

Ranked Voting on Social Networks*

Ariel D. Procaccia
Carnegie Mellon University
arielpro@cs.cmu.edu

Nisarg Shah
Carnegie Mellon University
nkshah@cs.cmu.edu

Eric Sodomka
Facebook
sodomka@fb.com

Abstract

Classic social choice theory assumes that votes are independent (but possibly conditioned on an underlying objective ground truth). This assumption is unrealistic in settings where the voters are connected via an underlying social network structure, as social interactions lead to correlated votes. We establish a general framework — based on random utility theory — for ranked voting on a social network with arbitrarily many alternatives (in contrast to previous work, which is restricted to two alternatives). We identify a family of voting rules which, without knowledge of the social network structure, are guaranteed to recover the ground truth with high probability in large networks, with respect to a wide range of models of correlation among input votes.

1 Introduction

Social choice theory (and *computational* social choice, in particular) typically views votes — represented as *rankings* of a set of alternatives — as manifestations of subjective *preferences*. But an alternative viewpoint has been gaining steam in recent years: The votes are seen as objective *noisy estimates* of the true quality of alternatives. This viewpoint actually dates back to the early beginnings of social choice theory in the 18th Century; its newfound popularity is due in part to potential applications in human computation [Mao *et al.*, 2013] and multiagent systems [Jiang *et al.*, 2014].

If we indeed assume that some alternatives are truly better than others, and that voters are communicating possibly inaccurate information about this ground truth, then the goal of a *voting rule* — which aggregates the reported votes into a single ranking — is indisputable: to uncover the truth. Caragiannis *et al.* [2013] formalize this abstract goal by asking for voting rules that are *accurate in the limit*:¹ The rule should return a ranking that reflects the ground truth with high probability when the electorate is large, i.e., with probability that goes to 1 as the number of submitted votes goes to infinity.

*This work was partially supported by NSF grants CCF-1215883 and IIS-1350598, and by a Sloan Research Fellowship.

¹This is known as *consistency* in statistics, but this term has long been used in social choice theory to refer to a different property.

They pinpoint families of voting rules that exhibit robustness: they are accurate in the limit with respect to a wide range of *noise models*, which govern the way noisy votes are generated, given the ground truth [Caragiannis *et al.*, 2013; 2014].

While these results are promising, they rely on a crucial modeling assumption: votes are independent. This assumption is clearly satisfied in some settings — when votes are submitted by computer Go programs [Jiang *et al.*, 2014], say. However, in many other settings — especially when the voters are people — votes are likely to be correlated through social interactions. We refer to the structure of these interactions as a *social network*, interpreted in the broadest possible sense: any form of interaction qualifies for an edge. From this broad viewpoint, the structure of the social network cannot be known, and, hence, votes are correlated in an unpredictable way. Inspired by the robustness approach of Caragiannis *et al.* [2013; 2014], our goal is to

... model the generation of noisy rankings on a social network given a ground truth, and identify voting rules that are accurate in the limit with respect to any network structure and (almost) any choice of model parameters.

Our Model and Results. Our starting point is the recently-introduced *independent conversations model* [Conitzer, 2013]. In this model, there are only two alternatives: one is “correct” (stronger) and one is “incorrect” (weaker). Each edge of the social network is an independent conversation between two voters, whose result (which is independent of the results on other edges — hence the name of the model) is the correct alternative with probability $p > 1/2$, and the incorrect alternative with probability $1 - p$. Then, each voter aggregates the results on the incident edges using the majority rule, and submits the resulting alternative (i.e., the final vote) to the voting rule. Note that if two voters are neighbors in the network, their votes are *not* independent. The voting rule only observes the final votes submitted by the voters (and not the results of conversations on the edges), and must aggregate these votes to find the correct alternative.

Conitzer acknowledges that his goal is to “give a simple model that helps to illustrate which phenomena we are likely to encounter as we move to more complex models” [Conitzer, 2013, p. 1483]. We are indeed interested in a more com-

plex model, which supports multiple alternatives and rests on richer probabilistic foundations. In our model, we assume that each alternative a has a true quality μ_a . The result of an independent conversation on an edge is a noisy quality estimate for each alternative a sampled from a Gaussian distribution with mean μ_a . Each voter assigns a weight to each incident edge, and computes an aggregate quality estimate for each alternative a by taking a weighted average of the noisy quality estimates of a on the incident edges. The voter submits a ranking of the alternatives by their aggregate quality estimates.

We analyze the performance of two disjoint families of voting rules — PM-c rules and PD-c rules — proposed by Caragiannis et al. [2013], which together include most popular voting rules, along with the performance of another voting rule — the modal ranking rule — which has been shown to exhibit extreme robustness properties under independent noisy votes [Caragiannis et al., 2014]. Under a mild condition on the weights placed by the voters on their incident edges, we show that all PM-c rules, an important subset of PD-c rules, and the modal ranking rule are accurate in the limit when all the Gaussian distributions have equal variance (Section 5). However, when the Gaussians can have unequal variance, many PD-c rules and the modal ranking rule are no longer accurate in the limit, whereas all PM-c rules stay accurate in the limit (Section 6). Therefore, PM-c rules exhibit qualitatively more robustness than PD-c rules and the modal ranking rule.

2 Related Work

Our paper is closely related to two papers by Conitzer [2012; 2013]. The independent conversations model of the latter paper was discussed above. Importantly, the challenge Conitzer [2013] addresses is quite different from ours: he is interested in finding the *maximum likelihood estimator (MLE)* for the ground truth, i.e., he wants to know which of the two alternatives is more likely to be correct, given the observed (binary) votes. The answer strongly depends on social network structure, and his main result is that, in fact, the problem is #P-hard. In an earlier, brief note, Conitzer [2012] is also interested in the maximum likelihood approach to noisy voting on a social network. While the model he introduces also extends to the case of more than two alternatives, the assumptions of the model are such that the (known) network structure is essentially irrelevant, that is, the maximum likelihood estimator is invariant to network structure.

While the above papers are, to our knowledge, the only papers that deal with the MLE approach to voting on a social network, there is a substantial body of work on the MLE approach to voting more generally [Young, 1988; Conitzer and Sandholm, 2005; Conitzer et al., 2009; Elkind et al., 2010; Xia et al., 2010; Xia and Conitzer, 2011; Lu and Boutilier, 2011; Procaccia et al., 2012; Azari Soufiani et al., 2012; 2013; 2014]. However, all of these papers assume that votes are drawn i.i.d. from a noise model.

A bit further afield, there is a large body of work that studies the diffusion of opinions, votes, technologies, or products (but not ranked estimates) in a social network. An espe-

cially pertinent example is the work of Mossel et al. [2014], where at each time step voters adopt the most popular opinion among their neighbors, and at some point opinions are aggregated via the plurality rule. Other popular diffusion models include the *independent cascade model*, the *linear threshold model*, and the DeGroot model [DeGroot, 1974]; see the survey by Kleinberg [2007] for a fascinating overview.

3 Preliminaries

Let A denote a set of alternatives, where $|A| = m$. Let $\mathcal{L}(A)$ denote the set of rankings (linear orders) over A . A vote σ is a ranking in $\mathcal{L}(A)$, and a profile π is a collection of votes. We use n to denote the number of votes. Let $a \succ_{\sigma} b$ denote that alternative a is preferred to alternative b in ranking σ , and let $\sigma(a)$ denote the rank of a in σ . A voting rule is formally a *social welfare function (SWF)* that maps every profile to a ranking. Caragiannis et al. [2013] define two general families of voting rules that capture most prominent voting rules.

- *PM-c rules*: For a profile π , the *pairwise-majority (PM) graph* is a directed graph whose vertices are the alternatives, and there exists an edge from $a \in A$ to $b \in A$ if a strict majority of the voters prefer a to b . A voting rule f is called *pairwise-majority consistent (PM-c)* if for every profile π with a complete acyclic PM graph whose vertices are ordered according to $\sigma \in \mathcal{L}(A)$, we have $f(\pi) = \sigma$. Prominent voting rules such as the Kemeny rule, the Slater rule, the ranked pairs method, Copeland’s method, and Schulze’s method are PM-c rules.
- *PD-c rules*: In a profile π , alternative a is said to *position-dominate* alternative b if for every $k \in \{1, \dots, m-1\}$, (strictly) more voters rank a in the first k positions than b . The *position-dominance (PD) graph* is a directed graph whose vertices are the alternatives, and there exists an edge from a to b if a position-dominates b . A voting rule f is called *position-dominance consistent (PD-c)* if for every profile π with a complete acyclic PD graph whose vertices are ordered according to $\sigma \in \mathcal{L}(A)$, we have $f(\pi) = \sigma$.

Positional scoring rules are a popular family of voting rules that are PD-c. A positional scoring rule is given by a score vector $(\alpha_1, \dots, \alpha_m)$ where $\alpha_i \geq \alpha_{i+1}$ for $1 \leq i < m$. Given a profile π , for each vote $\sigma \in \pi$ and each position $k \in \{1, \dots, m\}$, the rule awards α_k points to the alternative placed at position k in vote σ , and returns a ranking by sorting the alternatives according to their total scores. Plurality, Borda count, and veto are examples of well-known positional scoring rules. A positional scoring rule is *strict* if the score vector satisfies $\alpha_i > \alpha_{i+1}$ for $1 \leq i < m$ (Borda count is strict). In addition to positional scoring rules, Bucklin’s rule is also PD-c.

Caragiannis et al. [2014] introduce another voting rule — the *modal ranking rule* — which is neither a PM-c rule nor a PD-c rule. Given a profile, the modal ranking rule simply returns the ranking that appears the largest number of times in the profile. Our analysis focuses on the performance of PM-c rules, PD-c rules, and the modal ranking rule.

4 Our Model

As discussed in Section 1, we extend the *independent conversations model* [Conitzer, 2013], which is restricted to two alternatives. In our extended model, there is a set of (arbitrarily many) alternatives A . We assume that the voters are connected via an underlying social network structure, represented as an undirected graph $G = (V, E)$ (here, V is the set of voters). We use the notation $e \downarrow v$ to denote that edge e is incident on voter v . Let $E(v) = \{e \in E \mid e \downarrow v\}$ denote the set of edges incident on v , and let $d_v = |E(v)|$ denote the degree of v in the network G . As we explain in Section 1, the social network structure may be unknown to us. Our model has four key components.

- *Ground truth.* We assume that each alternative $a \in A$ has a “true quality” denoted by μ_a . The ground truth ranking of the alternatives σ^* ranks the alternatives by their true qualities. We assume that for some constant $\Delta > 0$, we have $|\mu_a - \mu_b| \leq \Delta$ for all distinct $a, b \in A$.
- *Quality estimates.* When voters v and v' share an edge, they have an independent discussion. We represent the result of this discussion as a quality estimate for each alternative. Specifically, we associate a random variable $X_{e,a}$ to each edge e for the quality estimate of each alternative a . Crucially, we assume that all $\{X_{e,a}\}_{e \in E, a \in A}$ are mutually independent.
- *Aggregation rules.* We assume that voter v uses an aggregation rule $g_v : \mathbb{R}^{d_v} \rightarrow \mathbb{R}$ to derive an aggregate quality estimate $Y_{v,a} = g(\{X_{e,a}\}_{e \in E(v)})$ for each alternative $a \in A$. In the tradition of random utility theory, his submitted vote σ_v is a ranking of the alternatives by their aggregate quality estimates.
- *Voting rule.* The only information we observe is the set of rankings (votes) submitted by the voters. In particular, we are unaware of the quality estimates sampled on the edges (i.e., values of $X_{e,a}$), or the aggregate quality estimates derived by the voters (i.e., values of $Y_{v,a}$). Moreover, we assume that the distributions of the independent conversations on the edges, the aggregation rules used by the voters, the identities of the voters, and their social network structure are also unknown to us. We use an *anonymous* voting rule $f : \mathcal{L}(A)^n \rightarrow \mathcal{L}(A)$ to aggregate the submitted ranked votes into a final ranking of the alternatives. Our goal is to be accurate in the limit, i.e., produce the ground truth ranking σ^* with probability 1 as the number of voters n goes to infinity.

In the next two sections, we instantiate this general model by considering specific distributions of the quality estimates on the edges ($X_{e,a}$) and specific choices of the aggregation rules used by the voters.

5 Equal Variance

Let us focus on the following model of independent conversations and aggregation rules.

Quality estimates. Our choice is inspired by the classic *Thurstone-Mosteller model* [Thurstone, 1927; Mosteller, 1951], in which a quality estimate is derived by taking a

sample from a Gaussian distribution centered around the true quality. This model is member of the more general class of *random utility models* (see [Azari Soufiani *et al.*, 2012] for their use in social choice) in which the distribution need not be Gaussian. In our setting, for each edge $e \in E$ and alternative $a \in A$ we assume $X_{e,a} \sim \mathcal{N}(\mu_a, \nu^2)$, which is a Gaussian distribution with mean μ_a , variance ν^2 , and probability density function

$$p(x) = \frac{1}{\sqrt{2\pi\nu^2}} e^{-\frac{(x-\mu_a)^2}{2\nu^2}}.$$

Crucially, we assume that the variance of all the Gaussians is equal, i.e., the noise present in the quality estimates is random noise that is not dependent on the voters or on the alternatives. This is not a weak assumption; we relax it in Section 6.

Aggregation rules. We assume that voters aggregate the quality estimates of the alternatives on their incident edges by computing a weighted mean. Specifically, assume that each voter v places a weight $w_v(e) \in \mathbb{R}_{\geq 0}$ on each incident edge $e = (v, v') \in E(v)$, which represents how much the voter weights or believes in the conversation with voter v' . Without loss of generality, let the weights be normalized such that $\sum_{e \in E(v)} w_v(e) = 1$ for all $v \in V$. Then, the aggregate quality estimate derived by voter v for alternative a is given by $Y_{v,a} = \sum_{e \in E(v)} w_v(e) X_{e,a}$.

We aim to find voting rules that provide accuracy in the limit for any social network structure G , and for a wide range of choices of the unknown parameters: the true qualities of the alternatives $\{\mu_a\}_{a \in A}$, the variance of the Gaussian distributions ν^2 , and the weights assigned by voters to their incident edges $\{w_v(e)\}_{v \in V, e \in E(v)}$. The main difficulty is that the votes of two voters may be correlated when they share an edge in the social network, but the network is unknown to the voting rule. To this end, we first prove a result that shows that under certain conditions, the correlation has negligible effect on the final outcome. We later leverage this result to identify anonymous voting rules that are accurate in the limit.

Lemma 1. *Let $Z_v^1, Z_v^2 \in [-\xi, \xi]$ be two bounded random variables associated with each voter $v \in V$, where $\xi > 0$ is a constant. For $i, j \in \{1, 2\}$ and $v, v' \in V$, assume Z_v^i and $Z_{v'}^j$ are independent unless $v = v'$ or $(v, v') \in E$. If there exist positive constants C, γ, δ , and ϵ such that for all $v \in V$,*

1. $\mathbb{E}[Z_v^1] - \mathbb{E}[Z_v^2] \geq \gamma$, and
2. $\Pr[Z_v^1 \leq Z_v^2 + \delta] \leq C/(d_v)^{1+\epsilon}$,

then $\lim_{n \rightarrow \infty} \Pr[\sum_{v \in V} Z_v^1 > \sum_{v \in V} Z_v^2] = 1$.

Before we dive into the proof, note that if the random variables were independent, condition 1 and Hoeffding’s inequality would have implied the required result. For correlated variables, the intuition is as follows. If d_v is small, then Z_v^1 and Z_v^2 are correlated with only a few other random variables. If d_v is large, then $Z_v^1 > Z_v^2$ holds with high probability anyway. As we later see in Theorem 1, this is because voters with large degrees produce accurate votes by assimilating a large amount of independent information from incident edges.

Proof. Partition the set of voters V into two subsets:

$$V_1 = \left\{ v \in V \mid d_v \leq n^{\frac{1+0.5\epsilon}{1+\epsilon}} \right\} \text{ and } V_2 = V \setminus V_1.$$

Define $Z_{V_i}^j = \sum_{v \in V_i} Z_v^j$ and $Z_V^j = Z_{V_1}^j + Z_{V_2}^j$ for $i, j \in \{1, 2\}$. We wish to prove that $Z_V^1 > Z_V^2$ holds with high probability. We focus on the relations between $Z_{V_1}^1$ and $Z_{V_1}^2$, and between $Z_{V_2}^1$ and $Z_{V_2}^2$ separately, and later combine the two results to prove the required result.

Voters in V_1 . Observe that $\mathbb{E}[Z_{V_1}^1 - Z_{V_1}^2] = \sum_{v \in V_1} \mathbb{E}[Z_v^1] - \mathbb{E}[Z_v^2] \geq |V_1| \cdot \gamma$. As previously mentioned, we cannot simply use Hoeffding's inequality because the indicator random variables are correlated. We instead use Chebyshev's inequality.

$$\begin{aligned} \Pr[Z_{V_1}^1 \leq Z_{V_1}^2] &\leq \Pr[|(Z_{V_1}^1 - Z_{V_1}^2) - \mathbb{E}[Z_{V_1}^1 - Z_{V_1}^2]| \geq |V_1| \cdot \gamma] \\ &\leq \frac{\text{Var}(Z_{V_1}^1 - Z_{V_1}^2)}{|V_1|^2 \cdot \gamma^2}. \end{aligned} \quad (1)$$

Here, $\text{Var}(\cdot)$ denotes the variance of a random variable. To derive an upper bound on $\text{Var}(Z_{V_1}^1 - Z_{V_1}^2)$, we use the fact that for $i, j \in \{1, 2\}$ and $v, v' \in V_1$, indicator random variables Z_v^i and $Z_{v'}^j$ are only correlated if $v = v'$ or v and v' share an edge (i.e., $(v, v') \in E$). Thus, the random variables corresponding to voter v can be correlated with the random variables corresponding to at most $1 + d_v$ voters. Further, when they are correlated, their covariance satisfies

$\text{Cov}(Z_v^i, Z_{v'}^j) \leq \sqrt{\text{Var}(Z_v^i) \cdot \text{Var}(Z_{v'}^j)} \leq \xi^2$, where the last transition holds because the variance of a $[-\xi, \xi]$ -bounded random variable is at most ξ^2 due to Popoviciu's inequality. Hence,

$$\begin{aligned} \text{Var}(Z_{V_1}^1 - Z_{V_1}^2) &= \sum_{i, j \in \{1, 2\}} \sum_{v \in V_1} \sum_{\substack{v' \in V_1: \\ [v'=v] \vee [(v, v') \in E]}} \text{Cov}(Z_v^i, Z_{v'}^j) \\ &\leq \xi^2 \cdot \sum_{v \in V_1} 4 \cdot (1 + d_v) \leq 4 \cdot \xi^2 \cdot |V_1| \cdot \left(1 + n^{\frac{1+0.5\epsilon}{1+\epsilon}}\right), \end{aligned}$$

where the last transition holds because $d_v \leq n^{\frac{1+0.5\epsilon}{1+\epsilon}}$ for all $v \in V_1$. Substituting this into Equation (1),

$$\Pr[Z_{V_1}^1 \leq Z_{V_1}^2] \leq 4 \cdot \xi^2 \cdot \frac{1 + n^{\frac{1+0.5\epsilon}{1+\epsilon}}}{|V_1| \cdot \gamma^2} \quad (2)$$

Note that this probability could be high when $|V_1|$ is small.

Voters in V_2 . Fix $v \in V_2$. Then, $d_v \geq n^{\frac{1+0.5\epsilon}{1+\epsilon}}$ by the definition of V_2 . Hence,

$$\Pr[Z_v^1 \leq Z_v^2 + \delta] \leq \frac{C}{(d_v)^{1+\epsilon}} \leq \frac{C}{n^{1+0.5\epsilon}}, \quad (3)$$

where the first transition follows from the second condition assumed in the lemma. Now,

$$\begin{aligned} \Pr[Z_{V_2}^1 \leq Z_{V_2}^2 + |V_2| \cdot \delta] &\leq \sum_{v \in V_2} \Pr[Z_v^1 \leq Z_v^2 + \delta] \\ &\leq \frac{C \cdot |V_2|}{n^{1+0.5\epsilon}} \leq \frac{C}{n^{0.5\epsilon}}, \end{aligned} \quad (4)$$

where the first transition follows from the Pigeonhole principle, the second transition follows from Equation (3), and the last transition holds because $|V_2| \leq n$. Note that this probability must go to 0 as $n \rightarrow \infty$, unlike $\Pr[Z_{V_1}^1 \leq Z_{V_1}^2]$.

We now consider two cases to combine our results.

1. Suppose $|V_2| \geq n \cdot 2\xi / (2\xi + \delta)$. Then, $|V_1| \leq n \cdot \delta / (2\xi + \delta)$. Observe that we always have $Z_{V_1}^1 - Z_{V_1}^2 \geq -|V_1| \cdot 2\xi$. If it holds that $Z_{V_2}^1 - Z_{V_2}^2 > |V_2| \cdot \delta$, then $Z_V^1 > Z_V^2$ follows by adding the two inequalities and substituting the bounds of $|V_1|$ and $|V_2|$. Hence, $\Pr[Z_V^1 \leq Z_V^2] \leq \Pr[Z_{V_2}^1 \leq Z_{V_2}^2 + |V_2| \cdot \delta]$, which goes to 0 as n goes to infinity due to Equation (4).
2. Suppose $|V_2| \leq n \cdot 2\xi / (2\xi + \delta)$. Then, $|V_1| \geq n \cdot \delta / (2\xi + \delta)$. Substituting this into Equation (2), we see that $\Pr[Z_{V_1}^1 \leq Z_{V_1}^2]$ approaches 0 as n goes to infinity. Equation (4) already shows that $\Pr[Z_{V_2}^1 \leq Z_{V_2}^2] \leq \Pr[Z_{V_2}^1 \leq Z_{V_2}^2 + |V_2| \cdot \delta]$ approaches 0 as n goes to infinity. Hence, $\Pr[Z_V^1 \leq Z_V^2] \leq \Pr[Z_{V_1}^1 \leq Z_{V_1}^2] + \Pr[Z_{V_2}^1 \leq Z_{V_2}^2]$ goes to 0 as n goes to infinity.

Thus, in both cases we have the desired result. ■

We now use Lemma 1 to derive our main result.

Theorem 1. *If there exists a universal constant $D \in \mathbb{N}$ such that $\sum_{e \in E(v)} [w_v(e)]^2 \leq \Delta^2 / (8\nu^2 \ln d_v)$ for all voters v with degree $d_v \geq D$, then all PM-c rules, the modal ranking rule, and all strict positional scoring rules are accurate in the limit irrespective of the choices of the unknown parameters: the social network structure G , the true qualities $\{\mu_a\}_{a \in A}$, the variance ν^2 , and the weights $\{w_v(e)\}_{v \in V, e \in E(v)}$.*

Before we prove the result, we remark that the bound on $\sum_{e \in E(v)} [w_v(e)]^2$ is a mild restriction. In our setting with normalized weights ($\sum_{e \in E(v)} w_v(e) = 1$), the unweighted mean has $\sum_{e \in E(v)} [w_v(e)]^2 = 1/d_v$ which is much smaller than our required bound. More generally, the condition is satisfied if no voter v places an excessive weight — specifically, a weight greater than $\Delta / (4\nu\sqrt{d_v \ln d_v})$ — on any single incident edge. Below, we present the proof for PM-c rules and the modal ranking rule; the proof for strict positional scoring rules appears in the full version of the paper.²

Proof of Theorem 1 (partial). Let us begin with PM-c rules.

PM-c Rules. Recall that PM-c rules are guaranteed to return the ground truth ranking σ^* if the pairwise majority graph is consistent with σ^* . We wish to use Lemma 1 to show that for every pair of alternatives $a, b \in A$ such that $a \succ_{\sigma^*} b$, there is an edge from a to b in the pairwise majority graph of the profile consisting of the votes submitted by the voters with probability 1 as n goes to infinity. Applying the union bound over all pairs of alternatives implies that the entire pairwise majority graph would be consistent with σ^* with probability 1 as n goes to infinity.

Now, for voter $v \in V$ and alternative $a \in A$, the aggregate quality estimate $Y_{v,a} = \sum_{e \in E(v)} w_v(e) X_{e,a}$ follows the distribution $\mathcal{N}(\mu_a, \nu^2 \sum_{e \in E(v)} [w_v(e)]^2)$ because each $X_{e,a} \sim \mathcal{N}(\mu_a, \nu^2)$. Let $(W_v)^2 = \sum_{e \in E(v)} [w_v(e)]^2$.

Fix alternatives $a, b \in A$ such that $a \succ_{\sigma^*} b$ (thus, $\mu_a > \mu_b$). Note that $Y_{v,a} - Y_{v,b} \sim \mathcal{N}(\mu_a - \mu_b, 2\nu^2 (W_v)^2)$. Now, recall that there is an edge from a to b in the pairwise majority

²Available at: <http://www.cs.cmu.edu/~arielpro/papers.html>

graph if a strict majority of the voters prefer a to b , i.e., if $\sum_{v \in V} \mathbb{I}[Y_{v,a} > Y_{v,b}] > n/2$ (where \mathbb{I} is the indicator random variable). Hence, in Lemma 1 we take $Z_v^1 = \mathbb{I}[Y_{v,a} > Y_{v,b}]$ and $Z_v^2 = \mathbb{I}[Y_{v,a} \leq Y_{v,b}]$. Finally, we complete the proof by showing that the two conditions required by Lemma 1 hold.

Condition 1: $\mathbb{E}[Z_v^1] - \mathbb{E}[Z_v^2] \geq \gamma$, where $\gamma > 0$ is a constant. Note that $\mathbb{E}[Z_v^1] - \mathbb{E}[Z_v^2] = 2 \cdot \Pr[Y_{v,a} > Y_{v,b}] - 1$. Since $Y_{v,a} - Y_{v,b} \sim \mathcal{N}(\mu_a - \mu_b, 2\nu^2 (W_v)^2)$, we have that

$$\Pr[Y_{v,a} - Y_{v,b} > 0] = \Phi\left(\frac{\mu_a - \mu_b}{\sqrt{2}\nu W_v}\right) \geq \Phi\left(\frac{\Delta}{\sqrt{2}\nu}\right) \geq \frac{1}{2} + \gamma',$$

where $\gamma' > 0$ is a constant. Here, the second transition holds because $W_v \leq 1$ and $\mu_a - \mu_b \geq \Delta$, and the final transition is a standard property of the Gaussian distribution. Hence, condition 1 holds with $\gamma = 2\gamma'$.

Condition 2: $\Pr[Z_v^1 \leq Z_v^2 + \delta] = O(1/(d_v)^{1+\epsilon})$, where $\epsilon, \delta > 0$ are constants. Take $\delta = 0.5$, and recall that Z_v^1 and Z_v^2 are indicator random variables. Then,

$$\Pr[Z_v^1 \leq Z_v^2 + \delta] = \Pr[Z_v^1 = 0 \vee Z_v^2 = 1] = \Pr[Y_{v,a} \leq Y_{v,b}].$$

Since $Y_{v,a} - Y_{v,b} \sim \mathcal{N}(\mu_a - \mu_b, 2\nu^2 (W_v)^2)$, we have

$$\Pr[Y_{v,a} - Y_{v,b} \leq 0] = 1 - \Phi(\lambda) \leq \frac{1}{\sqrt{2\pi} \cdot \lambda} e^{-\lambda^2/2},$$

where $\lambda = (\mu_a - \mu_b)/(\sqrt{2}\nu W_v)$, and the last transition is a standard upper bound for Gaussian distributions. Substituting our assumption that $(W_v)^2 \leq \Delta^2/(8\nu^2 \ln d_v)$ and simplifying, we obtain that the probability is $O(1/(d_v)^2)$. Hence, condition 2 holds with $\epsilon = 1$.

Since both conditions are satisfied, Lemma 1 implies that every PM-c rule is accurate in the limit.

Modal Ranking Rule. Recall that the modal ranking rule chooses the most frequent ranking in the profile. Thus, we need to show that the ground truth ranking appears more frequently than any other ranking. Fix a ranking $\sigma \neq \sigma^*$, and define $Z_v^1 = \mathbb{I}[\sigma_v = \sigma^*]$ and $Z_v^2 = \mathbb{I}[\sigma_v = \sigma]$. Then, we wish to use Lemma 1 to show that the number of occurrences of σ^* in the profile is larger than the number of occurrences of σ with probability 1 as n goes to infinity. Applying the union bound over all rankings $\sigma \neq \sigma^*$ would imply that σ^* would be the most frequent ranking in the profile with probability 1 as n goes to infinity. Thus, the modal ranking rule would be accurate in the limit.

Next, we show that the two conditions of Lemma 1 hold.

Condition 1: $\mathbb{E}[Z_v^1] - \mathbb{E}[Z_v^2] \geq \gamma$, where $\gamma > 0$ is a constant. To derive this, we leverage a result by Jiang et al. [2014]. Using techniques from the proof of their Theorem 2, it can be shown that if we obtain a ranking σ by sampling utilities from Gaussians and ordering the alternatives by their sampled utilities, then for any ranking $\tau \in \mathcal{L}(A)$ and alternatives $a, b \in A$ such that $a \succ_{\sigma^*} b$ and $a \succ_{\tau} b$, we have $\Pr[\sigma = \tau] - \Pr[\sigma = \tau_{a \leftrightarrow b}]$ is at least a positive constant, where $\tau_{a \leftrightarrow b}$ denotes the ranking obtained by swapping alternatives a and b in τ . That is, swapping two alternatives

to match their order as in σ^* increases the probability of the ranking being sampled by at least a positive constant. However, this result uses a lower bound on the variances of the Gaussian distributions from which quality estimates are sampled. In our case, no such lower bound may exist for vertices with high degree. However, in the absence of such a lower bound one can still show that for the ranking σ_v of voter v , we have that $\Pr[\sigma_v = \tau] - \Pr[\sigma_v = \tau_{a \leftrightarrow b}]$ is non-negative for every $\tau \in \mathcal{L}(A)$ (with $a \succ_{\tau} b$), and is at least a positive constant γ' when $\tau = \sigma^*$; this is formally presented as a lemma in the full version of the paper.

Finally, to show that $\Pr[\sigma_v = \sigma^*] - \Pr[\sigma_v = \sigma] \geq \gamma$ (where $\gamma > 0$ is a constant), we start from ranking σ and perform ‘‘bubble sort’’ to convert it into σ^* . That is, in each iteration we find a pair that is ordered differently than in σ^* , and swap the pair. Note that this process converges to σ^* in at most m^2 iterations, and the probability of the ranking never decreases, and increases by at least γ' in the last iteration. This proves that condition 1 holds with $\gamma = \gamma'$.

Condition 2: $\Pr[Z_v^1 \leq Z_v^2 + \delta] = O(1/(d_v)^{1+\epsilon})$ for constants $\delta, \epsilon > 0$. This condition is very easy to establish. Again, take $\delta = 0.5$. Then,

$$\begin{aligned} \Pr[Z_v^1 \leq Z_v^2 + \delta] &= \Pr[Z_v^1 = 0 \vee Z_v^2 = 1] = \Pr[\sigma_v \neq \sigma^*] \\ &\leq \sum_{a,b \in A : a \succ_{\sigma^*} b} \Pr[Y_{v,a} \leq Y_{v,b}], \end{aligned}$$

where the last transition holds because if σ_v does not match σ^* , then there exist alternatives a and b such that $a \succ_{\sigma^*} b$ but $b \succ_{\sigma_v} a$ (thus, $Y_{v,a} \leq Y_{v,b}$). However, note that this probability is at most m^2 times the probability obtained in condition 2 for PM-c rules, which was $O(1/(d_v)^2)$. Because the number of alternatives m is a constant in our model, multiplying by m^2 does not increase the order in terms of d_v . Hence, condition 2 also holds with $\epsilon = 1$.

In conclusion, Lemma 1 implies that the modal ranking rule is accurate in the limit, as required. ■

While all strict positional scoring rules are accurate in the limit irrespective of the social network structure, one can show that other positional scoring rules such as plurality are not always accurate in the limit; an example is presented in the full version of the paper.

6 Unequal Variance

In the previous section we showed that PM-c rules, the modal ranking rule, and strict scoring rules are accurate in the limit when the independent conversations on the edges produce quality estimates from Gaussian distributions (with equal variance) and voters aggregate them using a weighted mean. The equal variance assumption is perhaps the most restrictive assumption in the model of Section 5. In this section, we analyze a more general model, which is identical to the model of Section 5, except for allowing Gaussians with different variance. Formally, we instantiate our general model using the following model of quality estimates.

Quality estimates. For each edge $e \in E$ and alternative $a \in A$, assume $X_{e,a} \sim \mathcal{N}(\mu_a, (\nu_{e,a})^2)$. Crucially, we assume

that all $(\nu_{e,a})^2$ are upper bounded by a global constant. For notational convenience, denote this constant by ν^2 . Hence, $(\nu_{e,a})^2 \leq \nu^2$ for all $e \in E$ and $a \in A$.

Computer-based simulations provided non-trivial counterexamples (presented in the full version) showing that unequal variance invalidates Theorem 1 with respect to strict positional scoring rules and the modal ranking rule.

Theorem 2. *There exist a social network graph $G = (V, E)$, true qualities of alternatives $\{\mu_a\}_{a \in A}$, and Gaussian random variables $X_{e,a}$ for each edge $e \in E$ and alternative $a \in A$ whose variances depend on the alternative a , for which the modal ranking rule is not accurate in the limit, and there exists a strict scoring rule (in particular, Borda count) which is not accurate in the limit.*

In a nutshell, the key insight is that we find a Gaussian distribution for each alternative $a \in A$ such that ranking the alternatives based on a quality estimate sampled from their Gaussian distribution leads to: (i) a ranking other than the true ranking is returned with a probability higher than that of the true ranking itself, which causes the modal ranking rule to fail to achieve accuracy in the limit, and (ii) the probabilities of different alternatives being placed in various positions is such that between two alternatives, the less preferred alternative in the true ranking has greater expected Borda score than the more preferred alternative, causing Borda count to violate accuracy in the limit. Despite these counterintuitive phenomena, it holds that the top alternative in the true ranking is ranked higher than the alternative ranked second in the true ranking with probability strictly greater than $1/2$, and a similar statement also holds for all other pairs of alternatives, thereby ensuring that PM-c rules are accurate in the limit.

Happily, the success of PM-c rules is not a coincidence. Indeed, note that in our proof of Theorem 1 we leverage the results of Jiang et al. [2014] to prove condition 1 of Lemma 1 for PD-c rules and the modal ranking rule. Jiang et al. crucially assume that all distributions have equal variance, and their results break down when this assumption is violated. On the other hand, our proof for the accuracy in the limit of PM-c rules does not rely on their results, and, in fact, does not make use of the equal variance assumption. Specifically, with unequal variance we have that $Y_{v,a} - Y_{v,b}$ follows the Gaussian distribution $\mathcal{N}(\mu_a - \mu_b, \sum_{e \in E(v)} [w_v(e)]^2 \cdot ((\nu_{e,a})^2 + (\nu_{e,b})^2))$. Note that our proof only uses an upper bound on the variance of this Gaussian distribution, and the variance is still upper bounded by $2\nu^2 (W_v)^2$. Hence, for PM-c rules, the proof of Theorem 1 goes through even with unequal variance, and shows that all PM-c rules are accurate in the limit.

Theorem 3. *Assume that there exists a constant ν such that $(\nu_{e,a})^2 \leq \nu^2$ for all $e \in E$ and $a \in A$, and a universal constant $D \in \mathbb{N}$ such that $\sum_{e \in E(v)} [w_v(e)]^2 \leq \Delta^2 / (8\nu^2 \ln d_v)$ for all voters v with degree $d_v \geq D$. Then, all PM-c rules are accurate in the limit irrespective of the choices of the unknown parameters: the true qualities $\{\mu_a\}_{a \in A}$, the variances $\{(\nu_{e,a})^2\}_{e \in E, a \in A}$, and the weights $\{w_v(e)\}_{v \in V, e \in E(v)}$.*

Theorem 3 establishes that PM-c rules are qualitatively more robust than PD-c rules and the modal ranking rule in our setting: While PD-c rules and the modal ranking rule

lose their accuracy in the limit when relaxing the equal variance assumption, PM-c rules still guarantee accuracy in the limit irrespective of all unknown parameters. In fact, observe that in our proof for PM-c rules, we only require that for every pair of alternatives $a, b \in A$ with $a \succ_{\sigma^*} b$, we have (i) $\Pr[Y_{v,a} > Y_{v,b}] > 1/2$, and (ii) both $Y_{v,a}$ and $Y_{v,b}$ are *sufficiently concentrated* around their respective means μ_a and μ_b so that $\Pr[Y_{v,a} \leq Y_{v,b}] = o(1/d_v)$. Using this observation, we can extend the robustness of PM-c rules beyond the restrictions imposed by Theorem 3 in both dimensions: the possible distributions on the edges and the possible aggregation rules used by the voters.

For example, leveraging an elegant extension of the classic McDiarmid inequality by Kontorovich [2013], we can show that PM-c rules are accurate in the limit when the distributions on the edges have finite “subgaussian diameter” (this includes all distributions with bounded support and all Gaussian distributions) and voters use weighted mean aggregation. On the other hand, using a concentration inequality for medians, one can show that when the distributions on the edges are Gaussians with bounded variance, then the voters could also use weighted median (instead of weighted mean) aggregation, and PM-c rules would remain accurate in the limit.

7 Discussion

Let us briefly discuss several pertinent issues.

Temporal dimension. While in our model each voter performs a one-time, synchronous aggregation of information from its incident edges, in general voters may perform multiple and/or asynchronous updates. After k updates, the information possessed by a voter would be a weighted aggregation of the information from all nodes up to distance k from the voter, although the weight associated with another voter at distance k would presumably be exponentially small in k . Deriving positive robustness results in this model seems to require making our simple covariance bounds more sensitive to weights. We believe that Gaussian hypercontractivity results [Mossel, 2010] may be helpful in this context.

Opinions on vertices. The *independence* part of our extension of the independent conversations model seems to be a restrictive assumption because the conversations of a voter with two other voters are likely to be positively correlated through the beliefs of the voter. In this sense, it seems more natural to consider a model where the opinions are attached to vertices rather than edges. Specifically, one might consider a model where the prior opinion of each voter is first drawn from a distribution, and then voters are allowed to aggregate opinions from their neighbors. This leads to immediate impossibilities. Indeed, consider a star network where all peripheral voters give weight 1 to the central voter and 0 to themselves (this does not violate the conditions of Theorem 1). At the end, all peripheral voters would have perfectly correlated votes, coinciding with the prior opinion of the central voter which is inaccurate with a significant probability. It follows that any reasonable anonymous voting rule, which would output this opinion, would not be accurate in the limit. Interestingly, we can circumvent this impossibility easily if we know the social

network structure: We can simply return the vote submitted by the central voter, which is guaranteed to be accurate as the central voter assimilates information from many sources.

Ground truth and opinion formats. Finally, we assume that the ground truth is a true quality for each alternative, which led to a random utility based model. Another compelling alternative is to assume that the ground truth is only an ordinal ranking of the alternatives. In this case, the samples on the edges would also be rankings (instead of noisy quality estimates), and voters would aggregate rankings on their incident edges using their own local voting rules. This model gives rise to many counterintuitive phenomena. For example, using Borda count to aggregate two rankings sampled from the popular Mallows' model [Mallows, 1957] with noise parameters $\varphi = 0.1$ and $\varphi = 0.9$ leads to a ranking that is not the ground truth being returned with higher probability than the ground truth itself, ultimately showing that Borda count would not be accurate in the limit. Remarkably, popular PM-c rules seem to be robust against such examples, hinting at the possibility that PM-c rules may also possess compelling robustness properties in this model.

References

- [Azari Soufiani *et al.*, 2012] H. Azari Soufiani, D. C. Parkes, and L. Xia. Random utility theory for social choice. In *Proc. of 26th NIPS*, pages 126–134, 2012.
- [Azari Soufiani *et al.*, 2013] H. Azari Soufiani, D. C. Parkes, and L. Xia. Preference elicitation for general random utility models. In *Proc. of 29th UAI*, pages 596–605, 2013.
- [Azari Soufiani *et al.*, 2014] H. Azari Soufiani, D. C. Parkes, and L. Xia. Computing parametric ranking models via rank-breaking. In *Proc. of 31st ICML*, pages 360–368, 2014.
- [Caragiannis *et al.*, 2013] I. Caragiannis, A. D. Procaccia, and N. Shah. When do noisy votes reveal the truth? In *Proc. of 14th EC*, pages 143–160, 2013.
- [Caragiannis *et al.*, 2014] I. Caragiannis, A. D. Procaccia, and N. Shah. Modal ranking: A uniquely robust voting rule. In *Proc. of 28th AAAI*, pages 616–622, 2014.
- [Conitzer and Sandholm, 2005] V. Conitzer and T. Sandholm. Communication complexity of common voting rules. In *Proc. of 6th EC*, pages 78–87, 2005.
- [Conitzer *et al.*, 2009] V. Conitzer, M. Rognlie, and L. Xia. Preference functions that score rankings and maximum likelihood estimation. In *Proc. of 21st IJCAI*, pages 109–115, 2009.
- [Conitzer, 2012] V. Conitzer. Should social network structure be taken into account in elections? *Mathematical Social Sciences*, 64(1):100–102, 2012.
- [Conitzer, 2013] V. Conitzer. The maximum likelihood approach to voting on social networks. In *Proc. of 51st Allerton*, pages 1482–1487, 2013.
- [DeGroot, 1974] M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- [Elkind *et al.*, 2010] E. Elkind, P. Faliszewski, and A. Slinko. Good rationalizations of voting rules. In *Proc. of 24th AAAI*, pages 774–779, 2010.
- [Jiang *et al.*, 2014] A. X. Jiang, L. S. Marcolino, A. D. Procaccia, T. Sadholm, N. Shah, and M. Tambe. Diverse randomized agents vote to win. In *Proc. of 28th NIPS*, pages 2573–2581, 2014.
- [Kleinberg, 2007] J. Kleinberg. Cascading behavior in networks: Algorithmic and economic issues. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 24. Cambridge University Press, 2007.
- [Kontorovich, 2013] A. Kontorovich. Concentration in unbounded metric spaces and algorithmic stability. arXiv:1309.1007, 2013.
- [Lu and Boutilier, 2011] T. Lu and C. Boutilier. Robust approximation and incremental elicitation in voting protocols. In *Proc. of 22nd IJCAI*, pages 287–293, 2011.
- [Mallows, 1957] C. L. Mallows. Non-null ranking models. *Biometrika*, 44:114–130, 1957.
- [Mao *et al.*, 2013] A. Mao, A. D. Procaccia, and Y. Chen. Better human computation through principled voting. In *Proc. of 27th AAAI*, pages 1142–1148, 2013.
- [Mossel *et al.*, 2014] E. Mossel, J. Neeman, and O. Tamuz. Majority dynamics and aggregation of information in social networks. *Autonomous Agents and Multi-Agent Systems*, 28(3):408–429, 2014.
- [Mossel, 2010] E. Mossel. Gaussian bounds for noise correlation of functions. *Geometric and Functional Analysis*, 19(6):1713–1756, 2010.
- [Mosteller, 1951] F. Mosteller. Remarks on the method of paired comparisons: I. the least squares solution assuming equal standard deviations and equal correlations. *Psychometrika*, 16(1):3–9, 1951.
- [Procaccia *et al.*, 2012] A. D. Procaccia, S. J. Reddi, and N. Shah. A maximum likelihood approach for selecting sets of alternatives. In *Proc. of 28th UAI*, pages 695–704, 2012.
- [Thurstone, 1927] L. L. Thurstone. A law of comparative judgement. *Psychological Review*, 34:273–286, 1927.
- [Xia and Conitzer, 2011] L. Xia and V. Conitzer. A maximum likelihood approach towards aggregating partial orders. In *Proc. of 22nd IJCAI*, pages 446–451, 2011.
- [Xia *et al.*, 2010] L. Xia, V. Conitzer, and J. Lang. Aggregating preferences in multi-issue domains by using maximum likelihood estimators. In *Proc. of 9th AAMAS*, pages 399–408, 2010.
- [Young, 1988] H. P. Young. Condorcet's theory of voting. *The American Political Science Review*, 82(4):1231–1244, 1988.