 - \text{EL}(b|P)$, which means that $\text{EL}(a|P) < \text{EL}(b|P)$ and a is the unique winner.

The “only if part”: suppose $\frac{\varphi(1-\varphi^{m-1})}{1-\varphi} > 1$. For any odd number $k \in \mathbb{N}$ we consider the a profile P_k whose weighted majority graph is the same as in Figure 1. The existence of P^* is guaranteed by McGarvey's theorem [23]. More precisely, in Figure 1 the weight on the edges from a to all other alternatives is 1; the weight on the edges from b to all other alternatives (except a) is k ; for any $3 \leq i_1 < i_2 < m$, the weight on $a_{i_1} \rightarrow a_{i_2}$ is k .




Figure 1: The WMG of P_k for odd k .

Let $V_a = a \succ b \succ a_3 \succ \dots \succ a_m$ and for each $2 \leq i \leq m$, let V_b^i be the ranking obtained from V_a by moving a to the i -th position. It is not hard to check that $\lim_{k \rightarrow \infty} \frac{\sum_{V \in \mathcal{L}_a} \varphi^{\text{Kd}(P_k, V)}}{\varphi^{\text{Kd}(P_k, V_a)}} = 1$ and

³Definitions of these rules except Copeland and maximin can be found in [18], where it was proved that they satisfy strict Condorcet criterion.

$\lim_{k \rightarrow \infty} \frac{\sum_{V \in \mathcal{L}_b} \varphi^{\text{Kd}(P_k, V)}}{\sum_{i=2}^m \varphi^{\text{Kd}(P_k, V_b^i)}} = 1$. We note that for each $2 \leq i \leq m$, $\text{Kd}(P_k, V_b^i) = \text{Kd}(P_k, V_a) + i - 1$. This means that

$$\begin{aligned} \lim_{k \rightarrow \infty} \frac{\sum_{V \in \mathcal{L}_b} \varphi^{\text{Kd}(P_k, V)}}{\sum_{V \in \mathcal{L}_a} \varphi^{\text{Kd}(P_k, V)}} &= \lim_{k \rightarrow \infty} \frac{\sum_{i=2}^m \varphi^{\text{Kd}(P_k, V_b^i)}}{\varphi^{\text{Kd}(P_k, V_a)}} \\ &= \varphi + \varphi^2 + \dots + \varphi^{m-1} = \frac{\varphi(1 - \varphi^{m-1})}{1 - \varphi} > 1 \end{aligned}$$

Therefore, there exists odd $k \in \mathbb{N}$ such that $\text{EL}(b|P_k) < \text{EL}(a|P_k)$, which means that a cannot be the winner. Because a is the Condorcet winner in P_k , $f_{\text{Ma}, \varphi}^{\text{Top}}$ does not satisfy Condorcet criterion. \square

Theorem 6. For any profile P , $f_{\text{Co}, \varphi}^{\text{Borda}}(P) = \arg \max_{a \in \mathcal{A}} \sum_{c \neq a} \frac{1}{1 + \varphi^{w_P(a, c)}}$.

Proof: For any $W \in \mathcal{B}(\mathcal{A})$ and any pair of alternatives a, b , let $I_W(a \succ b) = 1$ if $a \succ_W b$; otherwise $I_W(a \succ b) = 0$. It follows that $m - 1 - L_{\text{Borda}}(W, a) = \sum_{b \neq a} I_W(a \succ b)$. Let $\mathcal{B}_{a \succ b}$ denote the set of all rankings over \mathcal{A} where $a \succ b$.

$$\begin{aligned} &m - 1 - \text{EL}(a|P) \\ &= \sum_{W \in \mathcal{B}(\mathcal{A})} \Pr(W|P)(m - 1 - L_{\text{Borda}}(W, a)) \\ &= \sum_{W \in \mathcal{B}(\mathcal{A})} \Pr(W|P) \sum_{c \neq a} I_W(a \succ c) \\ &= \sum_{c \neq a} \sum_{W \in \mathcal{B}_{a \succ c}} \Pr(W|P) \end{aligned}$$

Following similar calculations as in [16, 7], we have

$$\begin{aligned} &\sum_{c \neq a} \sum_{W \in \mathcal{B}_{a \succ c}} \Pr(W|P) \\ &\propto \sum_{c \neq a} \varphi^{P(c \succ a)} \prod_{\{b, d\}: \{b, d\} \neq \{a, c\}} (\varphi^{P(b \succ d)} + \varphi^{P(d \succ b)}) \\ &\propto \sum_{c \neq a} \frac{\varphi^{P(c \succ a)}}{\varphi^{P(c \succ a)} + \varphi^{P(a \succ c)}} = \sum_{c \neq a} \frac{1}{1 + \varphi^{w_P(a, c)}} \end{aligned}$$

Therefore, for any pair of alternatives (a, b) , $\text{EL}(a|P) \leq \text{EL}(b|P)$ if and only if $\sum_{c \neq a} \frac{1}{1 + \varphi^{w_P(a, c)}} \geq \sum_{c \neq b} \frac{1}{1 + \varphi^{w_P(b, c)}}$. This proves the theorem. \square

Proposition 1. For all $0 < \varphi < 1$, $f_{\text{Co}, \varphi}^{\text{Borda}}$ satisfies monotonicity.

Proof: For any profile P , any $a \in f_{\text{Co}, \varphi}^{\text{Borda}}(P)$ and any profile P' obtained from P by raising the positions of a without changing relative positions of other alternatives. It is not hard to check that for any $b \neq a$, $w_{P'}(a, b) > w_P(a, b)$, and the weights of edges not involving a do not change. Therefore, for any $b \neq a$, $\sum_{c \neq a} \frac{1}{1 + \varphi^{w_{P'}(a, c)}} > \sum_{c \neq a} \frac{1}{1 + \varphi^{w_P(a, c)}} \geq \sum_{c \neq b} \frac{1}{1 + \varphi^{w_P(b, c)}} > \sum_{c \neq b} \frac{1}{1 + \varphi^{w_{P'}(b, c)}}$. It follows from Theorem 6 that $a \in f_{\text{Co}, \varphi}^{\text{Borda}}(P')$. \square

Theorem 7. $f_{\text{Co}, \varphi}^{\text{Borda}}$ satisfies the Condorcet criterion if and only if $\varphi \leq \frac{1}{m-1}$.

Proof: The ‘‘if’’ part. Let P be any profile where a is the Condorcet winner. This means that for any $c \neq a$, $w_P(a, c) \geq 1$. By Theorem 6 we have $\sum_{c \neq a} \frac{1}{1 + \varphi^{w_P(a, c)}} \geq \frac{m-1}{1 + \varphi}$. For any $b \neq a$, we have $\sum_{c \neq a} \frac{1}{1 + \varphi^{w_P(a, c)}} < \frac{1}{1 + \varphi^{-1}} + m - 2$. When $\varphi \leq \frac{1}{m-1}$, we have $\frac{1}{1 + \varphi^{-1}} + m - 2 \leq \frac{m-1}{1 + \varphi}$. Therefore, a is the unique winner.

The ‘‘only if’’ part is proved by considering the profile P_k whose WMG is in Figure 1 and let $k \rightarrow \infty$. \square

Theorem 8. As $\varphi \rightarrow 0$, $f_{\text{Co}, \varphi}^{\text{Borda}}$ converges to a refinement of Copeland_{0.5}. As $\varphi \rightarrow 1$, $f_{\text{Co}, \varphi}^{\text{Borda}}$ converges to a refinement of Borda.

Proof: For any profile P any pair of alternatives a, b we have

$$\lim_{\varphi \rightarrow 0} \frac{1}{1 + \varphi^{w_P(a, b)}} = \begin{cases} 1 & \text{if } w_P(a, b) > 0 \\ 0.5 & \text{if } w_P(a, b) = 0 \\ 0 & \text{otherwise} \end{cases}$$

Therefore, for any alternative a , $\lim_{\varphi \rightarrow 0} \sum_{c \neq a} \frac{1}{1 + \varphi^{w_P(a, c)}}$ is its Copeland_{0.5} score, which means that the winners must also be winners under Copeland_{0.5}.

For any $k > 0$, when $\epsilon \rightarrow 0$, we have $\frac{1}{1 + (1 - \epsilon)^k} = \frac{1}{2}(1 + \frac{k\epsilon}{2} + o(\epsilon))$ and $\frac{1}{1 + (1 - \epsilon)^{-k}} = \frac{1}{2}(1 - \frac{k\epsilon}{2} + o(\epsilon))$. Therefore, for any alternative a , $\sum_{c \neq a} \frac{1}{1 + \varphi^{w_P(a, c)}} = \frac{m-1}{2} + \frac{1}{4}(\sum_{c \neq a} w_P(a, c))(1 - \varphi) + o(1 - \varphi)$. We note that $\sum_{c \neq a} w_P(a, c)$ equals to twice the Borda score of a in P minus $n(m - 1)$. Therefore, as $\varphi \rightarrow 1$, all $f_{\text{Co}, \varphi}^{\text{Borda}}$ winners must be Borda winners. \square

We propose a new class of ranking model and framework as follows.

Definition 5. For any $0 < \varphi < 1$ we define $\mathcal{M}_{\text{Pair}, \varphi}$ as follows. The parameter space $\Theta = \{\theta_{bc} : b \neq c \in \mathcal{A}\}$. For any $V \in \mathcal{L}(\mathcal{A})$ we let $\pi_{\theta_{bc}}(V) \propto \begin{cases} 1 & \text{if } b \succ_V c \\ \varphi & \text{otherwise} \end{cases}$.

$$\text{Let } L_1(\theta_{bc}, a) = \begin{cases} 1 & \text{if } a = c \\ 0 & \text{otherwise} \end{cases} \text{ and}$$

$$L_2(\theta_{bc}, a) = \begin{cases} 0 & \text{if } a = b \\ 1 & \text{otherwise} \end{cases}.$$

$$\text{Let } \mathcal{F}_{\text{Pair}, \varphi}^1 = (\mathcal{M}_{\text{Pair}, \varphi}, L_1) \text{ and } \mathcal{F}_{\text{Pair}, \varphi}^2 = (\mathcal{M}_{\text{Pair}, \varphi}, L_2).$$

That is, the parameters in $\mathcal{M}_{\text{Pair}, \varphi}$ correspond to pairwise comparisons between alternatives. A parameter θ_{bc} can be interpreted as ‘‘ $b \succ c$ is the strongest pairwise comparison’’. The first loss function states that the loss of a is 1 if and only if a is the less preferred alternative in the parameter. The second loss function states that the loss of a is 0 if and only if a is the preferred alternative in the parameter.

$\mathcal{M}_{\text{Pair}, \varphi}$ might be of independent interest. In this paper we

focus on the satisfaction of axioms for the two Bayesian estimators and leave further exploration of the model for future work. We note that given φ , the normalization factor for all θ_{bc} are the same.

Theorem 9. *The Bayesian estimator $f_{Pair,\varphi}^1$ of $\mathcal{F}_{Pair,\varphi}^1$ is $\arg \min_{a \in \Theta} \sum_{b \neq a} \varphi^{w_P(a,b)/2}$. The Bayesian estimator $f_{Pair,\varphi}^2$ of $\mathcal{F}_{Pair,\varphi}^2$ is $\arg \max_{a \in \Theta} \sum_{b \neq a} \varphi^{w_P(b,a)/2}$.*

It is easy to check that both $\mathcal{F}_{Pair,\varphi}^1$ and $\mathcal{F}_{Pair,\varphi}^2$ satisfy neutrality and parameter connectivity. Therefore, their Bayesian estimators satisfy neutrality and minimaxity.

Corollary 3. *$f_{Pair,\varphi}^1$ and $f_{Pair,\varphi}^2$ satisfy neutrality and minimaxity (w.r.t. to $\mathcal{F}_{Pair,\varphi}^1$ and $\mathcal{F}_{Pair,\varphi}^2$, respectively).*

Proposition 2. *$f_{Pair,\varphi}^1$ and $f_{Pair,\varphi}^2$ satisfy monotonicity.*

Proof: The proof is similar to the proof of Theorem 1. We note that raising the position of a will increase the weight on some edges $a \rightarrow b$. Weights on other edges do not change. Monotonicity of both rules can be verified by applying Theorem 9. \square

Theorem 10. *$f_{Pair,\varphi}^1$ satisfies the Condorcet criterion if and only if $\varphi \leq \frac{1}{m-1}$. For all $0 < \varphi < 1$, $f_{Pair,\varphi}^2$ does not satisfy Condorcet criterion.*

Proof: The “if” part for $f_{Pair,\varphi}^1$ follows after Theorem 11 because when $\varphi < \frac{1}{m-1}$, $f_{Pair,\varphi}^1$ is a refinement of maximin and any refinement of maximin satisfies the Condorcet criterion.

The “only if” part for $f_{Pair,\varphi}^1$ and the non-satisfaction for $f_{Pair,\varphi}^2$ are proved by considering the profile P_k whose weighted majority graph is in Figure 1 and let $k \rightarrow \infty$. \square

For any profile P , the *maximax* rule chooses all alternatives with the maximum weight on at least one outgoing edge in the weighted majority graph. That is, the rule is $\arg \max_a \max_b w_P(a,b)$.

Theorem 11. *For any $\varphi \leq \frac{1}{m-1}$, $f_{Pair,\varphi}^1$ is a refinement of maximin, and $f_{Pair,\varphi}^2$ is a refinement of maximax. As $\varphi \rightarrow 1$, both rules converge to refinements of Borda.*

Proof: By Theorem 9, for any $\varphi \leq \frac{1}{m-1}$, for any alternative a , $\sum_{b \neq a} \varphi^{w_P(a,b)/2}$ is mainly determined by $\min_{b \neq a} w_P(a,b)/2$, which is half of a ’s min-score. It follows that all winners under $f_{Pair,\varphi}^1$ must be maximin winners. Similarly, $\sum_{b \neq a} \varphi^{w_P(b,a)/2}$ is determined by $\min_{b \neq a} w_P(b,a)/2$, which corresponds to $\max_{b \neq a} w_P(a,b)/2$. It follows that the winner under $f_{Pair,\varphi}^2$ must be maximax winners.

For any $\epsilon > 0$, we have $\sum_{b \neq a} (1 - \epsilon)^{w_P(a,b)/2} = m - 1 - \sum_{b \neq a} \frac{w_P(a,b)}{2} \epsilon + o(\epsilon)$. Similar to the proof of Theorem 8, the minimizers of this function as $\epsilon \rightarrow 0$, which is $f_{Pair,\varphi}^1$

as $\varphi \rightarrow 1$, must be Borda winners. The proof for $f_{Pair,\varphi}^2$ is similar. \square

6 SUMMARY AND FUTURE WORK

We characterized neutrality and proved an impossibility theorem about strict Condorcet criterion for Bayesian estimators. We also proposed new frameworks to obtain new BEs and showed that some of them satisfy many desirable axioms and can be computed in polynomial time.

There are many directions for future work. Can we answer Q2 in the Introduction for other desirable axioms such as homogeneity? How about axioms for other types preferences such as ratings? Are there any BEs that are refinements of other commonly studied rules, especially Copeland $_\alpha$ for $\alpha \notin \{0.5, 1\}$, STV, and ranked pairs? What are other natural frameworks and which axioms do their BEs satisfy?

ACKNOWLEDGMENTS

This work is supported by the National Science Foundation under grant IIS-1453542 and a Simons-Berkeley research fellowship. We thank Felix Brandt, Jean-Francois Laslie, Bill Zwicker, participants of the Economics and Computation Program at Simons Institute, and anonymous reviewers for helpful comments and suggestions.

References

- [1] IMDb Votes/Ratings Top Frequently Asked Questions, 2016. URL http://www.imdb.com/help/show_leaf?votestopfaq.
- [2] Practical example of bayes estimators, 2016. URL https://en.wikipedia.org/wiki/Bayes_estimator#Practical_example_of_Bayes_estimators.
- [3] Alon Altman and Moshe Tennenholtz. An axiomatic approach to personalized ranking systems. *Journal of the ACM*, 57(4), 2010. Article 26.
- [4] Kenneth Arrow. *Social choice and individual values*. New Haven: Cowles Foundation, 2nd edition, 1963. 1st edition 1951.
- [5] Hossein Azari Soufiani, David C. Parkes, and Lirong Xia. Random utility theory for social choice. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, pages 126–134, Lake Tahoe, NV, USA, 2012.
- [6] Hossein Azari Soufiani, David C. Parkes, and Lirong Xia. Preference Elicitation For General Random Utility Models. In *Proceedings of Uncertainty in Artificial Intelligence (UAI)*, Bellevue, Washington, USA, 2013.
- [7] Hossein Azari Soufiani, David C. Parkes, and Lirong Xia. Statistical decision theory approaches to social choice. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, Montreal, Quebec, Canada, 2014.
- [8] James O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer, 2nd edition, 1985.

- [9] Christian Borgs, Jennifer Chayes, Adrian Marple, and Shang-Hua Teng. An Axiomatic Approach to Community Detection. In *Proceedings of ITCS*, 2016.
- [10] Mark Braverman and Elchanan Mossel. Noisy Sorting Without Resampling. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 268–276, Philadelphia, PA, USA, 2008.
- [11] Ioannis Caragiannis, Ariel Procaccia, and Nisarg Shah. When do noisy votes reveal the truth? In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, Philadelphia, PA, 2013.
- [12] Christine Choirat and Raffaello Seri. Estimation in discrete parameter models. *Statistical Science*, 27(2):278–293, 2012.
- [13] Marquis de Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Paris: L'Imprimerie Royale, 1785.
- [14] Vincent Conitzer and Tuomas Sandholm. Common voting rules as maximum likelihood estimators. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 145–152, Edinburgh, UK, 2005.
- [15] Vincent Conitzer, Matthew Rognlie, and Lirong Xia. Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*, pages 109–115, Pasadena, CA, USA, 2009.
- [16] Edith Elkind and Nisarg Shah. Electing the Most Probable Without Eliminating the Irrational: Voting Over Intransitive Domains. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, pages 182–191, 2014.
- [17] Edith Elkind, Piotr Faliszewski, and Arkadii Slinko. Good rationalizations of voting rules. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 774–779, Atlanta, GA, USA, 2010.
- [18] Peter C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.
- [19] David Hughes, Kevin Hwang, and Lirong Xia. Computing Optimal Bayesian Decisions for Rank Aggregation via MCMC Sampling. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2015.
- [20] Tyler Lu and Craig Boutilier. Learning mallows models with pairwise preferences. In *Proceedings of the Twenty-Eighth International Conference on Machine Learning (ICML 2011)*, pages 145–152, Bellevue, WA, USA, 2011.
- [21] Colin L. Mallows. Non-null ranking model. *Biometrika*, 44(1/2):114–130, 1957.
- [22] Andrew Mao, Ariel D. Procaccia, and Yiling Chen. Better human computation through principled voting. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Bellevue, WA, USA, 2013.
- [23] David C. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21(4):608–610, 1953.
- [24] Sahand Negahban, Sewoong Oh, and Devavrat Shah. Iterative ranking from pair-wise comparisons. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, pages 2483–2491, Lake Tahoe, NV, USA, 2012.
- [25] David M Pennock, Eric Horvitz, and C. Lee Giles. Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 729–734, Austin, TX, USA, 2000.
- [26] Marcus Pivato. Voting rules as statistical estimators. *Social Choice and Welfare*, 40(2):581–630, 2013.
- [27] Ariel D. Procaccia, Sashank J. Reddi, and Nisarg Shah. A maximum likelihood approach for selecting sets of alternatives. In *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence*, 2012.
- [28] Lirong Xia. Statistical Properties of Social Choice Mechanisms. In *Proceedings of the Fifth International Workshop on Computational Social Choice*, Pittsburgh, Pennsylvania, USA, 2014.
- [29] Lirong Xia and Vincent Conitzer. A maximum likelihood approach towards aggregating partial orders. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, pages 446–451, Barcelona, Catalonia, Spain, 2011.
- [30] H. Peyton Young. Condorcet's theory of voting. *American Political Science Review*, 82:1231–1244, 1988.