Correlated Voting

Correlated Voting

Debmalya Mandal
dmandal@g.harvard.edu
Harvard University

David C. Parkes
parkes@eecs.harvard.edu
Harvard University

Abstract

We study the social choice problem where a group of n voters report their preferences over alternatives and a voting rule is used to select an alternative. We show that when the preferences of voters are positively correlated according to the Kendall-Tau distance, the probability that any scoring rule is not *ex post* incentive compatible (EPIC) goes to zero exponentially fast with the number of voters, improving over the previously known rate of $1/\sqrt{n}$ for independent preferences. Motivated by rank-order models from machine learning, we introduce two examples of positively-correlated models, namely *Conditional Mallows* and *Conditional Plackett-Luce*. Conditional Mallows satisfies Kendall-Tau correlation and fits our positive result. We also prove that Conditional Plackett-Luce becomes EPIC exponentially quickly.

1 Introduction

Social choice theory studies how to aggregate preferences and select an outcome. A canonical problem is voting, where reports are preferences over a list of candidates and a voting rule selects the winning candidate.

The problem of voting is ubiquitous systems in which multiple agents interact with each other. Several tools and ideas from social choice theory have found applications in different areas of multi-agent systems including resource allocation [Parkes *et al.*, 2015], rank aggregation [Altman and Tennenholtz, 2010], recommender systems [Pennock *et al.*, 2000], choosing between multiple issues [Lang and Xia, 2009], and crowdsourcing [Mao *et al.*, 2013]. Voting also has strong connections to rank aggregation models in machine learning. These models view rank data as a noisy version of a true underlying rank, and try to recover the true rank.

It is typical to ignore the possibility that data is misreported in order to change the aggregate view. Consider, for example, the problem of a conference review system where a reviewer gives a ranking over a subset of submitted papers. The committee wants to aggregate truthful reports of opinion on submitted papers while different subcommunities may want to push decisions in one direction or another.

Indeed, the problem of incentive-aligned social choice is generally hopeless. Gibbard [1973] and Satterthwaite [1975] prove that when the number of alternatives is at least three, a voting rule is unanimous and strategy-proof if and only if it is dictatorial. There are two standard ways of circumventing the Gibbard-Satterthwaite impossibility result. First, the preferences of voters may arise from a restricted preference class [Barberà, 2011]. Second, some have appealed to worst-case and average-case computational intractability of the problem of strategic manipulation, arguing that this provides robustness to strategic behavior ([Faliszewski and Procaccia, 2010] and for a recent survey, see [Conitzer and Walsh, 2016]).

However, these approaches do not consider a probabilistic model of social choice, with preferences sampled from a distribution. The relevant incentive question is to consider the situation that each voter may have information about how likely the preferences of other voters' are. Our particular interest is in studying the problem of incentive-aligned preference aggregation when voters' preferences are not necessarily independent. Whereas it is typical to assume that individual rank preferences of voters are independent, this is not always true. For example, in a conference review mechanism, once a reviewer realizes that paper A is significantly better than paper B, they may be more likely to place more probability on others thinking the same about the papers.

Majumdar and Sen [2004] initiated the study of *ordinally Bayesian incentive compatible (OBIC)* voting rules for the setting in which voters have incomplete information about each others' preferences. OBIC is a weaker form of strategy-proofness, stating that truthful reporting is a Bayes-Nash equilibrium for any cardinal utility consistent with agent preferences. Later, Sen et al. [2015] prove that the Gibbard-Satterthwaite impossibility result breaks down if beliefs are positively correlated. These authors start with a social choice function f satisfying certain properties and exhibit a class of positively correlated distributions for which f is OBIC.

Instead of OBIC, we adopt the notion of *ex post incentive compatibility* (EPIC). EPIC is a weaker solution concept than strategy-proofness, and requires truth-telling to be a best response to every preference profile of others provided they are also truthful. Our goal is to understand the following question: *given positively correlated beliefs that are neither uniform nor extremely correlated, how do common voting rules,* *and in particular scoring rules, perform with regard to EPIC?*

We take an asymptotic view along the lines of Baharad and Neeman [2002], who prove that when the preferences of agents are drawn uniformly and independently at random the probability that any scoring rule is not EPIC goes to zero at a rate proportional to $1/\sqrt{n}$ for n agents.[1] Their result also holds for small, local correlations among preferences, but fails when the correlation is global or large enough. We first prove that positive correlation helps dramatically. Our main result is that the probability that any scoring rule is not EPIC goes to zero exponentially fast when an agent's belief is that the preferences of others are positively correlated with his own preference order according to the Kendall-Tau distance. We also establish a general result for any *conditionally independent and identical beliefs*, showing convergence of scoring rules to EPIC at a rate proportional to $1/\sqrt{n}$.

Motivated by rank-order models from machine learning [Marden, 1996], we introduce two examples of positively-correlated belief systems, namely *Conditional Mallows* and *Conditional Plackett-Luce*. These two families of belief systems span a wide range of positively correlated beliefs— from being arbitrarily close to uniform to being extremely correlated. Conditional Mallows is Kendall-Tau correlated and fits our main positive result. Conditional Plackett-Luce model is not, but we provide a different proof of exponential convergence to EPIC for this model.

2 Setup

The set of alternatives is $A = \{1, \ldots, m\}$ and the set of voters is $N = \{1, \ldots, n\}$. Let $\mathbb{P}$ be the set of all linear orders over A. Voter i has a preference order $P_i \in \mathbb{P}$ over the m alternatives. Any voting rule is represented by a *social choice function* $f : \mathbb{P}^n \to A$. We will write P_{-i} to denote a preference profile of all the voters other than i. We will use aP_ib to denote that alternative a is preferred over alternative b according to the preference order P_i.

Given preference P_i, let $\mu_i(P_{-i}|P_i)$ denote the probability that voter i ascribes to voters other than i having preference profile $P_{-i} \in \mathbb{P}^{n-1}$. Note that like Sen et al. [2015], we do not insist that the conditional probabilities $\mu_i(P_{-i}|P_i)$ should be generated from a given underlying common prior over the entire profile of n voters. We will call a collection of conditional probability distributions $\mu_1, \ldots, \mu_n$ a *belief system*.

We adopt *ex post* incentive compatibility (EPIC) as a solution concept. A voting rule is *strategy-proof* if truthful reporting is a dominant strategy for each voter. EPIC provides truthful reporting as a best response to any preference profile of others, provided they are also truthful.

Definition 2.1. A social choice function $f : \mathbb{P}^n \to A$ is EPIC with respect to the belief system $\{\mu_i\}_{i=1}^n$, if for each agent i and $\forall P_i, P_i'$

$$f(P_i, P_{-i})P_i f(P_i', P_{-i}) \;\forall P_{-i} \text{ in support of } \mu_i(\cdot|P_i) \quad (1)$$

For $\mu_i(\cdot|P_i)$ with full support over $\mathbb{P}^{n-1}$ for every P_i, then f is EPIC iff it is also strategy-proof. However, EPIC and strategy-proof social choice functions become very different when we compute the probability that an agent can manipulate. Suppose agent i has preference P_i and $L(P_i) = \{P_{-i} \in \mathbb{P}^{n-1} : \exists P_i' \; f(P_i', P_{-i})P_i f(P_i, P_{-i})\}$, the set of all preference profiles of others such that agent i wants to deviate. Then the probability that i can manipulate in the sense of a dominant strategy is either 0 or 1 (it is 0 if $L(P_i) = \emptyset$, and 1 otherwise). On the other hand, the probability that i can manipulate in the sense of an *ex post* Nash equilibrium is $\sum_{P_{-i} \in L(P_i)} \mu_i(P_{-i})$. We are interested in showing that the last probability becomes exponentially small in the number of agents when the belief system is positively correlated.

2.1 Positively Correlated Preferences

Sen et al. [2015] introduce two types of correlation for a given belief system $\{\mu_i\}_{i=1}^n$: *Top-Set correlation* and *Kendall-Tau correlation*. Let $B_k(P_i)$ denote the set of top k alternatives as ranked by P_i. A belief system is *Top-Set correlated* (TS-correlated) if every voter i believes that the event where every other voter has the same top-k set of alternatives as i is strictly more likely than the event where every other voter has some other subset T as their top-k set.

Definition 2.2 (Top-Set Correlated)**.** Belief system $\{\mu_i\}_{i=1}^n$ is *TS-correlated* if $\forall k \in \{1, \ldots, m-1\}, \forall i, \forall P_i, \; \forall T \neq B_k(P_i)$ and $|T| = k$:

$$\sum_{\substack{\{P_{-i}:\,\forall j\neq i \\ B_k(P_j)=B_k(P_i)\}}} \mu_i(P_{-i}|P_i) > \sum_{\substack{\{P_{-i}:\,\forall j\neq i \\ B_k(P_j)=T\}}} \mu_i(P_{-i}|P_i) \quad (2)$$

A belief system is *Kendall-Tau correlated* (KT-correlated) if every voter i believes that the preference profiles of other voters are ordered in decreasing probability according to increasing sum Kendall-Tau distance between the preferences of others and voter i's true preference P_i. Let $d(P_i, P_j)$ be the *Kendall-Tau distance* (i.e., the number of pairwise disagreements) between preferences P_i and P_j.

Definition 2.3 (Kendall-Tau Correlated)**.** Belief system $\{\mu_i\}_{i=1}^n$ is *KT-correlated* if for all P_i, P_{-i}, P_{-i}',

$$\mu_i(P_{-i}|P_i) > \mu_i(P_{-i}'|P_i) \text{ if} \sum_{j\neq i} d(P_j, P_i) < \sum_{j\neq i} d(P_j', P_i).$$

We use $D(P_{-i}|P_i)$ in place of $\sum_{j\neq i} d(P_j, P_i)$. It is easy to show that any KT-correlated belief system is also TS-correlated, but not vice-versa. [Bhargava *et al.*, 2011]

Definition 2.4. A belief system $\{\mu\}_{i=1}^n$ is *conditionally independent and identically distributed* (c.i.i.d.) if

$$\forall i,\; \forall P_i,\; \forall P_{-i}: \quad \mu_i(P_{-i}|P_i) = \prod_{j\neq i} \nu(P_j|P_i) \quad (3)$$

where $\nu(\cdot|P_i)$ is a probability distribution over orders $\mathbb{P}$.

We work with c.i.i.d. belief systems. We will see examples of different families of positively correlated and c.i.i.d. belief systems in Section 4.

[1] Although Baharad and Neeman [2002] use the term "asymptotic strategyproofness", they actually consider EPIC as the notion of equilibrium

3 Scoring Rules

Scoring rules [Young, 1975] are defined in the following way. Fix a non-decreasing sequence of real numbers, $s_1 \geq s_2 \geq \ldots \geq s_m$, such that $s_1 > s_m$. If a voter ranks an alternative x at position j then x gets a score of s_j. We will write $sc(j, P_i)$ to denote the score of an alternative j according to the preference P_i. The score of an alternative is the sum of the scores received from all the voters. The alternative with the highest score is chosen as the outcome of the election. In case there is a tie, a winning alternative is selected according to some tie-breaking rule. We insist that $s_1 > s_m$, otherwise if $s_1 = s_m$, every alternative receives the same score and the resulting social choice function just depends on the tie-breaking rule. Some popular scoring rules are: Plurality $(1, 0, 0, \ldots, 0)$, Borda $(m-1, m-2, \ldots, 1, 0)$ and Veto $(1, 1, 1, \ldots, 1, 0)$.

Baharad and Neeman [2002] prove that the probability that any scoring rule is not EPIC goes to zero at rate of $1/\sqrt{n}$ with the number of agents n. They consider the following setting: (1) each voter is equally likely to have any preference order, i.e. the marginal distribution is uniform, and (2) the preferences of the voters are locally correlated, i.e. the voters can be numbered in a sequence such that the further they are apart, the more independent their preferences become. In an independent work, Slinko [2002] shows that the number of manipulable preference profiles (profiles where some agent can benefit by deviating unilaterally) goes to zero at a rate of $1/\sqrt{n}$ for any scoring rule. This result also proves the same rate of convergence to EPIC for uniform i.i.d. preferences. We prove that when the preferences of the voters are KT-correlated and c.i.i.d., the probability that any given scoring rule is not EPIC goes to zero exponentially fast. Our result strengthens what is known about the asymptotic non-manipulability of scoring rules.

3.1 EPIC Convergence for KT-correlation

Suppose $\{\mu_i\}_{i=1}^n$ is a c.i.i.d. belief system, that is, for all i, we have $\mu_i(P_{-i}|P_i) = \prod_{j \neq i} \nu(P_j|P_i)$. Consider any preference $P_i \in \mathbb{P}$ and any two alternatives a and b such that aP_ib. Let

$$\mu_{a,b}(P_i) = \mathrm{E}_{P \sim \nu(\cdot|P_i)}\left[sc(a, P_i) - sc(b, P_i)\right] \qquad (4)$$

denote the expected difference of scores between alternatives a and b for a random preference order of some other agent, this preference drawn according to the conditional distribution $\nu(\cdot|P_i)$. Let $\sigma^2_{a,b}(P_i)$ denote the variance of the difference in score between a and b. We now state our main result.

Theorem 1. *Suppose a belief system $\{\mu_i\}_{i=1}^n$ is c.i.i.d. and KT-correlated. Then the probability that any given scoring rule is not EPIC w.r.t. $\{\mu_i\}_{i=1}^n$ goes to zero at a rate proportional to $O\left(e^{-cn}\right)$, for some constant $c > 0$.*

A useful lemma establishes that $\mu_{a,b}(P_i) > 0$, for any preference P_i and alternatives a, b such that aP_ib.

Lemma 2. *Suppose the belief system $\{\mu_i\}_{i=1}^n$ is c.i.i.d. and KT-correlated. Consider a preference ordering P_i, and alternatives a and b such that aP_ib. Then $\mu_{a,b}(P_i) > 0$.*

Proof. Select two preference orderings P_j and P'_j such that $d(P_j, P_i) < d(P'_j, P_i)$. Now consider the following two preference profiles for voters other then i:

1. $P^1_{-i} = (P_i, \ldots, P_i, P_j, P_i, \ldots, P_i)$
2. $P^2_{-i} = (P_i, \ldots, P_i, P'_j, P_i, \ldots, P_i)$

Then $D(P^1_{-i}, P_i) = d(P_j, P_i) < d(P'_j, P_i) = D(P^2_{-i}, P_i)$. Since $\{\mu_i\}_{i=1}^n$ is KT-correlated, we have $\mu_i(P^1_{-i}|P_i) > \mu_i(P^2_{-i}|P_i)$. Furthermore, $\{\mu_i\}_{i=1}^n$ is c.i.i.d., therefore there exists a function ν such that $\mu_i(P_{-i}|P_i) = \prod_{j \neq i} \nu(P_j|P_i)$. $\mu_i(P^1_{-i}|P_i) > \mu(P^2_{-i}|P_i)$ implies

$$\nu(P_i|P_i)^{n-2}\nu(P_j|P_i) > \nu(P_i|P_i)^{n-2}\nu(P'_j|P_i).$$

Therefore, we have $d(P_j, P_i) < d(P'_j, P_i)$ implies $\nu(P_j|P_i) > \nu(P'_j|P_i)$. Now, partition the set of all preference orderings $\mathbb{P}$ into $\mathbb{P}_{a \succ b}$ and $\mathbb{P}_{b \succ a}$. Every preference ordering P in $\mathbb{P}_{a \succ b}$ ranks a above b and vice versa for $\mathbb{P}_{b \succ a}$. Note that there exists a one-to-one mapping $f : \mathbb{P}_{a \succ b} \to \mathbb{P}_{b \succ a}$, namely $f(P)$ be the same as P except positions of alternatives a and b exchanged. Then

$$\begin{aligned}\mu_{a,b}(P_i) &= \sum_{P \in \mathbb{P}} (sc(a, P) - sc(b, P))\nu(P|P_i) \\ &= \sum_{P \in \mathbb{P}_{a \succ b}} (sc(a, P) - sc(b, P))(\nu(P|P_i) - \nu(f(P)|P_i)).\end{aligned}$$

Now for all $P \in \mathbb{P}_{a \succ b}$, $d(P, P_i) < d(f(P), P_i)$ and therefore $\nu(P|P_i) > \nu(f(P)|P_i)$. Moreover there exists at least one P such that $sc(a, P) > sc(b, P)$ (e.g. when P places a at top and b at bottom). This proves that $\mu_{a,b}(P_i) > 0$. □

We will also use the Chernoff-Hoeffding inequality for bounded random variables.

Lemma 3. *(Theorem 2,[Hoeffding, 1963]) Let $X_1, X_2, \ldots, X_n$ be independent and for all i, $a_i \leq X_i \leq b_i$. Let $\bar{X}_n = \frac{1}{n}\sum_{i=1}^n X_i$ and $\mu = \mathrm{E}\left[\frac{1}{n}\sum_{i=1}^n X_i\right]$. Then*

$$\Pr\left[\bar{X}_n \leq \mu - t\right] \leq \exp\left\{-\frac{2n^2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right\}$$

.

Proof. (of Theorem 1.) Suppose voter 1 observes a preference order x_1. Let X_i be the random variable corresponding to the preference ordering seen by agent i. Conditioned on the event $X_1 = x_1$, voter 1 believes that the random variables $X_2, X_3, \ldots$ are i.i.d. with distribution $\nu(\cdot|x_1)$.

Fix any two alternatives a and b such that $a\ x_1\ b$. We show that the probability that agent 1 can improve the ordering of a vs b in the social ranking, through making some misreport of his preference order, falls exponentially quickly in the number of agents n. For $i = 2, 3, \ldots$, let $Z^i_{a,b} = sc(a, X_i) - sc(b, X_i)$ be the difference of scores between alternatives a and b assigned by voter i's true preference X_i. Since $X_2, X_3, \ldots$ are i.i.d., so are $Z^2_{a,b}, Z^3_{a,b}, \ldots$. As introduced earlier, $\mathrm{E}\left[Z^i_{a,b}\right] = \mu_{a,b}(x_1)$ and $\mathrm{Var}(Z^i_{a,b}) = \sigma^2_{a,b}(x_1)$. Define $Z^n_{a,b} = \sum_{i=2}^{n+1} Z^i_{a,b}$ (where voters 2 through $n+1$ each have c.i.i.d. preferences).

Now voter 1 wants to manipulate and report a preference ordering different from x_1 only if $s_m - s_1 \leq Z^n_{a,b} \leq$

$sc(b,x_1) - sc(a,x_1)$. To see this, first suppose $Z^n_{a,b} < s_m - s_1 < 0$. Then even by placing a at top and b at bottom, voter 1 can increase the difference in scores between a and b by at most $s_1 - s_m$ and this still fails to cause a to rank higher than b. So, in this case, voter 1 is happy to report x_1.

On the other hand, suppose $Z^n_{a,b} > sc(b,x_1) - sc(a,x_1)$. If 1 reports truthfully then the net difference of scores between a and b is $Z^n_{a,b} + sc(a,x_1) - sc(b,x_1) > 0$ and a and b are already ordered according to 1's preferences.

Let $\Delta_{a,b} = sc(b,x_1) - sc(a,x_1)$. Note that, $\Delta_{a,b} \le 0$ since $a\ x_1\ b$. Then the probability of manipulation by voter 1 is bounded by $\Pr\left[s_m - s_1 \le Z^n_{a,b} \le \Delta_{a,b}\right]$, which we bound using Chernoff-Hoeffding inequality (Lemma 3) as follows :

$$\Pr\left[s_m - s_1 \le Z^n_{a,b} \le \Delta_{a,b}\right] \le \Pr\left[Z^n_{a,b} \le \Delta_{a,b}\right]$$
$$= \Pr\left[\frac{1}{n}\sum_{i=2}^{n+1} Z^i_{a,b} \le \frac{1}{n}\Delta_{a,b}\right]$$
$$= \Pr\left[\frac{1}{n}\sum_{i=2}^{n+1} Z^i_{a,b} \le \mu_{a,b}(x_1) - t\right].$$

Where $t = \mu_{a,b}(x_1) - \frac{1}{n}\Delta_{a,b}$. Since $Z^i_{a,b}$ is the difference of scores between alternatives a and b according to X_i, we have $Z^i_{a,b} \in [s_m - s_1, s_1 - s_m]$. Now applying Chernoff-Hoeffding inequality we get an upper bound of

$$\exp\left\{-\frac{2n^2t^2}{4n(s_1-s_m)^2}\right\} = \exp\left\{-\frac{n}{2}\left(\frac{\mu_{a,b}(x_1)}{s_1-s_m}\right)^2 + \frac{\mu_{a,b}(x_1)\Delta_{a,b}}{(s_1-s_m)^2} - \frac{1}{2n}\left(\frac{\Delta_{a,b}}{s_1-s_m}\right)^2\right\}.$$

As n goes to infinity, $e^{-\gamma/n}$ goes to 1 for any constant $\gamma > 0$. Therefore, for large enough n, the probability of manipulation is $O\left(e^{-c_{a,b}(x_1)n}\right)$ where $c_{a,b}(x_1) = \frac{1}{2}\left(\frac{\mu_{a,b}(x_1)}{s_1-s_m}\right)^2$. Since the preference ordering x_1 and the alternatives a and b can arbitrary, using a union bound we get the actual probability of manipulation is e^{-cn} where the constant c takes the worst-case, and

$$c = \min_{x\in\mathbb{P}} \min_{a\ x_1\ b} \frac{1}{2}\left(\frac{\mu_{a,b}(x_1)}{s_1-s_m}\right)^2. \tag{5}$$

Lemma 2 proves that for any KT-correlated and c.i.i.d. belief system, $\mu_{a,b}(x_1) > 0$ for any x_1 and any a and b such that $a\ x_1\ b$. This proves that $c > 0$, and finishes the proof. ☐

3.2 EPIC Convergence for TS-correlation

Lemma 2 need not be true for a TS-correlated belief system. Here is a counter-example. Consider a situation with 2 voters, 3 alternatives, and Borda scoring rule $(2,1,0)$. Suppose voter 1's preference ordering is $1\ P_1\ 2\ P_1\ 3$ (in short 123). We now construct a TS-correlated belief system for which $\mu_{1,2}(123) = \mathrm{E}_{X_2\sim\nu(\cdot|123)}\left[Z^2_{1,2}\right] < 0$.

Any TS-correlated belief system needs to satisfy the following system of inequalities (we drop the conditioning on preference ordering 123 for notational convenience):

$$\mu_1(123) + \mu_1(132) > \mu_1(213) + \mu_1(231)$$
$$\mu_1(123) + \mu_1(132) > \mu_1(312) + \mu_1(321)$$
$$\mu_1(123) + \mu_1(213) > \mu_1(132) + \mu_1(312)$$
$$\mu_1(123) + \mu_1(213) > \mu_1(231) + \mu_1(321)$$

Let us choose the following distribution: $\mu_1(123) = 1 - 6\varepsilon, \mu_1(132) = \mu_1(213) = \mu_1(312) = \mu(321) = \varepsilon, \mu_1(231) = 2\varepsilon$. It can be easily verified that μ_1 satisfies the inequalities above as long as $\varepsilon < 1/8$. Now $\mu_{1,2}(123) = 1 - 9\varepsilon < 0$ if $\varepsilon > 1/9$. To get a counter-example, one can choose any $\varepsilon \in (1/9, 1/8)$.

However, we now prove that the rate of convergence is not worse than $O\left(1/\sqrt{n}\right)$ for any c.i.i.d. belief system. Our proof uses the Berry-Esseen theorem which quantifies the rate of convergence of the central limit theorem.

Lemma 4. *(Berry-Esseen) Let $X_1, X_2, \ldots, X_n$ be i.i.d. with $\mathrm{E}[X_i] = 0, \mathrm{E}\left[X_i^2\right] = \sigma^2, \mathrm{E}\left[|X_i|^3\right] = \rho < \infty$. If $F_n(x)$ is the distribution function of $(X_1 + \ldots + X_n)/\sigma\sqrt{n}$ and $\Phi(x)$ is the distribution function of standard normal random variable then, $\forall x \in \mathbb{R}$,*

$$|F_n(x) - \Phi(x)| \le \frac{3\rho}{\sigma^3\sqrt{n}}.$$

See [Durrett, 2010] (Theorem 3.4.9) for a proof.

Theorem 5. *Suppose a belief system $\{\mu_i\}_{i=1}^n$ is c.i.i.d. Then the probability that any given scoring rule is not EPIC w.r.t. $\{\mu_i\}_{i=1}^n$ goes to zero at a rate proportional to $O\left(\frac{1}{\sqrt{n}}\right)$.*

Proof. As proved in Theorem 1, the probability of manipulation is bounded by $\Pr\left[s_m - s_1 \le Z^n_{a,b} \le \Delta_{a,b}\right]$. Now, we use the central limit theorem instead of Chernoff-Hoeffding inequality to bound the probability of manipulation. $Z^2_{a,b}, Z^3_{a,b}, \ldots$ are iid with mean $\mu_{a,b}(x_1)$ and variance $\sigma^2_{a,b}(x_1)$. Therefore, by the central limit theorem

$$\frac{Z^n_{a,b} - n\mu_{a,b}(x_1)}{\sigma_{a,b}(x_1)\sqrt{n}} \xrightarrow{d} \mathcal{N}(0,1). \tag{6}$$

Let us write $T_n = \frac{Z^n_{a,b} - n\mu_{a,b}(x_1)}{\sigma_{a,b}(x_1)\sqrt{n}}$. Then

$$\Pr\left[s_m - s_1 \le Z^n_{a,b} \le \Delta_{a,b}\right]$$
$$= \Pr\left[\frac{s_m - s_1 - n\mu_{a,b}(x_1)}{\sigma_{a,b}(x_1)\sqrt{n}} \le T_n \le \frac{\Delta_{a,b} - n\mu_{a,b}(x_1)}{\sigma_{a,b}(x_1)\sqrt{n}}\right]$$
$$= F_{T_n}\left(\frac{\Delta_{a,b} - n\mu_{a,b}(x_1)}{\sigma_{a,b}(x_1)\sqrt{n}}\right) - F_{T_n}\left(\frac{s_m - s_1 - n\mu_{a,b}(x_1)}{\sigma_{a,b}(x_1)\sqrt{n}}\right)$$

Here F_{T_n} is the distribution function of T_n. Now we use the following relation :

$$\forall l \le u,\ F_{T_n}(u) - F_{T_n}(l) \le |F_{T_n}(u) - F_{T_n}(l)|$$
$$\le |F_{T_n}(u) - \Phi(u)| + |\Phi(u) - \Phi(l)| + |\Phi(l) - F_{T_n}(l)|.$$

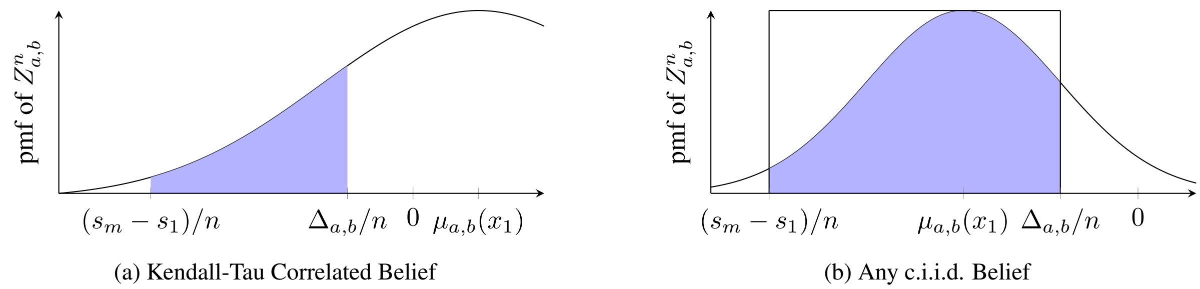

Figure 1: The blue regions determine the probability of manipulation. For a KT-correlated belief (1a) we use the tail probability and for any general c.i.i.d. belief (1b) we use the area of the surrounding rectangle.

We can bound the first and the third term by the Berry-Esseen Theorem (Lemma 4) (since $Z^i_{a,b}$ is a discrete distribution, its third absolute moment $\rho_{a,b}(x_1)$ is finite) to get

$$F_{T_n}(u) - F_{T_n}(l) \leq \frac{6\rho_{a,b}(x_1)}{\sigma^3_{a,b}(x_1)\sqrt{n}} + \int_l^u \phi(t)dt. \quad (7)$$

Using the last relation, we can finally prove the following bound on probability of manipulation :

$$\frac{6\rho_{a,b}(x_1)}{\sigma^3_{a,b}(x_1)\sqrt{n}} + \int_{\frac{s_m - s_1 - n\mu_{a,b}(x_1)}{\sigma_{a,b}(x_1)\sqrt{n}}}^{\frac{\Delta_{a,b} - n\mu_{a,b}(x_1)}{\sigma_{a,b}(x_1)\sqrt{n}}} \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt$$

Since $e^{-t^2/2} \leq 1$, the last integral can be bounded above by $\frac{1}{\sqrt{2\pi}} \frac{\Delta_{a,b} + s_1 - s_m}{\sigma_{a,b}(x_1)\sqrt{n}}$. This proves that the probability of manipulation is $O\left(1/\sqrt{n}\right)$. □

4 Rank-Order Models

Now we present some examples of belief systems that are c.i.i.d. and positively correlated. Our main source of examples is the rank-order models from machine learning.

Mallows Model

The Mallows [1957] model was originally defined with respect to a fixed (latent) preference ordering σ and a *dispersion parameter* $r \in (0, 1]$. Let π be any preference ordering, then the Mallows model specifies $\Pr\left[\pi|\sigma, r\right] \propto r^{d(\pi,\sigma)}$, i.e. the probability of observing the ordering π decays exponentially with its Kendall-Tau distance from σ. We are not aware of any adaptation of the Mallows model for a conditional belief system. However any such adaptation should capture the idea that once a voter i observes a preference ordering P_i, then she believes that the preference ordering of any other agent is a noisy version of P_i, with its probability decaying exponentially with the Kendall-Tau distance from P_i. This motivates the following belief system:

Definition 4.1. A belief system $\{\mu_i\}_{i=1}^n$ is a *Conditional Mallows model* if $\exists r \in (0, 1]$ such that

$$\mu_i(P_{-i}|P_i) \propto r^{\sum_{j\neq i} d(P_j, P_i)}. \quad (8)$$

Theorem 6. *Every Conditional Mallows model is c.i.i.d. and KT-correlated.*

The proof is trivial since if $\{\mu_i\}_{i=1}^n$ is a conditional belief system then $\nu(P_j|P_i) \propto r^{d(P_j,P_i)}$. As a consequence of Theorem 6, if $\{\mu_i\}_{i=1}^n$ is a conditional Mallows model, then any scoring rule becomes EPIC at a rate proportional to $O\left(e^{-cn}\right)$ where c is as defined in Equation (5).

Plackett-Luce Model

The Plackett-Luce model [Plackett, 1975; Luce, 1959] on a set of alternatives $\{1, \ldots, m\}$ has m parameters: $\gamma_j > 0$ for each alternative j, such that $\sum_{j=1}^m \gamma_j = 1$. For a permutation π of the m alternatives, let $\pi[k]$ denote the alternative placed at position k. The probability of permutation π is given as:

$$\Pr\left[\pi|\{\gamma_j\}_{j=1}^m\right] = \frac{\gamma_{\pi[1]}}{\gamma_{\pi[1]} + \ldots + \gamma_{\pi[m]}} \cdot \frac{\gamma_{\pi[2]}}{\gamma_{\pi[2]} + \ldots + \gamma_{\pi[m]}} \cdots \frac{\gamma_{\pi[m-1]}}{\gamma_{\pi[m-1]} + \gamma_{\pi[m]}} \cdot \frac{\gamma_{\pi[m]}}{\gamma_{\pi[m]}}. \quad (9)$$

We propose the following belief system $\{\mu_i\}_{i=1}^n$ based on the Plackett-Luce Model.

Definition 4.2. A belief system $\{\mu_i\}_{i=1}^n$ is a *Conditional Plackett-Luce model* if there exists m parameters $\gamma_1 > \gamma_2 > \ldots > \gamma_m$ and $\sum_{l=1}^m \gamma_l = 1$ such that $\mu_i(P_{-i}|P_i) = \prod_{j\neq i} \nu(P_j|P_i)$, where

$$\nu(P_j|P_i) = \Pr\left[P_j|\{\gamma_{P_i[l]}\}_{l=1}^m\right]. \quad (10)$$

To compute $\nu(P_j|P_i)$, we first permute the m parameters $\{\gamma_l\}_{l=1}^m$ according to P_i so that $\gamma_{P_i[1]} > \gamma_{P_i[2]} > \ldots > \gamma_{P_i[m]}$, and then use the standard Plackett-Luce model.

The Conditional Plackett-Luce model is not KT-correlated. Suppose there are two voters ($n = 2$) and four alternatives ($m = 4$). We will write 1234 to denote the preference ordering $1\ P_1\ 2\ P_1\ 3\ P_1\ 4$. We now construct a set of parameters $\{\gamma_l\}_{l=1}^4$ such that $\mu_1(1342 \mid 1234) > \mu_1(2134 \mid 1234)$. Since $d(1342, 1234) = 2 > 1 = d(2134, 1234)$, $\mu_1(\cdot)$ is not KT-correlated. Let $\gamma_1 = 1 - 6\varepsilon, \gamma_2 = 3\varepsilon, \gamma_3 = 2\varepsilon, \gamma_4 = \varepsilon$. For, $\gamma_1 > \gamma_2$, we require $1 - 6\varepsilon > 3\varepsilon$ or $\varepsilon < 1/9$. Now,

$$\mu_1(1342 \mid 1234) > \mu_1(2134 \mid 1234)$$
$$\Leftrightarrow \gamma_1 \frac{\gamma_3}{\gamma_3 + \gamma_4 + \gamma_2} \frac{\gamma_4}{\gamma_4 + \gamma_2} > \gamma_2 \frac{\gamma_1}{\gamma_1 + \gamma_3 + \gamma_4} \frac{\gamma_3}{\gamma_3 + \gamma_4}$$
$$\Leftrightarrow 1/12 > 2\varepsilon/(1 - 3\varepsilon) \Leftrightarrow \varepsilon < 1/27.$$

Therefore, as long as $\varepsilon < 1/27$, the belief system fails to be KT-correlated. However, we now prove that the Conditional PL model is TS-correlated. We use the following lemma about the likelihoods in the Conditional PL model.

Lemma 7. *For all k, all $1 \leq k \leq m$, such that $l_1, \ldots, l_k$ distinct,* $\sum_{\pi:\pi[1]=l_1,\ldots,\pi[k]=l_k} \Pr[\pi|\{\gamma_l\}_{l=1}^m] = \gamma_{l_1} \frac{\gamma_{l_2}}{1-\gamma_{l_1}} \cdots \frac{\gamma_{l_k}}{1-\gamma_{l_1}-\ldots\gamma_{l_{k-1}}}$.

See Hunter [2004] for an outline of the proof.

Theorem 8. *Every Conditional PL Model is TS-correlated.*

Proof. Let $\{\mu_i\}_{i=1}^n$ be a Conditional PL belief system with parameters $\{\gamma_l\}_{l=1}^m$ where $\gamma_1 > \gamma_2 > \ldots > \gamma_m > 0$ and $\sum_{l=1}^m \gamma_l = 1$. For any set T of size k $(1 \leq k \leq m-1)$,

$$\sum_{\substack{\{P_{-i}\ :\ B_k(P_j)=T \\ \forall j \neq i\}}} \mu_i(P_{-i}|P_i) = \left(\sum_{\{P:B_k(P)=T\}} \nu(P|P_i) \right)^{n-1}.$$

This follows from two observations : (1) $\{\mu_i\}_{i=1}^n$ is c.i.i.d., so $\mu_i(P_{-i}|P_i) = \prod_{j\neq i} \nu(P_j|P_i)$, and (2) Using the multinomial expansion, $\left(\sum_{\{P:B_k(P)=T\}} \nu(P|P_i)\right)^{n-1}$ can be written as

$$\sum_{\forall j \neq i\ P_j \in \{P:B_k(P)=T\}} \prod_{j\neq i} \nu(P_j|P_i) \tag{11}$$

Now, from the definition of a TS-correlated belief system (2), it is enough to show that $\sum_{P:B_k(P)=B_k(P_i)} \nu(P|P_i) > \sum_{P:B_k(P)=T} \nu(P|P_i)$. Without loss of generality, we can assume that P_i places i-th alternative at position i. Let $T = \{t_1, \ldots, t_k\}$ such that $\gamma_{t_1} > \gamma_{t_2} > \ldots > \gamma_{t_k}$. We will write $\mathcal{S}([k])$ to denote the set of all permutations over the set $\{1, \ldots, k\}$. We have:

$$\sum_{P:B_k(P)=B_k(P_i)} \nu(P|P_i) > \sum_{P:B_k(P)=T} \nu(P|P_i)$$

$$\Leftrightarrow \sum_{\pi\in\mathcal{S}([k])} \sum_{P:P[j]=\pi[j],\ 1\leq j\leq k} \Pr[P|\{\gamma_l\}_{l=1}^m] >$$

$$\sum_{\pi\in\mathcal{S}([k])} \sum_{P:P[j]=t_{\pi[j]},\ 1\leq j\leq k} \Pr[P|\{\gamma_l\}_{l=1}^m]$$

{using Lemma 7}

$$\Leftrightarrow \sum_{\pi\in\mathcal{S}([k])} \left\{ \gamma_{\pi[1]} \frac{\gamma_{\pi[2]}}{1-\gamma_{\pi[1]}} \cdots \frac{\gamma_{\pi[k]}}{1-\gamma_{\pi[1]}-\ldots\gamma_{\pi[k-1]}} \right.$$

$$\left. - \gamma_{t_{\pi[1]}} \frac{\gamma_{t_{\pi[2]}}}{1-\gamma_{t_{\pi[1]}}} \cdots \frac{\gamma_{t_{\pi[k]}}}{1-\gamma_{t_{\pi[1]}}-\ldots\gamma_{t_{\pi[k-1]}}} \right\} > 0 \tag{12}$$

Since T is different from $B_k(P)$, the set of top-k alternatives in preference ordering P, we have for any $\pi \in \mathcal{S}([k])$, $\gamma_{\pi[j]} \geq \gamma_{t_{\pi[j]}}$ for $j = 1, \ldots, k$ and $\exists t\ \gamma_{\pi[t]} > \gamma_{t_{\pi[t]}}$. This implies that for all $l \geq t$

$$\frac{\gamma_{\pi[l]}}{1-\gamma_{\pi[1]}-\ldots-\gamma_{\pi[l-1]}} > \frac{\gamma_{t_{\pi[l]}}}{1-\gamma_{t_{\pi[1]}}-\ldots-\gamma_{t_{\pi[l-1]}}},$$

which is sufficient to guarantee (12). □

Since the Conditional PL model is not KT-correlated, we cannot use Theorem 1 to claim exponential EPIC convergence of any scoring rule under this model. However, the next theorem shows that the convergence is indeed exponential, and that this is true for any scoring rule.

Theorem 9. *The probability that any scoring rule is not EPIC w.r.t. a Conditional Plackett-Luce model goes to zero at a rate proportional to $O(e^{-cn})$, for some constant $c > 0$.*

Proof. Let the belief system $\{\mu_i\}_{i=1}^n$ be a Conditional Plackett-Luce model with parameters $\{\gamma_l\}_{l=1}^m$ where $\gamma_1 > \gamma_2 > \ldots > \gamma_m > 0$ and $\sum_{l=1}^m \gamma_l = 1$.

Consider a voter i with preference ordering P_i and two alternatives a and b such that aP_ib. As we have argued in Theorem 1, it is sufficient to show that $\mu_{a,b}(P_i) > 0$. Without loss of generality we can assume that P_i places the i-th alternative at position i. We have:

$$\mu_{a,b}(P_i) = \mathrm{E}_{X_j\sim\nu(\cdot|P_i)}[sc(a,X_j) - sc(b,X_j)]$$
$$= \sum_{l_1<l_2} (s_{l_1} - s_{l_2}) \{\Pr[X_j[l_1]=a, X_j[l_2]=b|\{\gamma_t\}_{t=1}^m]$$
$$-\Pr[X_j[l_1]=b, X_j[l_2]=a|\{\gamma_t\}_{t=1}^m]\}$$

Therefore, it is enough to show that for any $l_2 > l_1$, $\Pr[X_j[l_1]=a, X_j[l_2]=b|\{\gamma_l\}_{l=1}^m] > \Pr[X_j[l_1]=b, X_j[l_2]=a|\{\gamma_l\}_{l=1}^m]$. Now, suppose alternatives a and b are placed respectively at positions r and s of P_i. Therefore, $\gamma_r > \gamma_s$. Then by Lemma 7 we have,

$$\Pr[X_j[l_1]=a, X_j[l_2]=b|\{\gamma_l\}_{l=1}^m]$$
$$= \sum_{i_1,\ldots,i_{l_2-2}\in[m]\setminus\{a,b\}} \sum_{\sigma\in\mathcal{S}(\{i_1,\ldots,i_{l_2-2}\})}$$

$f(\gamma_{\sigma[i_1]}, \ldots, \gamma_{\sigma[i_{l_1-1}]}, \gamma_r, \gamma_{\sigma[i_{l_1}]}, \ldots, \gamma_{\sigma[i_{l_2-2}]}, \gamma_s)$ where

$$f(\lambda_1, \ldots, \lambda_k) = \lambda_1 \frac{\lambda_2}{1-\lambda_1} \cdots \frac{\lambda_k}{1-\lambda_1-\ldots-\lambda_{k-1}}.$$

It is easy to see that $\gamma_r > \gamma_s$ implies $f(\gamma_{\sigma[i_1]}, \ldots, \gamma_{\sigma[i_{l_1-1}]}, \gamma_r, \gamma_{\sigma[i_{l_1}]}, \ldots, \gamma_{\sigma[i_{l_2-2}]}, \gamma_s) > f(\gamma_{\sigma[i_1]}, \ldots, \gamma_{\sigma[i_{l_1-1}]}, \gamma_s, \gamma_{\sigma[i_{l_1}]}, \ldots, \gamma_{\sigma[i_{l_2-2}]}, \gamma_r)$, as $\gamma_r\gamma_s$ cancels out from both sides, and for all j such that $l_1 \leq j \leq l_2 - 2$, we have $(1-\gamma_{\sigma[i_1]}-\ldots-\gamma_{\sigma[i_j]}-\gamma_r)^{-1} > (1-\gamma_{\sigma[i_1]}-\ldots-\gamma_{\sigma[i_j]}-\gamma_s)^{-1}$. □

5 Conclusion

We have provided a positive result for incentive-aligned social choice in the presence of correlation between voter preferences. When the beliefs of agents are positively correlated in terms of Kendall-Tau distance, we show that all scoring rules become EPIC exponentially quickly in the number of voters. We have instantiated the general framework to conditional variations on the popular Mallows and Plackett-Luce models for rank aggregation. One question for future work is whether exponential, EPIC convergence holds for correlated distributions that are not conditionally independent and identical. We also leave open the question of closing the gap between exponential convergence and $1/\sqrt{n}$ for the class of TS-correlated beliefs.

6 Acknowledgements

This work is supported in part by NSF grant CCF-1301976 and a grant from the FLI. We also thank Yang Liu for technical feedback on an earlier version.

References

[Altman and Tennenholtz, 2010] Alon Altman and Moshe Tennenholtz. An Axiomatic Approach to Personalized Ranking Systems. *J. ACM*, 57(4), 2010.

[Baharad and Neeman, 2002] Eyal Baharad and Zvika Neeman. The Asymptotic Strategyproofness of Scoring and Condorcet Consistent Rules. *Review of Economic Design*, 7(3):331–340, 2002.

[Barberà, 2011] Salvador Barberà. Strategyproof Social Choice. *Handbook of social Choice and Welfare*, 2:731–831, 2011.

[Bhargava *et al.*, 2011] Mohit Bhargava, Dipjyoti Majumdar, and Arunava Sen. Incentive-Compatible Voting Rules with Positively Correlated Beliefs. *Available at* `http://www.isid.ac.in/~asen/papers/ms.pdf/`*, version : April 10*, 2011.

[Conitzer and Walsh, 2016] Vincent Conitzer and Toby Walsh. Barriers to Manipulation in Voting. *Handbook of Computational Social Choice, Forthcoming*, 2016.

[Durrett, 2010] Rick Durrett. *Probability: Theory and Examples*. Cambridge university press, 2010.

[Faliszewski and Procaccia, 2010] Piotr Faliszewski and Ariel D Procaccia. AI's War on Manipulation: Are We Winning? *AI Magazine*, 31(4):53–64, 2010.

[Gibbard, 1973] Allan Gibbard. Manipulation of Voting Schemes: A General Result. *Econometrica*, pages 587–601, 1973.

[Hoeffding, 1963] Wassily Hoeffding. Probability Inequalities for Sums of Bounded Random Variables. *Journal of the American statistical association*, 58(301):13–30, 1963.

[Hunter, 2004] David R Hunter. MM Algorithms for Generalized Bradley-Terry Models. *Annals of Statistics*, pages 384–406, 2004.

[Lang and Xia, 2009] Jérôme Lang and Lirong Xia. Sequential Composition of Voting Rules in Multi-Issue Domains. *Mathematical Social Sciences*, 57(3):304–324, 2009.

[Luce, 1959] R.. Ducan Luce. *Individual Choice Behavior : A Theoretical Analysis*. John Wiley and sons, 1959.

[Majumdar and Sen, 2004] Dipjyoti Majumdar and Arunava Sen. Ordinally Bayesian Incentive Compatible Voting Rules. *Econometrica*, 72(2):523–540, 2004.

[Mallows, 1957] Colin L Mallows. Non-Null Ranking Models. *Biometrika*, pages 114–130, 1957.

[Mao *et al.*, 2013] Andrew Mao, Ariel D. Procaccia, and Yiling Chen. Better Human Computation Through Principled Voting. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence, July 14-18, 2013, Bellevue, Washington, USA.*, 2013.

[Marden, 1996] John I Marden. *Analyzing and Modeling Rank Data*. CRC Press, 1996.

[Parkes *et al.*, 2015] David C. Parkes, Ariel D. Procaccia, and Nisarg Shah. Beyond Dominant Resource Fairness: Extensions, Limitations, and Indivisibilities. *ACM Transactions on Economics and Computation*, 3(1):3, 2015.

[Pennock *et al.*, 2000] 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 Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, pages 729–734. AAAI Press, 2000.

[Plackett, 1975] Robin L Plackett. The Analysis of Permutations. *Applied Statistics*, pages 193–202, 1975.

[Satterthwaite, 1975] Mark Allen Satterthwaite. Strategy-Proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *Journal of Economic Theory*, 10(2):187–217, 1975.

[Sen *et al.*, 2015] Arunava Sen, Mohit Bhargava, and Dipjyoti Majumdar. Incentive-Compatible Voting Rules with Positively Correlated Beliefs. *Theoretical Economics*, 10(3):867–885, 2015.

[Slinko, 2002] Arkadii Slinko. On Asymptotic Strategy-Proofness of Classical Social Choice Rules. *Theory and Decision*, 52(4):389–398, 2002.

[Young, 1975] H Peyton Young. Social Choice Scoring Functions. *SIAM Journal on Applied Mathematics*, 28(4):824–838, 1975.