

Objective Social Choice: Using Auxiliary Information to Improve Voting Outcomes

Silviu Pitis

University of Toronto, Vector Institute
Toronto, Ontario, Canada
spitis@cs.toronto.edu

Michael R. Zhang

University of Toronto, Vector Institute
Toronto, Ontario, Canada
michael@cs.toronto.edu

ABSTRACT

How should one combine noisy information from diverse sources to make an inference about an objective ground truth? This frequently recurring, normative question lies at the core of statistics, machine learning, policy-making, and everyday life. It has been called “combining forecasts”, “meta-analysis”, “ensembling”, and the “MLE approach to voting”, among other names. Past studies typically assume that noisy votes are identically and independently distributed (i.i.d.), but this assumption is often unrealistic. Instead, we assume that votes are independent but not necessarily identically distributed and that our ensembling algorithm has access to certain auxiliary information related to the underlying model governing the noise in each vote. In our present work, we: (1) define our problem and argue that it reflects common and socially relevant real world scenarios, (2) propose a multi-arm bandit noise model and count-based auxiliary information set, (3) derive maximum likelihood aggregation rules for ranked and cardinal votes under our noise model, (4) propose, alternatively, to learn an aggregation rule using an order-invariant neural network, and (5) empirically compare our rules to common voting rules and naive experience-weighted modifications. We find that our rules successfully use auxiliary information to outperform the naive baselines.*

KEYWORDS

objective social choice; ensemble methods; combining forecasts

ACM Reference Format:

Silviu Pitis and Michael R. Zhang. 2020. Objective Social Choice: Using Auxiliary Information to Improve Voting Outcomes. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 8 pages.

1 INTRODUCTION

Many collective decision making processes aggregate noisy good faith opinions in order to make an inference about some underlying ground truth. In cooperative policy making, for example, each party advocates for the policy they believe is objectively best. Similarly, in academic peer review, a meta-reviewer combines good faith reviewer opinions about a submitted paper. Other examples are easy to come by. We refer to this setting as *objective* social choice, to contrast it with the typical *subjective* social choice setting [35], where the optimal choice is defined in terms of the voter utilities

*Code available at https://github.com/spitis/objective_social_choice

rather than a ground truth. Whereas subjective social choice can be viewed as *collective compromise*, objective social choice can be viewed as *collective estimation*.

Unlike the subjective setting, where it is natural to consider each source or voter equally—an axiom known as “anonymity” [32]—objective analysis suggests otherwise: diverse and more informed opinions should be valued more. Many sensible, real-world settings involve asymmetric (non-anonymous) voting, making this a relevant line of analysis. Academic review is one. Another is corporate governance, where different stakeholder classes have varying voting powers, depending on the issue. In such cases, varying voter weights are natural, and one can evaluate the quality of social choices via other avenues (e.g., direct evaluation [27] or *ex post* analysis [19]). In other settings, such as national elections, the objective approach raises ethical concerns of fairness, and the objective approach may be inappropriate.

Although objective social choice has been the subject of numerous studies in social choice [6, 10, 12, 44], forecasting [3, 9, 13], statistics [16, 18] and machine learning [15, 37] (Section 2), to our knowledge, no prior work has dealt with the case of non-i.i.d. ordinal feedback (i.e., ranked preferences). Yet this is the case in many practical applications. During peer review, for instance, two of three reviewers might share primary areas of expertise, but being human, cannot share comparable cardinal estimates. Or consider a robot that must aggregate feedback from human principals. Once again, the different principals will draw upon diverse background to form their opinions, which can only be shared as ordinal preferences. In each case, how should the non-i.i.d. feedback be aggregated?

Our work is intended as a first step toward answering this question. To narrow the scope of our inquiry, we make several modeling assumptions (Section 3), which we hope can be relaxed in future work. In particular, we assume that (1) the underlying ground truth and noise generating process is modeled as a k -armed bandit problem, where the different arms represent different alternatives, (2) different voters see different pulls of the arms, and (3) the social decision rule sees how many pulls each voter saw (but not their outcomes). We solve for the maximum likelihood social choice in a series of cases (Section 4). As our derived rules rest on strong assumptions about the noise generation process, we also propose to learn a more flexible aggregation rule using an order-invariant neural network (Section 4.2). We empirically compare our derived and learned rules to classical voting rules (Section 5). Our results confirm the intuition that objective estimation can be improved by up-weighting opinions from diverse and more informed sources.

2 RELATED WORK

Social Choice. The fundamental question of social choice asks: how should we combine the preferences of many into a social preference? [2, 38]. While the usual approach evaluates the social preference in terms of the preferences of individuals [1, 22, 35] (“subjective” social choice), a line of papers frame individual preferences as noisy reflections of an underlying ground truth and evaluate the social preference by comparing to the ground truth (“objective” social choice). Perhaps the first is due to Condorcet [10], who studied the problem where voters rank two alternatives correctly with some probability $p > \frac{1}{2}$. This simple noise model is a special case of the n -alternative Mallows model [31], according to which each voter ranks each pair of alternatives correctly with probability $p > \frac{1}{2}$ (and votes are redrawn if a cycle forms). Young [44] generalized Condorcet’s analysis to general Mallows noise ($n > 2$ alternatives), and showed that the Kemeny voting rule returns the maximum likelihood (MLE) estimate of the truth for this noise model. Conitzer and Sandholm [12] further extended this “maximum likelihood” analysis to other voting rules and i.i.d. noise models; their main results include a proof that *any* so-called “scoring rule” (e.g., plurality, Borda count, veto) is the MLE estimator for some i.i.d. noise model, as well as proofs that certain other voting rules (e.g., Copeland) are not MLE estimators for *any* i.i.d. noise model. Caragiannis et al. [5, 6] consider the sample complexity necessary to ensure high likelihood reconstructions of the ground truth under Mallow’s noise. Also related is the independent conversations model in social networks, where independent pairs of voters receive information about some ground truth. Conitzer [11] introduces the model for two alternatives and constructs the maximum likelihood estimator, which he shows to be #P-hard. Procaccia et al. [36] extend and analyze this model for multiple alternatives.

Our work is unique in two respects. First, we do not assume i.i.d. noise. Rather, we use a k -armed bandit noise model, which provides a basic but plausible noise generating process that can account for diversity in the subjective experiences of voters. Second, our approach is cardinal: rather than apply noise directly to ranked preferences, we apply noise to a cardinal ground truth. As Procaccia and Rosenschein [35] introduced cardinal analysis into subjective social choice, we do so in the objective case.

Forecasting, Statistics and Machine Learning. Numerous papers in forecasting and statistics have examined the combination of estimates. Bates and Granger [3] provided an early derivation of optimal weights for linearly combining two cardinal estimates. Their analysis was extended to the n estimate case by Dickinson [13, 14] and improved by Granger and Ramanathan [21], among many others [9, 20, 41]. While the literature on combining forecasts typically deals with point estimates (or time series thereof), significant work has also been done on combining probability distributions [17, 18, 26, 33]. In empirical statistics, the combination of experimental results is known as meta-analysis [16]. Almost all work on combining cardinal estimates considers linear combinations; this can be justified by an appeal to Harsanyi’s theorem [22, 42], which states (roughly) that any cardinal—in the VNM expected utility sense [40]—combination of cardinal estimates that satisfies Pareto indifference (i.e., the combination is a function of the estimates

and nothing else) can be expressed as a linear combination of the estimates.

Combining estimators through ensembles is a common technique used to improve inference performance in machine learning [15, 37]. Much like the literature on combining forecasts, Perrone and Cooper [34] and Tresp and Taniguchi [39] propose weighting schemes that ensemble estimators based on their variances.

Our proposal differs from the above works in that (1) we combine ordinal votes rather than cardinal predictions, and (2) we define an underlying noise model and use count-based information rather than empirical variances. Some recent works in reinforcement learning use ensembles in an ordinal setting [7, 8], but these works use naive ensembling techniques (majority vote and arithmetic mean).

The dueling bandit problem setting is similar to ours, in that ordinal comparisons are used to make an inference about an underlying, (potentially) cardinal bandit [45]. As it uses repeat online comparisons rather than count (or other similarity) information, the dueling bandits formulation is more suitable to interactive and online applications such as ad placement and recommender systems than one-shot votes. Our work could potentially be applied to initialize an online bandit when historical information is available.

3 MODEL

We first present a generic framework for objective social choice and then describe the modeling assumptions we make for our work.

3.1 Formal Setup

We assume the existence of a ground truth, cardinal objective function $V : \mathcal{A} \rightarrow \mathbb{R}$, where \mathcal{A} is a finite set of alternatives, and define $n \triangleq |\mathcal{A}|$. We represent V by the vector $[\mu_1, \mu_2, \dots, \mu_n]$, where μ_i is the “true quality” of alternative a_i , and denote the optimal alternative by $a^* \triangleq \arg \max_i \{\mu_i\}$. m voters partially observe this ground truth and provide our social choice rule f with their noisy votes. Each such set of noisy votes is an element of the voting or observation space \mathcal{X} , which can be seen as (part of) the input domain of f . In general, there are many ways in which voters could make their observations and provide their feedback. Regardless of the precise details, it seems plain that a rule with access to the votes, but to no other information ($\text{dom } f = \mathcal{X}$) should satisfy anonymity i.e., weigh each vote equally. It is also plain that for an anonymous voting rule, whether or not votes are i.i.d. is irrelevant. Therefore, our setting is only interesting when, in addition to votes, our voting rule has access to some auxiliary information or context $c \in C$, so that $\text{dom } f = \mathcal{X} \times C$. As is the case for \mathcal{X} , there are many options one could consider for C , and we make specific assumptions below.

The codomain of f may either be (1) \mathcal{V} , the set of valid ground truth functions (with $V \in \mathcal{V}$), so that f outputs cardinal prediction \hat{V} , (2) \mathcal{A} , so that f outputs a single best alternative \hat{a} , or (3) the set of ordinal rankings over the alternatives. Note that if the codomain is \mathcal{V} , one can consider this entire process as a sort of autoencoder: there is a noise model $g : \mathcal{V} \rightarrow \mathcal{X} \times C$ that produces the votes and auxiliary information, and the job of our rule $f : \mathcal{X} \times C \rightarrow \mathcal{V}$ is (roughly speaking) to reconstruct the input to g . Thus, optimal rules are closely tied to noise models; cf. [12]. Figure 1 summarizes the objective social choice framework.




Figure 1: A generic framework for objective social choice. The ground truth $V \in \mathcal{V}$ passes through noise model g to generate the votes $x \in \mathcal{X}$ and contexts $c \in \mathcal{C}$ for m voters. The rule f is applied to generate social choice $f(x, c)$.

In addition to specifying the noise model g , voting format \mathcal{X} , auxiliary information \mathcal{C} , and codomain of f , we must also specify an objective function: what makes a given rule or rule selection algorithm “good”? As usual, the answer will depend on the context. In our present work we seek the rule that corresponds to a maximum likelihood estimate (MLE) \hat{V} of the ground truth V ; that is, given data $(x, c) \in \mathcal{X} \times \mathcal{C}$, the output of f is consistent with the \hat{V} that is most likely to have generated votes x conditioned on the context c . It should be noted that where f returns a best alternative or ranking over alternatives, this problem formulation is different from finding the most likely best alternative or most likely ranking over alternatives, as done under ordinal (Mallows) noise [12, 44]—in our setting, the maximum likelihood alternative and ranking depends on a distributional estimate of \hat{V} . The MLE rule may not be the empirically best rule, and so to compare voting rules in Subsection 5, we will use the notion of *regret*, defined as $V(a^*) - V(\hat{a})$, where \hat{a} is the alternative most preferred by f .

3.2 Specific Modeling Assumptions

To narrow the scope of our present inquiry, we make the following assumptions about g and \mathcal{C} :

ASSUMPTION 1. *The voters observe the n -dimensional ground truth $V = [\mu_1, \mu_2, \dots, \mu_n]$ through an n -arm stochastic bandit [29]. Each arm reveals information about the corresponding dimension of V , and voters observe samples from arm i according to $r_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$, where σ_i^2 is the variance of arm i . To simplify analysis, we assume that the σ_i^2 are either known or equal.*

ASSUMPTION 2. *There are m voters, where the i -th voter sees the j -th arm pulled c_{ij} times. Let $\bar{c}_j = \sum_i c_{ij}$. Each voter sees different (independently sampled) pulls—thus, vote noise is independent, but not identically distributed.*

ASSUMPTION 3. *The auxiliary information $c \in \mathcal{C}$ consists of the observation counts for each voter. For each voter i , this is number of pulls for each arm: $[c_{i1}, c_{i2}, \dots, c_{in}]$.*

ASSUMPTION 4. *Voter i estimates V as $X_i = [x_{i1}, x_{i2}, \dots, x_{in}]$, which determines their vote (specific details below).*

The above assumptions leave open the voting format \mathcal{X} , and also the output (codomain) of our voting rule f . We explore different combinations of these in Section 4 below.

Although there are many alternatives to Assumptions 1-4, these basic assumptions strike us as a simple, yet flexible model. Many noise processes can be framed as bandits. Take peer review for instance: one could designate an arm for each paper under review (and accept the top k). Similarly, count information, which serves as a proxy for voter experience, provides a generic way of characterizing the “non-i.i.d.ness” of votes (an extension to our work might examine the case where some voters observe the *same* pulls, leading to dependent votes). Other interesting choices may include voter similarities, as specified by some kernel function (this would model the votes as a sample from a Gaussian process), or empirical covariance measurements (obtained by observing several votes). The assumption of Gaussian noise is relaxed in Subsection 4.2 and our experiments.

4 AGGREGATION RULES

4.1 Derived Rules

In this section we analyze MLE social choice under the specific modeling assumptions made above. We do this in a series of five cases of roughly increasing complexity where, in each case, we derive one or more *scoring rules* [12]. Scoring rules, such as the Borda, plurality and veto rules [4], compute for each alternative j a single aggregate score (or predicted utility, \hat{V}_j) by taking a simple sum across individual voter weights (i.e., $\hat{V}_j = \sum_i w_{ij}$, where the w_{ij} is the weight of voter i ’s vote for arm j). The alternatives are ranked according to these numbers and the top scoring alternative is selected. For example, the commonly used plurality rule assigns weight $w_{ij} = 1$ to voter i ’s top choice j , and $w_{ik} = 0$ for $k \neq j$, which results in selecting the alternative that is ranked first most often. In the two alternative cases below (cases 2 and 3), where the derived weights w_i do not have a j subscript, voter i ’s top choice gets weight w_i and their second choice gets weight $-w_i$ (or 0, since only relative weight matters).

The first two cases below, which use cardinal votes, simply recast known results into our setting. The latter three use ordinal votes and are novel contributions.

CASE 1 (MANY ALTERNATIVES, VOTES ARE CARDINAL MEANS). *There are n arms, and voter i provides their cardinal votes $\{x_{ij}\}$ for each arm j , where x_{ij} is the mean of i ’s observations for arm j .*

SOLUTION. Had our aggregation rule seen the pulls itself, its MLE estimate $\hat{\mu}_j$ of the true mean μ_j would be the mean observed reward, which can be computed directly from the available information:

$$\begin{aligned} \hat{\mu}_j &= (r_{1j} + r_{2j} + \dots + r_{mj}) / \bar{c}_j \\ &= (c_{1j}x_{1j} + c_{2j}x_{2j} + \dots + c_{mj}x_{mj}) / (\sum_i c_{ij}), \end{aligned}$$

so that $w_{ij} \propto c_{ij}$. □

Note that $x_{ij} \sim \mathcal{N}(\mu_j, \sigma_j^2 / c_{ij})$, so that each estimate is weighted *inversely proportional* to its variance, σ_j^2 / c_{ij} . The use of inverse variance to weight independent cardinal estimates is well known [3, 13, 16, 34, 39].

CASE 2 (2 ALTERNATIVES, VOTES ARE CARDINAL DIFFERENCES). *There are $n = 2$ arms, and each voter i provides their estimate $y_i = c_{i2} - c_{i1}$ of the cardinal difference between arms.*

SOLUTION. As x_{i1} and x_{i2} are independent, we have that $x_{i2} - x_{i1} \sim \mathcal{N}(\mu_2 - \mu_1, \frac{\sigma_2^2}{c_{i2}} + \frac{\sigma_1^2}{c_{i1}}) = \mathcal{N}(\mu_2 - \mu_1, \frac{\sigma_2^2 c_{i1} + \sigma_1^2 c_{i2}}{c_{i2} c_{i1}})$. To combine the votes we take the weighted mean with weights proportional to the inverse variances [3], so that $w_i \propto \frac{c_{i2} c_{i1}}{\sigma_2^2 c_{i1} + \sigma_1^2 c_{i2}}$. \square

Unlike Case 1, where the σ_j^2 was irrelevant to w_{ij} , the weights in Case 2 depend on σ_1^2/σ_2^2 . We assumed above that this ratio is known; if not, one might infer the ratio from data. An interesting corollary is that a voter that wishes to maximize the weight of her vote should pull each of the arms equally. If all voters adopt this strategy, we do not need estimates of the variances of the arms and can just weigh each vote in proportion to voter experience.

CASE 3 (2 ALTERNATIVES, VOTES ARE ORDINAL RANKS). *There are 2 arms, and each voter i provides an ordinal ranking (a_1, a_2) indicating that they value a_1 higher than a_2 (i.e., $x_{i1} \geq x_{i2}$).*

SOLUTION. As above, we have $x_{i2} - x_{i1} \sim \mathcal{N}(\mu_2 - \mu_1, \frac{\sigma_2^2 c_{i1} + \sigma_1^2 c_{i2}}{c_{i2} c_{i1}})$. Denoting the CDF of $x_{i2} - x_{i1}$ by Φ_i , and defining the binary variable $Y_i = \mathbb{1}_{x_{i2} - x_{i1} \leq 0}$, we have $Y_i \sim \mathcal{B}(\Phi_i(0))$ (the Bernoulli distribution parameterized by Φ_i evaluated at 0). Our votes $x \in \mathcal{X}$ consist of a set of samples $\{y_1 \sim Y_1, y_2 \sim Y_2, \dots, y_m \sim Y_m\}$. Since adding a constant to the underlying means has no effect on the likelihood, a direct inference about $V = (\mu_1, \mu_2)$ is impossible and we instead seek to estimate the difference $\mu_2 - \mu_1$. Defining, $s_i^2 = \frac{\sigma_2^2 c_{i1} + \sigma_1^2 c_{i2}}{c_{i2} c_{i1}}$, we want to choose $\Delta \triangleq \hat{\mu}_2 - \hat{\mu}_1$ to maximize the log-probability of the data (since Δ which maximizes the log likelihood also maximizes the likelihood):

$$\begin{aligned} \log \Pr(\mathcal{D}; \Delta) &= \sum_i \log \Phi_i(0)^{y_i} (1 - \Phi_i(0))^{1-y_i} \\ &= \sum_i y_i \log \left[\frac{1}{2} + \frac{1}{2} \operatorname{erf} \left(\frac{-\Delta}{s_i \sqrt{2}} \right) \right] + (1 - y_i) \log \left[\frac{1}{2} - \frac{1}{2} \operatorname{erf} \left(\frac{-\Delta}{s_i \sqrt{2}} \right) \right]. \end{aligned}$$

We could try to optimize directly with respect to Δ by setting $\frac{d}{d\Delta} \log \Pr(\mathcal{D}; \Delta) = 0$, but this appears analytically intractable:

$$0 = \sum_i y_i \frac{\frac{-2}{s_i \sqrt{2\pi}} \exp \frac{-\Delta^2}{2s_i^2}}{\left[1 + \operatorname{erf} \left(\frac{-\Delta}{s_i \sqrt{2}} \right) \right]} - (1 - y_i) \frac{\frac{-2}{s_i \sqrt{2\pi}} \exp \frac{-\Delta^2}{2s_i^2}}{\left[1 - \operatorname{erf} \left(\frac{-\Delta}{s_i \sqrt{2}} \right) \right]}.$$

However, since $\log \Pr(\mathcal{D}; \Delta)$ is concave (proof in Appendix), its gradient evaluated at $\Delta = 0$ points in the direction of the MLE solution and we can use this fact to find Y corresponding to MLE estimate of Δ by evaluating (see Appendix for details):

$$\frac{d}{d\Delta} \log \Pr(\mathcal{D}; \Delta)(0) \propto \sum_i [-y_i s_i^{-1} + (1 - y_i) s_i^{-1}],$$

so that $w_i \propto s_i^{-1} = \sqrt{\frac{c_{i2} c_{i1}}{\sigma_2^2 c_{i1} + \sigma_1^2 c_{i2}}}$. \square

CASE 4 (MANY ALTERNATIVES, VOTES ARE ORDINAL RANKS). *There are n arms, and voter i provides an ordinal ranking indicating whether they prefer a_j to a_k (i.e., whether $x_{ij} \geq x_{ik}$) for all pairs (a_j, a_k) .*

APPROXIMATE SOLUTION. Though we were unable to solve this case exactly, we take advantage of a naive independence assumption (a la Naive Bayes [30]) to arrive at a plausible, approximate aggregation rule. We will confirm in Section 5 that it empirically outperforms the baselines. As above, we define binary variable

$Y_{i,j < k} \triangleq \mathbb{1}_{x_{ij} - x_{ik} \leq 0}$ (indicating that k is preferred to j), so that votes are a set of samples $\{y_{i,j < k} \sim Y_{i,j < k}\}$, and assume:

ASSUMPTION (NAIVE INDEPENDENCE). *For all i and distinct pairs (j, k) and (s, t) , variables $Y_{i,j < k}$ and $Y_{i,s < t}$ are independent.*

This assumption is *never* true for $n > 2$. To see this, consider the alternatives j, k, ℓ , and note that $y_{i,j < k} = 1$ and $y_{i,k < \ell} = 1$ imply $y_{i,j < \ell} = 1$ (by transitivity of the underlying cardinal values), which violates independence. Nevertheless, by using this assumption, we can apply our Case 3 strategy by rewriting the probability of the data as a sum over the probabilities of the pairwise votes:

$$\log \Pr(\mathcal{D}; \hat{V}) = \sum_{j \neq k} \log \Pr(\mathcal{D}_{jk}; \Delta_{jk})$$

where $\hat{V} = [\hat{\mu}_1, \hat{\mu}_2, \dots, \hat{\mu}_n]$, $\Delta_{jk} \triangleq \hat{\mu}_j - \hat{\mu}_k$ as above, and $\Pr(\mathcal{D}_{jk}; \Delta_{jk})$ is the probability of observing the voters' pairwise comparisons between a_j and a_k given Δ_{jk} (ignoring other alternatives). Noting that $\frac{\partial \Delta_{jk}}{\partial \hat{\mu}_j} = 1$ and $\frac{\partial \Delta_{jk}}{\partial \hat{\mu}_k} = -1$, we can apply our Case 3 solution to find the partial derivatives of $\log \Pr(\mathcal{D}_{jk}; \hat{V})$, evaluated at $\hat{V} = 0$, with respect to $\hat{\mu}_j$ and $\hat{\mu}_k$. Summing across alternative pairs yields:

$$w_{ij} \propto \sum_{k \neq j} [-y_{i,j < k} + (1 - y_{i,j < k})] \sqrt{\frac{c_{ij} c_{ik}}{\sigma_k^2 c_{ij} + \sigma_j^2 c_{ik}}}. \quad (1)$$

The above solution might be improved by examining the following failure mode, which arises on account of the Naive Independence assumption. If the best alternative j is observed significantly less often than the second best alternative k —i.e., $\sum_i c_{ij} < \sum_i c_{ik}$ —the second best alternative will tend to receive more positive weight, even if all voters report the correct pairwise ordering. For example, in the case of one voter, if that voter reports the correct ordering for three alternatives with counts 1 (for the top alternative), 10, and 10, the above approximate solution will choose the second best alternative. This is obviously a bad outcome. To avoid it, we propose that each alternative's weight be normalized by the total absolute weight it would otherwise received, yielding normalized weights:

$$\bar{w}_{ij} = \frac{w_{ij}}{\sum_i \sum_{k \neq j} \sqrt{\frac{c_{ij} c_{ik}}{\sigma_k^2 c_{ij} + \sigma_j^2 c_{ik}}}},$$

where w_{ij} is defined as above. Our experiments test both the unnormalized (w_{ij}) and normalized (\bar{w}_{ij}) versions of the rule. \square

CASE 5 (MANY ALTERNATIVES, VOTES ARE TOP CHOICE ONLY). *There are n arms, and each voter i provides their top choice a_j indicating that they most prefer a_j (i.e., $x_{ij} \geq x_{ik}, \forall k \neq j$).*

APPROXIMATE SOLUTIONS. Let $\phi_{ij}(x)$ and $\Phi_{ij}(x)$ denote the Gaussian PDF and CDF for x_{ij} . We have that the probability of voter i selecting alternative j is equal to the probability that the largest order statistic of $X_{i,-j}$ (where $X_{i,-j}$ denotes the set $\{x_{ik} | k \neq j\}$) is less than x_{ij} :

$$\Pr(Y_{ij}) = \int_{-\infty}^{\infty} \Pr(x_j = s) \Pr(\max X_{i,-j} \leq s) ds = \mathbb{E}_{x \sim \phi_{ij}} \prod_{k \neq j} \Phi_{ik}(x).$$

The log-likelihood of the data is therefore:

$$\log \Pr(\mathcal{D}; \hat{V}) = \sum_i \log \mathbb{E}_{x \sim \phi_{ij}} \prod_{k \neq j} \Phi_{ik}(x)$$

where the expectation for each voter is taken with respect to that voter’s top choice a_j (we are abusing notation slightly, as a_j differs across voters). While there appears to be no way to maximize this analytically [23], we can compute the gradient with respect to \hat{V} :

$$\begin{aligned}
\nabla \log \Pr(\mathcal{D}; \hat{V}) &= \sum_{i=1}^n \nabla \log \mathbb{E}_{x \sim \phi_{ij}} \prod_{k \neq j} \Phi_{ik}(x) \\
&= \sum_{i=1}^n \frac{1}{f(i)} \nabla \mathbb{E}_{x \sim \phi_{ij}} \prod_{k \neq j} \Phi_{ik}(x) \\
&= \sum_{i=1}^n \frac{1}{f(i)} \int_{-\infty}^{\infty} \frac{\phi_{ij}(x)}{\phi_{ij}(x)} \nabla \left[\phi_{ij}(x) \prod_{k \neq j} \Phi_{ik}(x) \right] dx \\
&= \sum_{i=1}^n \frac{1}{f(i)} \mathbb{E}_{x \sim \phi_{ij}} \left[\begin{aligned} &(\nabla \log \phi_{ij}(x)) \prod_{k \neq j} \Phi_{ik}(x) + \nabla \prod_{k \neq j} \Phi_{ik}(x) \end{aligned} \right] \\
&= \sum_{i=1}^n \frac{1}{f(i)} \mathbb{E}_{x \sim \phi_{ij}} [g(i, x)s(i, x) + \nabla s(i, x)]
\end{aligned}$$

where the third and fourth equalities use the log derivative trick of the REINFORCE [43] gradient estimator together with the product rule, and $f(i)$, $g(i, x)$, and $s(i, x)$ are defined accordingly:

$$f(i) = \mathbb{E}_{x \sim \phi_{ij}} s(i, x) \quad s(i, x) = \prod_{k \neq j} \Phi_{ik}(x) \quad g(i, x) = \nabla \log \phi_{ij}(x)$$

We thus have a method for Monte Carlo estimation of the gradient of the log likelihood. Each term has an intuitive justification. $f(i)$ represents a weight for each voter. A voter whose vote is in line with the current guess of the underlying arm means has less weight on the gradient. $g(i, x)s(i, x)$ is part of the typical REINFORCE [43] objective and corresponds to increasing the probability in regions where the score (product of the CDFs of the remaining arms) is high. Finally $\nabla s(i, x)$ is the correction term that appears due to the dependence of the score function on the arm means. $\nabla s(i, x)$ incentivizes decreasing the means of the arms that are not voted for. Initializing $\hat{V} = 0$, we can either compute the gradient once and take the maximal component to be the winner (as in Case 3, but without the optimality guarantee) or use the gradient ascent algorithm to find an optimum. We will do the former, and call it the Case 5 “Monte Carlo approximation.”

In terms of implementation, $g(i, x)$, $s(i, x)$, and $\nabla s(i, x)$ are straightforward and can be done with a library that computes density functions for Gaussians. In particular, we note that each component of $\nabla s(i, x)$ consists of the product of CDFs and a single PDF.

As computing a good approximation using the Monte Carlo strategy can be expensive and requires known pull variance σ_i , we propose two analytical approximations that only require the ratio σ_i/σ_j to be known. First, noting that events $(x_{ij} \geq x_{ik})$ and $(x_{ij} \geq x_{ih})$ are positively correlated for all j, k, h , we have $\Pr(x_{ij} \geq$

$x_{ik} \mid x_{ij} \geq x_{ih}) \geq \Pr(x_{ij} \geq x_{ik})$, which gives the lower bound:

$$\begin{aligned}
\log \Pr(\mathcal{D}) &= \sum_i \log \prod_{k \neq j} \Pr \left[x_{ij} \geq x_{ik} \mid \bigcup_{1 \leq h < k, h \neq j} (x_{ij} \geq x_{ih}) \right] \\
&\geq \sum_i \sum_{k \neq j} \log \Pr(x_{ij} \geq x_{ik}).
\end{aligned}$$

We can now apply the same argument as in Case 4 and approximately optimize this lower bound by following its gradient at $\hat{V} = 0$. This leads to same weights as our unnormalized Case 4 rule (i.e., weigh votes according to equation 1) for observed comparisons (i.e., all pairs involving each voter’s top choice). We call this the Case 5 “lower bound approximation.”

A second approach makes the following simple observation: at $\hat{V} = 0$, a gradient step in the direction of maximizing $\Pr(x_{ij} \geq 0)$ also increases $\Pr(x_{ij} \geq \max_{i \neq j} X_{i,-j})$. To evaluate the gradient of $\log \Pr(x_{ij} \geq 0)$ at $\hat{V} = 0$, one can run through a computation similar to Case 3, or simply take the limit of the Case 3 weight as one of the counts goes to ∞ . This yields $w_{ij} = \sqrt{c_{ij}}$ for the top choices a_j each arm has equal observation variance (with $w_{ik} = 0$ for $k \neq j$). We call this the Case 5 “zero approximation.”

Both analytical approximations are a bit crude. The lower bound approximation ignores significant dependencies, and the zero approximation doesn’t factor in counts of non-selected alternatives. In both cases we use the gradient at $\hat{V} = 0$, but unlike in Case 3 where this is justified by concavity, there is no similarly strong justification here. Nevertheless, we will see in our experiments that both approximations improve over plurality baselines. \square

4.2 Learning an Aggregation Rule

Can we come up with a rule for the many alternative, ordinal rank case (Case 4) that does not rely on the Naive Independence assumption? Although we were unable to do so analytically, we propose to learn an aggregation rule from data. This rule will serve as a useful baseline for our derived rules, and the approach is flexible, in that it can be trained on data generated by any noise model (e.g., a k -armed bandit with uniform observation noise). As an additional benefit, the learned rule will output a distribution over outcomes (our derived rules output point estimates).

We require our learned rule to apply in the case of an arbitrary number of voters and alternatives. Ideally, our rule should be a function $f : \mathcal{X} \times \mathcal{C} \rightarrow \mathbb{R}^n$ that is *order invariant* with respect to voters, and *order equivariant* with respect to alternatives (permuting the alternatives permutes the results in the same way). Both properties were studied by Zaheer et al.’s work on Deep Sets [46], which investigated the expressiveness of the order invariant sum decomposition $\sigma(\sum_i h(z_i))$ and proposed simple neural network layers to model equivariant functions. An alternative approach to accommodating variable numbers of voters and alternatives would be to use recurrent architectures such as LSTMs [24] with respect to each dimension, but this would be sensitive to their orderings.

We adopt the Deep Set architecture $\sigma(\sum_i h(z_i))$, where the input z_i of the i -th voter is an $n \times k$ matrix, where n is the number of alternatives and k is the number of features representing each alternative’s count and vote information. We use equivariant functions (in terms of the alternatives) for both the encoder h and decoder




Figure 2: Architecture for our learned aggregation rule, based on Deep Sets [46]. For each voter i , vote features (votes and count information) $\{a_i, b_i, \dots\}$ are embedded via permutation equivariant h . Embedded votes are aggregated across voters using permutation invariant \oplus and passed through a permutation equivariant σ to produce alternative scores.

σ , and take the sum Σ_i across the voters. The decoder σ terminates in a softmax. This architecture satisfies all desiderata outlined above and outputs a proper distribution over outcomes. We train the network to minimize a negative log likelihood (cross entropy) loss where the targets are ground truth best outcome. Training was done via gradient descent for up to 5000 mini-batches of size 128, generated as described in our high variance experiment (Subsection 5.1) with a different random number of voters (sampled uniformly between 5 and 350) and different number of alternatives (sampled uniformly between 5 and 15) for each mini-batch. We tested 20 random hyperparameter configurations from a search space of 144, and kept the model with the lowest loss. See the Appendix for further details, including specific hyperparameters and a full description of our final architecture.

5 EXPERIMENTS

In this section we compare our derived and learned rules to common voting rules in settings of varying uncertainty. We find that our rules consistently outperform anonymous rules, even when there is significant count noise. Code to replicate the experiments is available online at https://github.com/spitis/objective_social_choice.

5.1 Ideal, High Variance Conditions

For this experiment, we generate 100,000 instances of the multi-armed bandit problem with 10 alternatives for different numbers of voters (3, 10, 30, 100, and 300). The ground truth mean of each arm is sampled as $\mu_i \sim \mathcal{N}(0, 1)$ and the voter counts c_{ij} are sampled uniformly between 1 to 50. Individual observations are sampled with high variance from $\mathcal{N}(\mu_i, 1000)$. The voters then report an ordinal ranking based on their estimated means for each alternative.

We compare the performance of our Case 4 and Case 5 rules as well as our learned voting rule to several baselines: basic plurality vote and Borda count [4], naively-modified Plurality vote and Borda count (“Plurality+” and “Borda+”), and a Case 1 oracle. The plurality baseline sets $w_{ij} = 1$ for voter i ’s top choice j , and $w_{ik} = 0$ for $k \neq j$. The Borda baseline sets $w_{ij} = \sum_{k \neq j} \mathbb{1}_{x_{ij} > x_{ik}}$. The Plurality+ and Borda+ baselines take the best performing modification of the basic Plurality and Borda baselines, where the modification uses the count information in an unjustified but plausible way. The tested modifications include weighing each voter’s scores by: the arithmetic mean of that voter’s counts $\{c_{ij}\}$, the harmonic mean of

the counts, or, in each case, the square root and logarithm thereof. The Case 1 oracle sees each voter’s cardinal estimate X and acts as an upper bound on performance. For our Case 5 Monte Carlo approximation we averaged 100 samples from ϕ_{ij} , which we found performed almost as well as 1000 samples and made simulations cheaper. In all cases, ties are broken by random selection.

Performance, as measured by regret, is shown in Table 1a. Relative performance in terms of accuracy (not shown) is approximately the same. The different voting rules are grouped according to access to votes and auxiliary information. Our rules consistently beat anonymous baselines. Among pairwise rules, we observe that our learned aggregation rule has the best overall non-oracle performance, but note that the two Case 4 rules are quite close in performance to the learned rule and are significantly cheaper to compute. The Case 5 results show that the zero approximation is consistently better than the lower bound, and very close to the Monte Carlo approximation (which should give near optimal performance). Finally, we note that all pairwise rules outperform all plurality rules (including Case 5). This is not surprising, as plurality rules use less information than pairwise rules.

5.2 Ideal, Low Variance Conditions

We now consider the same experiment as above under lower observation variance. Instead of sampling observations from $\mathcal{N}(\mu_i, 1000)$, we sample them from $\mathcal{N}(\mu_i, 10)$. The purpose here is two-fold. First, our learned aggregation rule, which was the best performing rule in high variance conditions, was trained in those exact conditions, and we hypothesize its performance will deteriorate out of domain. Second, we note that the failure mode of the unnormalized Case 4 is exacerbated by low variance, and hypothesize that the normalized rule will perform relatively better.

The results, shown in Table 1b, confirm our hypotheses. As compared to the high variance case, both the learned aggregation rule and the unnormalized version of our Case 4 rule do significantly worse relative to the Borda baseline. The normalized Case 4 rule does significantly better than other rules in low variance conditions. Interestingly, the Case 5 zero approximation does slightly better in low variance conditions than the Case 5 Monte Carlo approximation; this suggests that accurate Monte Carlo approximation requires more samples under low observation variance and that the zero approximation is near optimal.

5.3 Noisy Count Conditions

We now relax our assumption of perfect count information by introducing significant noise into the counts c_{ij} that are observed by our rules. This impacts all rules except the anonymous baselines (plurality and Borda). We experiment with two types of count noise: percentage noise applied to all counts, and resampled counts. In the percentage noise case, we adjust *all* reported counts by a percentage between -50% and $+50\%$ (sampled independently and uniformly), rounding to the nearest integer. In the resampled counts case, we replace one third of the reported counts with resampled values (i.e., an integer between 1 and 50). Otherwise, we follow the same procedure as before. To get an idea of how well we could do if the noise were to be expected, we retrain our learned aggregation rule on data generated according to the percentage noise case (but not

Num voters	3	10	30	100	300
Case 1 Oracle	1.1642	0.8625	0.5356	0.2390	0.0936
Borda	1.2116	0.9689	0.6629	0.3308	0.1385
Borda+	1.1863	0.9493	0.6555	0.3270	0.1369
Case 4	1.1760	0.9194	0.6069	0.2890	0.1177
Case 4 (normalized)	1.1879	0.9231	0.6058	0.2886	0.1173
Learned	1.1687	0.9086	0.5935	0.2788	0.1125
Plurality	1.3509	1.2116	1.0089	0.6905	0.3807
Plurality+	1.3232	1.1904	0.9888	0.6721	0.3680
Case 5 (lower bound)	1.2903	1.1547	0.9434	0.6224	0.3302
Case 5 (zero approx)	1.2847	1.1458	0.9297	0.6074	0.3193
Case 5 (Monte Carlo)	1.2848	1.1413	0.9278	0.6066	0.3178

(a) High Variance

Num voters	3	10	30	100	300
Case 1 Oracle	0.1075	0.0312	0.0102	0.0030	0.0011
Borda	0.1754	0.0631	0.0217	0.0068	0.0023
Borda+	0.1688	0.0590	0.0205	0.0062	0.0021
Case 4	0.1711	0.0603	0.0208	0.0064	0.0022
Case 4 (normalized)	0.1479	0.0487	0.0163	0.0050	0.0017
Learned	0.1767	0.0605	0.0211	0.0065	0.0022
Plurality	0.3147	0.0898	0.0290	0.0086	0.0029
Plurality+	0.2586	0.0839	0.0285	0.0086	0.0029
Case 5 (lower bound)	0.2112	0.0740	0.0253	0.0078	0.0026
Case 5 (zero approx)	0.2071	0.0726	0.0253	0.0077	0.0026
Case 5 (Monte Carlo)	0.2089	0.0739	0.0256	0.0078	0.0026

(b) Low Variance

Table 1: Average regret, $V(a^*) - V(\hat{a})$, in ideal conditions. Lower is better. Best non-Oracle rules of each type in bold.

	3	10	30	100	300
Case 1 Oracle	1.1726	0.8811	0.5579	0.2539	0.1002
Learned (noisy)	1.1728	0.9131	0.6012	0.2861	0.1157
Borda	1.2116	0.9689	0.6629	0.3308	0.1385
Borda+	1.2100	0.9643	0.6583	0.3281	0.1369
Case 4	1.1780	0.9194	0.6090	0.2918	0.1184
Case 4 (normalized)	1.1902	0.9234	0.6090	0.2912	0.1188
Learned	1.1793	0.9262	0.6161	0.2964	0.1210
Plurality	1.3509	1.2116	1.0089	0.6905	0.3807
Plurality+	1.3248	1.1889	0.9902	0.6740	0.3681
Case 5 (lower bound)	1.2931	1.1557	0.9466	0.6260	0.3312
Case 5 (zero approx)	1.2874	1.1485	0.9369	0.6160	0.3241
Case 5 (Monte Carlo)	1.2959	1.1548	0.9409	0.6167	0.3243

(a) 50% count noise

	3	10	30	100	300
Case 1 Oracle	1.2194	0.9883	0.6901	0.3492	0.1470
Learned (noisy)	1.1940	0.9433	0.6321	0.3080	0.1267
Borda	1.2116	0.9689	0.6629	0.3308	0.1385
Borda+	1.2100	0.9654	0.6589	0.3298	0.1374
Case 4	1.1926	0.9423	0.6319	0.3079	0.1270
Case 4 (normalized)	1.1996	0.9433	0.6323	0.3080	0.1269
Learned	1.1972	0.9481	0.6390	0.3131	0.1291
Plurality	1.3509	1.2116	1.0089	0.6905	0.3807
Plurality+	1.3331	1.1973	0.9971	0.6834	0.3751
Case 5 (lower bound)	1.3294	1.2165	1.0402	0.7405	0.4255
Case 5 (zero approx)	1.3113	1.1747	0.9701	0.6524	0.3504
Case 5 (Monte Carlo)	1.3067	1.1758	0.9780	0.6592	0.3548

(b) 33% count replacement

Table 2: Average regret, $V(a^*) - V(\hat{a})$, in noisy conditions. Lower is better. Best non-Oracle rules of each type in bold.

the resampled counts case). The experiments for this subsection utilize high observation variance ($\mathcal{N}(\mu_i, 1000)$). The results are shown in Tables 2a and 2b.

In case of 50% count noise, it is unsurprising that the neural network trained under those conditions does best. What is perhaps surprising is how robust the derived rules are to count noise. Both Case 4 and Case 5 rules beat their respective baseline by a respectable margin, even with inaccurate counts. The same trend continues in the case of 33% count replacement, where our Case 4 rules outperform the Oracle (which is no longer a true Oracle). We note, however, that performance declines more sharply in the count replacement case, which is to be expected since the per-count noise is biased. It is interesting to note that the Case 5 zero approximation is more robust to noise than the Monte Carlo approximation. Overall, the results indicate that even inaccurate count information can have significant value.

6 CONCLUSION

In this paper, we proposed a generic framework for objective social choice, which seeks to estimate a cardinal ground truth given noisy

votes. We considered a bandit-based noise model and proposed several voting rules that utilize auxiliary count information to improve inference relative to anonymous rules. Our empirical results confirm the efficacy of our rules relative to anonymous baselines and demonstrate robustness under noise in the auxiliary information.

The scope of the present work assumes that voters have independent information and is limited to a particular noise model and mode of auxiliary information (experience counts). It would be interesting to extend our objective social analysis to cases of dependent information, more general noise models (e.g., noise generated by a contextual bandit [29]), and other forms of auxiliary information (e.g., a similarity kernel between voters). Another extension might study group composition [25]: if we have some control over voter experience, how should we influence the group of voters to improve voting outcomes? We leave these angles to future work.

ACKNOWLEDGMENTS

We thank Nisarg Shah for his guidance throughout this project. We also thank Jimmy Ba, Harris Chan, Mufan Li and the anonymous referees for their helpful comments.

REFERENCES

- [1] Kenneth J Arrow. 2012. *Social choice and individual values*. Vol. 12. Yale university press.
- [2] Kenneth J Arrow, Amartya Sen, and Kotaro Suzumura. 2010. *Handbook of social choice and welfare*. Vol. 2. Elsevier.
- [3] John M Bates and Clive WJ Granger. 1969. The combination of forecasts. *Journal of the Operational Research Society* 20, 4 (1969).
- [4] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. 2016. *Handbook of computational social choice*. Cambridge University Press.
- [5] Ioannis Caragiannis and Evi Micha. 2017. Learning a Ground Truth Ranking Using Noisy Approval Votes.. In *IJCAI*. 149–155.
- [6] Ioannis Caragiannis, Ariel D Procaccia, and Nisarg Shah. 2016. When do noisy votes reveal the truth? *ACM Transactions on Economics and Computation (TEAC)* 4, 3 (2016), 15.
- [7] Richard Y Chen, Szymon Sidor, Pieter Abbeel, and John Schulman. 2017. UCB exploration via Q-ensembles. *arXiv preprint arXiv:1706.01502* (2017).
- [8] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*. 4299–4307.
- [9] Robert T Clemen. 1989. Combining forecasts: A review and annotated bibliography. *International journal of forecasting* 5, 4 (1989).
- [10] Marie J Condorcet. 1785. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. de l'Imprimerie Royale.
- [11] Vincent Conitzer. 2013. The maximum likelihood approach to voting on social networks. In *2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 1482–1487.
- [12] Vincent Conitzer and Tuomas Sandholm. 2005. Common Voting Rules As Maximum Likelihood Estimators. (2005), 8. <http://dl.acm.org/citation.cfm?id=3020336.3020354>
- [13] JP Dickinson. 1973. Some statistical results in the combination of forecasts. *Journal of the Operational Research Society* 24, 2 (1973).
- [14] JP Dickinson. 1975. Some comments on the combination of forecasts. *Journal of the Operational Research Society* 26, 1 (1975).
- [15] Thomas G Dietterich. 2000. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*. Springer.
- [16] JL Fleiss. 1993. Review papers: The statistical basis of meta-analysis. *Statistical methods in medical research* 2, 2 (1993).
- [17] Christian Genest and Kevin J McConway. 1990. Allocating the weights in the linear opinion pool. *Journal of Forecasting* 9, 1 (1990).
- [18] Christian Genest, James V Zidek, et al. 1986. Combining probability distributions: A critique and an annotated bibliography. *Statist. Sci.* 1, 1 (1986).
- [19] Paul Gompers, Joy Ishii, and Andrew Metrick. 2003. Corporate governance and equity prices. *The quarterly journal of economics* 118, 1 (2003), 107–156.
- [20] Clive WJ Granger. 1989. Invited review combining forecasts—Twenty years later. *Journal of Forecasting* 8, 3 (1989).
- [21] Clive WJ Granger and Ramu Ramanathan. 1984. Improved methods of combining forecasts. *Journal of forecasting* 3, 2 (1984).
- [22] John C Harsanyi. 1955. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of political economy* 63, 4 (1955).
- [23] Joshua E Hill. [n. d.]. The minimum of n independent normal distributions. ([n. d.]).
- [24] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [25] Lu Hong and Scott E Page. 2004. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences* 101, 46 (2004), 16385–16389.
- [26] Robert A Jacobs. 1995. Methods for combining experts' probability assessments. *Neural computation* 7, 5 (1995).
- [27] Dongyeop Kang, Waleed Ammar, Bhavana Dalvi, Madeleine van Zuylen, Sebastian Kohlmeier, Eduard Hovy, and Roy Schwartz. 2018. A Dataset of Peer Reviews (PeerRead): Collection, Insights and NLP Applications. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. Association for Computational Linguistics, New Orleans, Louisiana, 1647–1661. <https://doi.org/10.18653/v1/N18-1149>
- [28] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [29] Tor Lattimore and Csaba Szepesvári. 2018. Bandit algorithms. *preprint* (2018).
- [30] David D Lewis. 1998. Naive (Bayes) at forty: The independence assumption in information retrieval. In *European conference on machine learning*. Springer.
- [31] Colin L Mallows. 1957. Non-null ranking models. I. *Biometrika* 44, 1/2 (1957).
- [32] Kenneth O May. 1952. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica: Journal of the Econometric Society* (1952), 680–684.
- [33] Kevin J McConway. 1981. Marginalization and linear opinion pools. *J. Amer. Statist. Assoc.* 76, 374 (1981), 410–414.
- [34] Michael P Perrone and Leon N Cooper. 1992. When networks disagree: Ensemble methods for hybrid neural networks. (1992).
- [35] Ariel D Procaccia and Jeffrey S Rosenschein. 2006. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents*. Springer.
- [36] Ariel D Procaccia, Nisarg Shah, and Eric Sodomka. 2015. Ranked voting on social networks. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- [37] Lior Rokach. 2010. Ensemble-based classifiers. *Artificial Intelligence Review* 33, 1-2 (2010).
- [38] Amartya Sen. 2018. *Collective choice and social welfare*. Harvard University Press.
- [39] Volker Tresp and Michiaki Taniguchi. 1995. Combining estimators using non-constant weighting functions. In *Advances in neural information processing systems*. 419–426.
- [40] John Von Neumann and Oskar Morgenstern. 1953. Theory of games and economic behavior. (1953).
- [41] Kenneth F Wallis. 2011. Combining forecasts—forty years later. *Applied Financial Economics* 21, 1-2 (2011).
- [42] John A Weymark. 1991. A reconsideration of the Harsanyi–Sen debate on utilitarianism. *Interpersonal comparisons of well-being* 255 (1991).
- [43] Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8, 3-4 (1992), 229–256.
- [44] H Peyton Young. 1988. Condorcet's theory of voting. *American Political science review* 82, 4 (1988).
- [45] Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. 2012. The k-armed dueling bandits problem. *J. Comput. System Sci.* 78, 5 (2012), 1538–1556.
- [46] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan R Salakhutdinov, and Alexander J Smola. 2017. Deep sets. In *Advances in neural information processing systems*. 3391–3401.