Learning to Design Fair and Private Voting Rules (Extended Abstract)*

Farhad Mohsin[1] , **Ao Liu**[1] , **Pin-Yu Chen**[2] , **Francesca Rossi**[2] and **Lirong Xia**[1]
[1]Rensselaer Polytechnic Institute
[2]IBM Research

Abstract

Voting is used widely to aggregate preferences to make a collective decision. In this paper, we focus on evaluating and designing voting rules that support both the privacy of the voting agents and a notion of fairness over such agents. We introduce a novel notion of group fairness and adopt the existing notion of local differential privacy. We evaluate the level of group fairness in several existing voting rules, showing that it is not possible to always obtain maximal economic efficiency with high fairness. Then, we present both a machine learning and a constrained optimization approach to design new voting rules that are fair while maintaining a high level of economic efficiency. Finally, we empirically examine the effect of adding noise to create local differentially private voting rules and discuss the three-way trade-off between economic efficiency, fairness, and privacy.

1 Introduction

Voting is one of the most used and well-studied methods to make a collective decision [Brandt *et al.*, 2016]. Anonymity is considered an important axiom for preserving the "one person, one vote" principle. However, it may lead to the well-known "tyranny of the majority" [Mill, 1859] in some scenarios. While there exist other notions of fairness in voting (as discussed by [Bredereck *et al.*, 2018; Celis *et al.*, 2018; Chamberlin and Courant, 1983; Monroe, 1995] and more), we focus on the notion of *group fairness* [Chouldechova and Roth, 2020]. This draws motivation from the relevant literature in fair algorithmic decision-making and machine learning (ML) (see [Corbett-Davies *et al.*, 2017; Kleinberg, 2018; Verma and Rubin, 2018; Chouldechova and Roth, 2020] for an exposition). Due to bias in data or training methodologies, a system can be biased towards one group of people in terms of accuracy, positive predictive value, etc. To avoid this, fairness is defined over protected features (e.g., gender, race, etc.) that indicate group membership. For example, we may require that the prediction accuracy is equal for different groups, such as men and women. While voting does not have metrics like prediction accuracy, we can consider an analogous scenario where the average utility received by different groups of agents is equal. We, therefore, define a novel way to measure group fairness in voting (Definitions 1-2), focusing on how voting outcomes affect different groups of agents. We then investigate existing voting rules in terms of the trade-off between fairness and economic efficiency and find worst-case results for fairness (Table 2) that show that well-studied voting rules can be very unfair in the worst case.

However, this way of defining and achieving fairness over groups of agents needs to expose features of agents, since such features define the group that each agent belongs to. This means that the voting process is not anonymous, which leads to privacy concerns. To circumvent this, we employ the notion of *local differential privacy (local DP)* [Evfimievski *et al.*, 2003], which is a generalization of *differential privacy (DP)* [Dwork *et al.*, 2006], to ensure that an adversary cannot learn too much about the voting behaviors of the agents from the voting outcome. While there exist work that consider differentially private methods of voting, such as [Joseph *et al.*, 2018; Yan *et al.*, 2020; Wang *et al.*, 2019], we are not aware of a study on the three way-trade-off between fairness, economic efficiency, and privacy. We theoretically study this trade-off (Theorem 3) and show that- a high privacy requirement results in high efficiency loss, but we can have moderate privacy with only a small decrease in efficiency or fairness.

Finally, we present two automated frameworks to design voting rules with varying levels of fairness, privacy, and economic efficiency. Use of automatic mechanism design [Conitzer and Sandholm, 2003] was previously considered for voting by [Xia, 2013; Armstrong and Larson, 2019; Anil and Bao, 2021], albeit in different formats. For the first framework, we define a family of voting rules that maximize fairness under efficiency constraints and can be thought of as a natural extension of positional scoring rules such as Plurality, Borda, etc. The extension comes from looking at alternative scores as indicators for group utilities. The second framework employs a machine learning-based approach that allows us to design fair and efficient voting rules that go beyond just positional scoring rules and work with more general notions of economic efficiency and fairness. Experimentally, we show that the learned family of voting rules succeeds in achieving high fairness and efficiency satisfaction levels,

*Full version [Mohsin et al., 2022] published in Journal of Artificial Intelligence Research 75 (2022): 1139-1176 based on simulations on synthetic data. In particular, our newly designed voting rules are never dominated by a voting rule that focuses on purely economic efficiency or group fairness. Finally, we experimentally verify our theoretical results for the fairness-efficiency-privacy trade-off, showing that for moderate privacy requirements (when the noise level is not very high), the loss in efficiency and fairness is small.

2 Preliminaries

Voting Rules. Let $\mathcal{A} = \{a_1, \ldots, a_m\}$ be the set of m alternatives and $\mathcal{L}$ be the set of all rankings or linear orders (linear orders are anti-symmetric, transitive, and total binary relations) over $\mathcal{A}$. There are n agents (voters), each provides a full ranking, $R \in \mathcal{L}$ over $\mathcal{A}$ as her vote. In a ranking R, if alternative a is preferred to another alternative b, we write $a \succ_R b$. A collection of n votes, $P \in \mathcal{L}^n$ is called a preference profile. A voting rule is a mapping $r : \mathcal{L}^n \mapsto \mathcal{A}$ that chooses a winner from the preference profile. To indicate protected features, we assume each agent is a member of one of two disjoint groups with group sizes n_1, n_2 $(n_1 + n_2 = n)$. Each group has preference profile $P_k \in \mathcal{L}^{n_k}$ for $k = 1, 2$. To consider group memberships, we redefine *voting rule* as a mapping from a collection of preference profiles to a winner, $r : \mathcal{L}^{n_1} \times \mathcal{L}^{n_2} \mapsto \mathcal{A}$. We refer to P as the preference profile with all agents. For voting rules that do not use group membership, $r(P)$ and $r(P_1, P_2)$ are the same. Most of the theoretical and experimental work in the paper is presented for two groups, so we focus the preliminaries on two-group scenarios as well.

A common family of voting rules is **positional scoring rules**, which have score vector $\vec{s} = \langle s_1, \ldots, s_m \rangle$ such that $s_1 \geq \cdots \geq s_m$ and $s_1 > s_m$. For each ranking R, the j-th ranked alternative gets a score of s_j. Given a preference profile, the alternative with maximum total score will be the winner. Some popular scoring rules are: **Plurality**, with scoring vector $\langle 1, 0, \ldots, 0 \rangle$; **Borda**, with $\langle m-1, m-2, \ldots, 1, 0 \rangle$; **Veto**, with $\langle 1, 1, \ldots, 1, 0 \rangle$. **Condorcet rules** is a family of voting rules that are defined by a different measure of efficiency called the Condorcet criterion. For a preference profile, if an alternative beats all other alternatives in pairwise comparison, it is called the *Condorcet winner*. A voting rule satisfies the Condorcet criterion if it always selects the Condorcet winner whenever it exists. For example, the **Copeland** rule chooses the alternative that maximizes the number of alternatives that it beats in pairwise comparisons.

Economic Efficiency in Voting. In this paper, we consider two types of economic efficiency that are popular in the social choice literature, and both are related to the two families of voting rules that we mentioned before: Condorcet rules and positional scoring rules.

For a preference profile, P, the Condorcet winner exists only if there is an alternative that beats all other alternatives in pairwise comparison. We measure *Condorcet efficiency (CE)* as the fraction of preference profiles where a voting rule winner is identical to the Condorcet winner. This is an efficiency measure as efficient decisions (output of voting rules) should be preferred to all other alternatives. Condorcet rules like Copeland have a CE value of 1, whereas positional scoring rules have CE values less than one.

Our other efficiency notion takes a utilitarian view [Boutilier *et al.*, 2015], where each agent can receive different cardinal utilities from the alternatives. For an agent, a utility function $u : \mathcal{A} \times \mathcal{L} \mapsto \mathbb{R}$ defines alternative a's utility to agent i. For this paper, we limit ourselves to every agents having the same utility function, which is dependent only on the rank. Thus, we assume that a utility function u is defined by a vector $\vec{u} = \langle u_1, \ldots, u_m \rangle$ such that $u_1 \geq \cdots \geq u_m$ and $u_1 > u_m$. If an alternative a is ranked j-th in R, then $u(a, R) = u_j$.

All $u \in \mathcal{U}$ use a vector similar in definition to score vectors of positional scoring rules. To reduce confusion, we will use $\vec{s}$ for the score vectors and $\vec{u}$ for utility functions.[1] The *average utility* for alternative a is $W(a, P) = \frac{1}{n} \sum_{R \in P} u(a, R)$ and is a measure of economic efficiency.

Local Differential Privacy (DP). We adapt the formal definition of local differential privacy [Evfimievski *et al.*, 2003] to the domain of voting and state its difference from standard DP. A randomized voting rule r is said to be ε -local DP, if for any agent j, any alternative $a \in \mathcal{A}$, and any pair of rankings $R, R' \in \mathcal{L}$, the following holds: $\Pr[r(R, P_{-j}) = a] \leq \exp(\varepsilon) \cdot \Pr[r(R', P_{-j}) = a]$, where P_{-j} is the preference profile with all agents other than agent j. In particular, this indicates that the vote of any single agent will be hard to infer from the outcome. Hence, it gives a privacy guarantee to the agents. Smaller ε means stronger privacy guarantee. We note that voting rules with local DP (or standard DP) must be randomized. So, our local differentially private voting rules can only be used where some randomization would not be very problematic.

3 Group Fairness in Voting

In presence of pre-defined groups among agents, traditional voting rules that are anonymous and do not differentiate between different agents may be unfair to some groups. Consider the scenario presented in Table 1 with two groups of agents and three alternatives. Alternative A receives high utility from the larger group, G_1, and in turn has the highest average utility for the whole population. However, A has zero utility for the smaller group , making this an unfair decision. On the other hand, aternative B is not ranked lowest in terms of average utility to either group and can be viewed as a more fair decision than A.

We now present our formal definitions for group imbalance and imbalance-based group fairness. First, imbalance indicates the unfairness of a candidate given a preference profile.

Definition 1 (Group imbalance)**.** *Given a utility function u, an alternative $a \in \mathcal{A}$, and preference profiles P_1, P_2 for two groups of agents, imbalance between the two groups in terms of u for a is*

$$Imb(u, a, P_1, P_2) = \begin{cases} \frac{|W(a,P_1) - W(a,P_2)|}{W(a,P)} & \text{if } W(a, P) > 0, \\ 0 & \text{otherwise}. \end{cases}$$

[1] We want to emphasize that we use different vectors $\vec{s}$ and $\vec{u}$ for the score vectors and utilities intentionally. While the positional scoring rules has the notion of maximizing some sort of utility measure, the scoring vector used may be different from a true utility function.

	#agents	Ranking	Utility A	B	C
Group G_1	30	A $\succ$ B $\succ$ C	60	30	0
	30	B $\succ$ A $\succ$ C	30	60	0
	20	A $\succ$ C $\succ$ B	40	0	20
Group G_2	10	B $\succ$ C $\succ$ A	0	20	10
	10	C $\succ$ B $\succ$ A	0	10	20
G_1 average			1.625	1.125	0.25
G_2 average			0	1.5	1.5
Average			1.3	1.2	0.5

Table 1: Sample preference profile and group utilities

	u_{top}	u_{rank}
Plurality	0 if $z < m-1$ $1-1/z$, otherwise	0
Borda	0 if $m \geq 3$ $1-1/z$ otherwise	$1-1/z$
Copeland	0	$1-1/z$

Table 2: $F(r,u,n,z)$ for Plurality, Borda and Copeland under u_{top} and u_{rank}.

where P is the combined preference profile for all agents.

We define imbalance-based group fairness of a voting rule for a utility function in terms of the worst-case imbalance achieved by the voting rule winner for any preference profile. Group size information are summarized in parameters, group size ratio, $z = \frac{\max(n_1,n_2)}{\min(n_1,n_2)}$ and total agents, $n = n_1 + n_2$.

Definition 2 (Group imbalance-based fairness)**.** *Given a voting rule, r, utility function, u, n total agents, and group ratio, z, the group imbalance-based fairness is*

$$F(r,u,n,z) = 1 - \tfrac{1}{1+z} \cdot \max_{\substack{P_1 \in \mathcal{L}^{n_1} \\ P_2 \in \mathcal{L}^{n_2}}} \textit{Imb}(u, r(P_1,P_2), P_1, P_2).$$

Definition 2 gives a notion of fairness that is not specific to a preference profile, but rather looks at the worst-case. The fairness value is between 0 and 1 and a value of 0 indicates maximum possible unfairness.

Also, based on the definition of group imbalance, assuming what the utility function is, we can define maximum fair voting rules. For a utility function, u, the u-fair voting rule is guaranteed to have the highest fairness value.

Definition 3 (u-fair voting rules)**.** *For utility function u, the u-fair voting rule, r^u_{fair}, is a voting rule that chooses the alternative with minimum imbalance with respect to utility function u for any preference profile.*

4 Theoretical Results for Group Fairness

We proved a number of theoretical results for fairness values for well-known voting rules. Here, we present some results about popular voting rules, Plurality, Borda and Copeland. We consider the utility functions defined by $\vec{u}_{top} = \langle 1, 0, \ldots, 0 \rangle$ (top-1) and $\vec{u}_{rank} = \langle m-1, m-2, \ldots, 0 \rangle$ (rank). With the top-1 utility function, the agents only receive utility if their top ranked alternative wins. A fairness value of 0 for Borda means that, for some preference profile, the Borda winner will have the worst-case imbalance value, being highly unfair. But if we choose the Plurality winner, at least in some cases we are guaranteed a positive fairness value. On the other hand, with the rank utility function, where the utility linearly decreases along with rank, we see that the Borda winner has the same fairness value as the most fair voting rule, whereas Plurality's fairness value is 0. In both cases, the voting rule that maximizes the utility function has better fairness results. Table 2 for a summary of these results.

To see the more general results for any positional scoring rule and any Condorcet voting rule under general utility function notions, please refer to the full version of the paper. However, we want to present a somewhat negative result in that all traditional efficiency-maximizing voting rules may turn out to be unfair under some circumstance.

Theorem 1. *If r is any positional scoring rule or Condorcet voting rule, there exist some utility function $u \in \mathcal{U}$, and some group size parameters, n and z, such that*

$$\min_{\substack{u \in \mathcal{U} \\ n \in \mathbb{N}, z \geq 1}} F(r,u,n,z) = 0.$$

We also proved fairness result for u-fair voting rules.

Theorem 2. *Given n total agents, and group size ratio z,*

$$F(r^{top}_{fair}, u_{top}, n, z) = 1 - \frac{1}{z},$$
$$F(r^{rank}_{fair}, u_{rank}, n, z) = 1 - \frac{1}{z}.$$

5 Designing Fair and Efficient Voting Rules

We propose two frameworks for designing fair and efficient voting rules.

Framework 1: Utility-constrained Fair Voting Rules. Based on an assumed utility function, u, we can define constrained fair voting rules as a compromise between positional scoring rules and u-fair rules. As an example, we present fairness constrained version of Borda below in Definition 4.

Definition 4 (α-efficient fair Borda (α-FB))**.** *Given preference profiles P_1, P_2 for two groups of agents, $\alpha \in [0,1]$, the α-efficient Fair Borda winner, $r_{\alpha-FB}(P_1,P_2)$ is given by*

$$\underset{a \in \mathcal{A}}{\text{minimize}} \quad \textit{Imb}(u_{rank}, a, P_1, P_2)$$
$$\textit{subject to} \quad W_{rank}(a,P) \geq \alpha \cdot \max_{a' \in \mathcal{A}} W_{rank}(a',P).$$

Framework 2: ML-based Framework for Fair-efficient Voting Rules. We note that a voting rule r can be viewed as a multi-class classifier: the input is a preference profile P and the classes are the alternatives in $\mathcal{A}$. From this viewpoint, we propose a learning framework that generates synthetic data with random preference profiles. As feature vectors, we can use summary features for preference profiles like weighted majority graph or positional score matrix. We create a synthetic dataset, where part of the labels come from fair winners and the rest of the labels come from efficient winners. This mixing causes a classifier trained on this model to behave like a voting rule that is fairer than purely efficient voting rules and more efficient than purely fair voting rules. We use two methods for mixing two kinds of data in the synthetic dataset–one, sampling data from both fair and efficient voting rules and two, using soft probabilistic labels. Finally, we found best results while using the XGBoost model as a classifier [Chen and Guestrin, 2016]. We discuss experimental results for this in Section 7.

6 Adding Privacy to Collective Decision-making

We propose adding noise through the flipping-coin or randomized response method to achieve local DP. The process outputs the input vote without changing it with some probability p, and otherwise uniformly outputs any of the possible values. Consider $f_p(P)$ to be the noisy output for preference profile P. We find that, adding the noise to individual votes cause a trade-off with both fairness and utility.

Theorem 3 (Fairness-Privacy-utility Trade-off)**.** *For any ε-local DP requirement on making collective decisions with two groups, we have the following:*

$$(1)\ \Pr\left[\hat{W}(a,P) \geq W(a,P) - t\right] \geq 1 - \exp\left[-\frac{2t^2p^2n}{(\Delta u_{\max})^2}\right],$$

$$(2)\ \Pr\Bigg[\hat{Imb}(u,a,P_1,P_2) \leq Imb(u,a,P_1,P_2) + \frac{\Delta u_{\max}}{p}\cdot\frac{(Imb(u,a,P_1,P_2)+1)(n_1^{-0.3}+n_2^{-0.3})}{W(a,P)-\frac{\Delta u_{\max}}{p}\left(n_1^{-0.3}+n_2^{-0.3}\right)}\Bigg] \geq 1 - 2\exp\left(-2n_1^{0.4}\right) - 2\exp\left(-2n_2^{0.4}\right),$$

where $p = \frac{\exp(\varepsilon)-1}{|\mathcal{X}|+\exp(\varepsilon)-1}$ *and* $\Delta u_{\max} \triangleq u_1 - u_m$.

With high probability, the utility estimator is close to the actual utility. Additionally, if n_1, n_2 are high, alternatives with low imbalance (fair alternatives) and high utility (efficient alternatives) will have less noisy imbalance estimates.

7 Experimental Results

While our theoretical results for fairness deal with worst-case analysis, average-case analysis is also important. So, we do empirical analysis on synthetic data to get an idea about the average fairness-efficiency trade-off. To get synthetic data, we assume that all agent preferences for agents in the same group come from the same distribution, described using a statistical model. We use two types of models in our experiments. First, uniform or impartial culture, where everyone's vote is entirely random. This signifies the scenario where both group votes randomly and similarly. Second, using a Plackett-Luce model [Plackett, 1975; Luce, 1959]. For this, we define two groups with two sets of Plackett-Luce parameters and rankings sampled from there. We do this to simulate similar behavior between in-group votes and dissimilarity between across-group votes.

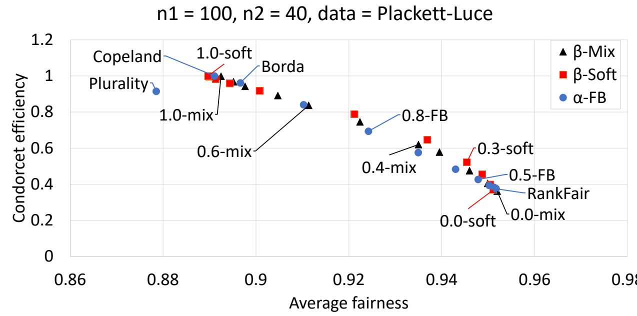

Figure 1: Trade-off between economic efficiency and fairness for various voting rules. Voting rules considered are Plurality, Borda, Copeland, α-efficient FB rules and the learned rules. RankFair indicates the u_{rank}-fair voting rule. The learned rules use the β-Mix and β-Soft method, both using Condorcet efficiency as the efficiency measure while training.

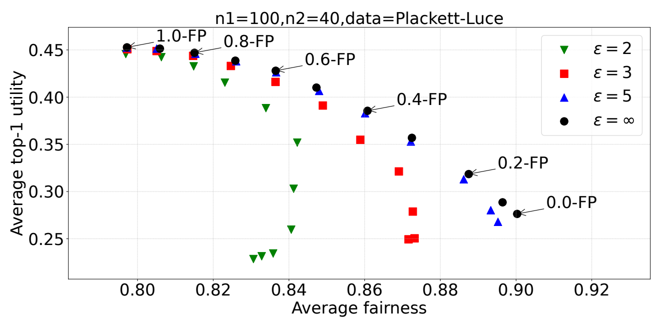

Figure 2: Efficiency (Average top-1 utility) and average fairness for various levels of privacy. High ε values indicate lower privacy requirements with $\varepsilon = \infty$ indicating non-private voting rule.

Trade-off between Fairness and Efficiency. In Figure 1, we see the trade-off between one notion of economic efficiency (Condorcet efficiency) and group fairness. As expected Copeland is the most efficient, being a Condorcet voting rule. However, it is highly unfair. On the other end, we see the rank-fair voting rule, which is most fair but highly inefficient. Both frameworks for designing new rules provide voting rules at different levels of trade-off between the two properties. Also, the learned voting rules from both learning methods, β-Mix and β-Soft , mostly dominate α-FB methods, and provide a good improvement in terms of fairness compared to Copeland (a Condorcet consistent rule) while achieving almost similar levels of Condorcet efficiency. We present experimental results for other efficiency measures in the full paper.

Trade-off between Fairness, Efficiency and Privacy. From the results in Figure 2, we see that for low privacy requirement ($\varepsilon = 5$), the loss in terms of average fairness and average utility is minimal. For higher privacy requirements, both fairness and efficiency suffer from the noise-adding mechanism. However, for $\varepsilon = 3$, which is a moderately standard setting, we see relatively low loss ($5 \sim 15\%$) for both average utility and fairness.

Acknowledgments

The authors are grateful to Michael Hind for discussions and the reviewers for their insightful comments and suggestions. This work was supported by the Rensselaer-IBM AI Research Collaboration, part of the IBM AI Horizons Network. Ao Liu is supported by the RPI-IBM AI Horizons scholarship. Lirong Xia acknowledges NSF 1453542, 1716333, 2007994, 2106983, and a Google gift fund for support.

Ethics Statement

We work on two important issues in collective decision-making – fairness and privacy and the problems we tackle in this paper has the broad goal of getting better collective decisions. The machine learning part of the experiments use synthetic data, generated from statistical models so there are no privacy concerns regarding the data.

References

[Anil and Bao, 2021] Cem Anil and Xuchan Bao. Learning to elect. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems*, pages 8006–8017. NeurIPS, 2021.

[Armstrong and Larson, 2019] Ben Armstrong and Kate Larson. Machine learning to strengthen democracy. In *AI for Social Good workshop at NeurIPS-2019*, 2019.

[Boutilier *et al.*, 2015] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D Procaccia, and Or Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.

[Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. *Handbook of computational social choice*. Cambridge University Press, 2016.

[Bredereck *et al.*, 2018] Robert Bredereck, Piotr Faliszewski, Ayumi Igarashi, Martin Lackner, and Piotr Skowron. Multiwinner elections with diversity constraints. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, pages 933–940. AAAI, 2018.

[Celis *et al.*, 2018] L Elisa Celis, Lingxiao Huang, and Nisheeth K Vishnoi. Multiwinner voting with fairness constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 144–151. IJCAI, 2018.

[Chamberlin and Courant, 1983] John R Chamberlin and Paul N Courant. Representative deliberations and representative decisions: Proportional representation and the borda rule. *American Political Science Review*, 77(3):718—-733, 1983.

[Chen and Guestrin, 2016] Tianqi Chen and Carlos Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 785–794. ACM, 2016.

[Chouldechova and Roth, 2020] Alexandra Chouldechova and Aaron Roth. A snapshot of the frontiers of fairness in machine learning. *Communications of the ACM*, 63(5), 2020.

[Conitzer and Sandholm, 2003] Vincent Conitzer and Tuomas Sandholm. Applications of automated mechanism design. In *Workshop on Bayesian Modeling Applications at UAI-2003*, 2003.

[Corbett-Davies *et al.*, 2017] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 797–806. ACM, 2017.

[Dwork *et al.*, 2006] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Theory of Cryptography Conference*, pages 265–284. Springer, 2006.

[Evfimievski *et al.*, 2003] Alexandre Evfimievski, Johannes Gehrke, and Ramakrishnan Srikant. Limiting privacy breaches in privacy preserving data mining. In *Proceedings of the 22nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, pages 211—-222. ACM, 2003.

[Joseph *et al.*, 2018] Matthew Joseph, Aaron Roth, Jonathan Ullman, and Bo Waggoner. Local differential privacy for evolving data. In *Proceedings of the 32nd Annual Conference Neural Information Processing Systems*, pages 2381—-2390. NeurIPS, 2018.

[Kleinberg, 2018] Jon Kleinberg. Inherent trade-offs in algorithmic fairness. In *Abstracts of the 2018 ACM International Conference on Measurement and Modeling of Computer Systems*, pages 40–40. ACM, 2018.

[Luce, 1959] R. Duncan. Luce. *Individual choice behavior - A Theoretical Analysis*. Wiley, 1959.

[Mill, 1859] John S Mill. *On liberty*. Longman, Roberts, & Green Co, 1859.

[Mohsin *et al.*, 2022] Farhad Mohsin, Ao Liu, Pin-Yu Chen, Francesca Rossi, and Lirong Xia. Learning to design fair and private voting rules. *Journal of Artificial Intelligence Research*, 75:1139–1176, 2022.

[Monroe, 1995] Burt L Monroe. Fully proportional representation. *American Political Science Review*, 89(4):925––940, 1995.

[Plackett, 1975] R. L. Plackett. The analysis of permutations. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 24, 1975.

[Verma and Rubin, 2018] Sahil Verma and Julia Rubin. Fairness definitions explained. In *Proceedings of the 2018 IEEE/ACM International Workshop on Software Fairness*. ACM, 2018.

[Wang *et al.*, 2019] Shaowei Wang, Jiachun Du, Wei Yang, Xinrong Diao, Zichun Liu, Yiwen Nie, Liusheng Huang, and Hongli Xu. Aggregating votes with local differential privacy: Usefulness, soundness vs. indistinguishability. arXiv preprint arXiv:1908.04920, 2019.

[Xia, 2013] Lirong Xia. Designing social choice mechanisms using machine learning. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems*, pages 471–474. IFAAMAS, 2013.

[Yan *et al.*, 2020] Ziqi Yan, Gang Li, and Jiqiang Liu. Private rank aggregation under local differential privacy. *International Journal of Intelligent Systems*, 35(10):1492–1519, 2020.