Differentially Private Condorcet Voting

Zhechen Li[1], Ao Liu[2], Lirong Xia[2], Yongzhi Cao[1*], Hanpin Wang[3,1]

[1]Key Laboratory of High Confidence Software Technologies (MOE), School of Computer Science, Peking University, China
[2]Department of Computer Science, Rensselaer Polytechnic Institute
[3]School of Computer Science and Cyber Engineering, Guangzhou University, China
lizhechen@pku.edu.cn, liua6@rpi.edu, xialirong@gmail.com, caoyz@pku.edu.cn, whpxhy@pku.edu.cn

Abstract

Designing private voting rules is an important and pressing problem for trustworthy democracy. In this paper, under the framework of differential privacy, we propose a novel famliy of randomized voting rules based on the well-known Condorcet method, and focus on three classes of voting rules in this family: Laplacian Condorcet method ($\mathsf{CM}_\lambda^{\text{LAP}}$), exponential Condorcet method ($\mathsf{CM}_\lambda^{\text{EXP}}$), and randomized response Condorcet method ($\mathsf{CM}_\lambda^{\text{RR}}$), where λ represents the level of noise. We prove that all of our rules satisfy absolute monotonicity, lexi-participation, probabilistic Pareto efficiency, approximate probabilistic Condorcet criterion, and approximate SD-strategyproofness. In addition, $\mathsf{CM}_\lambda^{\text{RR}}$ satisfies (non-approximate) probabilistic Condorcet criterion, while $\mathsf{CM}_\lambda^{\text{LAP}}$ and $\mathsf{CM}_\lambda^{\text{EXP}}$ satisfy strong lexi-participation. Finally, we regard differential privacy as a voting axiom, and discuss its relations to other axioms.

1 Introduction

Voting is a commonly used method for group decision making, where voters submit their preferences over a set of alternatives, and then a voting rule is applied to choose the winner. A major and classical paradigm behind the design and analysis of voting rules is the *axiomatic approach* (Plott 1976), under which voting rules are evaluated by their satisfaction to various normative properties, known as *(voting) axioms*. For example, the *Condorcet criterion* requires that whenever there exists a *Condorcet winner*, which is the alternative that beats all other alternatives in their head-to-head competitions, it must be selected as the winner.

Recently, privacy in voting has become a critical public concern. There are a series of works on examining the *differential privacy (DP)* (Dwork 2006) of voting (Shang et al. 2014; Hay, Elagina, and Miklau 2017; Yan, Li, and Liu 2020). These works mainly focused on applying several randomized mechanisms to existing voting rules, proving upper bounds on the privacy-preserving level (also called *privacy budget*, denoted by ϵ throughout the paper), and then evaluating the utility loss (measured by accuracy or mean square error) due to randomness. However, the upper bounds on privacy in most of them are not tight, which means that the exact privacy-preserving level of the mechanisms is unclear. Moreover, we are not aware of a previous work on making voting private while maintaining the satisfaction to desirable voting axioms beyond strategyproofness (Lee 2015). Therefore, the following question remains largely open.

How can we design private voting rules that satisfy desirable axiomatic properties?

Our contributions. We propose a novel class of randomized voting rules, denoted by $\mathsf{CM}_\lambda^{\texttt{Rand}}$, based on the celebrated *Condorcet method*, which chooses the Condorcet winner when it exists, where `Rand` is a randomized function (called a *mechanism* in DP literature) that introduces noises to pairwise comparisons between alternatives, and λ represents the level of noise. To choose a winner, $\mathsf{CM}_\lambda^{\texttt{Rand}}$ applies `Rand` with parameter λ to the pairwise comparisons for the input profile until a Condorcet winner appears, and then chooses it as the winner.

We focus on three classes voting rules in this family, namely $\mathsf{CM}_\lambda^{\text{LAP}}$, $\mathsf{CM}_\lambda^{\text{EXP}}$, and $\mathsf{CM}_\lambda^{\text{RR}}$, which are obtained by applying the Laplace mechanism, exponential mechanism, and randomized response mechanism, respectively. Under these mechanisms, while it may take exponentially many iterations to obtain the winner by definition, we show that the winner can be efficiently sampled (Lemma 1).

Our main technical contributions are three-fold. First, we prove that all the three classes of voting rules are differentially private by characterizing the upper and lower bounds on the privacy budget ϵ (Theorem 1). Second, we study the satisfaction of our voting rules to probabilistic variants to Condorcet criterion (p-Condorcet, requiring the winning rate of the Condorcet winner is not lower than the other alternatives), Pareto efficiency (p-Pareto, which requires the winning rate of a is not lower than b, if a Pareto dominates b), monotonicity (a-monotonicity, which ensures the winning rate of each alternative does not decrease when her ranking is lifted by any voter simply), strategyproofness (SD-strategyproofness, SD-SP for short, which ensures that no voter can benefit herself in the sense of stochastic dominance by changing her vote), and participation (lexi-participation, which ensures that no voter can improve the result of the voting lexicographically by withdrawing her vote). Besides, we consider the approximate version of p-Condorcet (α-p-Condorcet, Definition 5) and SD-SP (α-SD-SP, Definition

*Corresponding Author.

	p-Condorcet	α-p-Condorcet	p-Pareto	a-Mono.	α-SD-SP	Lexi-Par.	Strong Lexi-Par.
$\mathsf{CM}_\lambda^{\mathrm{RR}}$	✓	e^{λ}	✓	✓	$\mathrm{e}^{(2-2m)\lambda}$	✓	✗
$\mathsf{CM}_\lambda^{\mathrm{EXP}}$	✗	$\frac{1+\mathrm{e}^{\lambda/2}}{\left(1+\mathrm{e}^{-\lambda/2}\right)^{m-1}}$	✓	✓	$\mathrm{e}^{(2-2m)\lambda}$	✓	✓
$\mathsf{CM}_\lambda^{\mathrm{LAP}}$	✗	$2\mathrm{e}^{\lambda}\left(1-\frac{\mathrm{e}^{-\lambda}}{2}\right)^{m-1}$	✓	✓	$\mathrm{e}^{(2-2m)\lambda}$	✓	✓

Table 1: The satisfaction of our mechanisms to the voting axioms, where "✓" indicates that the row rule satisfies the column axiom, and "✗" indicates that the rule does not satisfy the axiom. The expressions in the table represent the level of satisfaction to the approximate axioms (the α in α-p-Condorcet and α-SD-SP).

7), and the strong version of lexi-participation (Definition 8). Our results suggest that $\mathsf{CM}_\lambda^{\mathrm{LAP}}$ outperforms $\mathsf{CM}_\lambda^{\mathrm{EXP}}$ in all aspects examined in the paper, while $\mathsf{CM}_\lambda^{\mathrm{RR}}$ sometimes achieves better p-Condorcet but only satisfies standard lexi-participation, instead of the strong version (Theorems 2 - 8). The results in the second part are summarized in Table 1. Third, we investigate the relations between DP and the voting axioms. We prove that Condorcet criterion and Pareto efficiency are incompatible with DP, and capture the upper bounds of satisfaction to p-Condorcet under ϵ-DP (Proposition 4 - 6). Besides, we show that DP guarantees a lower bound of satisfaction to SD-strategyproofness (Proposition 7). All of the missing proofs can be found in Appendix.

Related work and discussions. To the best of our knowledge, DP was first applied to the rank aggregation problem in (Shang et al. 2014). They analyzed the error rates and derived upper bounds on them. Lee proposed an algorithm which is both differentially private and robust to strategic manipulation for tournament voting rules (Lee 2015). Hay et al. used Laplace mechanism and exponential mechanism to improve the privacy of Quicksort and Kemeny-Young method (Hay, Elagina, and Miklau 2017). Kohli and Laskowski explored DP, strategyproofness, and anonymity for voting on single-peaked preferences (Kohli and Laskowski 2018). Torra analyzed the privacy-preserving level of random dictatorship with DP, which is a well-known randomized voting rule (Torra 2019). He investigated the condition where random dictatorship is differentially private, and improved the mechanism to achieve DP for general cases. Yan et al. made tradeoff between accuracy and privacy in rank aggregation to achieve local DP via Laplace mechanism and randomized response (Yan, Li, and Liu 2020).

Most of the above works did not consider the tradeoffs between privacy and those desirable properties, and the privacy bounds of them are usually not tight. Liu et al. proposed the exact version of distributional DP (Bassily et al. 2013) and studied the privacy-preserving level of several voting rules, but they did not investigate how to improve the privacy (Liu et al. 2020). Beyond social choice, DP has also been considered in other topics of economics, such as mechanism design (Pai and Roth 2013; Xiao 2013), and matching and resource allocation (Hsu et al. 2016; Kannan et al. 2018).

There is a large literature on the analysis of randomized voting (Brandt 2017), most of them studied the satisfaction to axiomatic properties, e.g., complexity of manipulation (Walsh and Xia 2012), strategyproofness (Aziz, Brandl, and Brandt 2014, 2015), Pareto efficiency (Brandl, Brandt, and Hofbauer 2015; Gross, Anshelevich, and Xia 2017), participation (Brandl, Brandt, and Hofbauer 2019) and monotonicity (Brandl, Brandt, and Stricker 2018). The fairness properties of sortition have also been investigated (Benadè, Gölz, and Procaccia 2019; Flanigan et al. 2020, 2021).

The approximation of those properties was also studied. Procaccia discussed how much a strategyproof randomized rule could approximate a deterministic rule (Procaccia 2010). Birrell and Pass explored the approximate strategyproofness for randomized voting rules (Birrell and Pass 2011). They bounded the difference of the expectations of the utility function with a parameter, but the ratio seems to be more natural for DP.

2 Preliminaries

Let $A = \{a_1, a_2, \ldots, a_m\}$ denote a set of $m \geqslant 2$ alternatives. For any $n \in \mathbb{N}$, let $N = \{1, 2, \ldots, n\}$ be a set of voters. For each $j \in N$, the vote of voter j is a linear order $\succ_j \in \mathcal{L}(A)$, where $\mathcal{L}(A)$ denotes the set of all linear orders over A, i.e., all transitive, antireflexive, antisymmetric, and complete binary relations. Let $P = \{\succ_1, \succ_2, \ldots, \succ_n\}$ denote the *(preference) profile*. For each $j \in N$, let P_{-j} denote the profile obtained from P by removing $\succ_j$. A (randomized) voting rule is a mapping $r\colon \mathcal{L}(A)^n \to \Delta(A)$, where $\Delta(A)$ denotes the set of all probability distributions on A.

Given a profile $P \in \mathcal{L}(A)^n$, let $S_P[a, b]$ denote the number of voters who prefer a to b, i.e., $S_P[a, b] = |\{j \in N : a \succ_j b\}|$. Let $w_P[a, b] = S_P[a, b] - S_P[b, a]$ be the *majority margin* of a over b. Then the *weighted majority graph (WMG)* of P can be defined: the vertices of WMG are alternatives in A and there is a directed edge from a to b with weight $w_P[a, b]$ if and only if $w_P[a, b] > 0$. Similarly, letting $U_P[a, b] = \mathrm{Sgn}(w_P[a, b])$, the *unweighted majority graph (UMG)* of P can also be defined: the set of vertices is A and there is an unweighted directed edge from a to b if and only if $U_P[a, b] = 1$, where Sgn denotes the sign function, i.e., $\mathrm{Sgn}(x) = x/|x|$ for all $x \neq 0$ and $\mathrm{Sgn}(0) = 0$. The *Condorcet winner* of P is an alternative $a \in A$, such that $U_P[a, b] = 1$ for all $b \in A \backslash \{a\}$, denoted by $\mathrm{CW}(P)$. Notice that the Condorcet winner is completely determined by the UMG, we also use $\mathrm{CW}(U_P)$ to denote the Condorcet winner claimed by the UMG.

Axioms of voting. A voting rule r satisfies *Condorcet criterion*, if $\mathbb{P}[r(P) = \mathrm{CW}(P)] = 1$ holds for all profile P that $\mathrm{CW}(P)$ exists. The rule r satisfies *Pareto effi-*

ciency, if $\mathbb{P}[r(P) = b] = 0$ for all profile P, where exists $a, b \in A$ that $a \succ_j b$ for all $j \in N$. And r satisfies *absolute monotonicity* (Brandl, Brandt, and Stricker 2018), if $\mathbb{P}[r(P) = a] \leqslant \mathbb{P}[r(P') = a]$ holds for all P, P', such that $P_{-j} = P'_{-j}$, $\succ_j \neq \succ'_j$, and $\succ'_j$ is a pushup of a in $\succ_j$, i.e., $\succ'_j$ raises the position of a in $\succ_j$, and keeps the relative position of other alternatives unchanged. A randomized rule r satisfies *SD-Strategyproofness* (Aziz, Brandt, and Brill 2013), if for all P, P' and $j \in N$ that $P_{-j} = P'_{-j}$ and $\succ_j \neq \succ'_j$, $\sum_{b \succ_j a} \mathbb{P}[r(P) = b] \geqslant \sum_{b \succ_j a} \mathbb{P}[r(P') = b]$, for all $a \in A$ [1]. A voting rule satisfies *lexi-participation* if for all P, P' that $P' = P \backslash \{\succ_j\}$, there does not exist $a \in A$, such that $\mathbb{P}[r(P) = a] < \mathbb{P}[r(P') = a]$ and $\mathbb{P}[r(P) = b] = \mathbb{P}[r(P') = b]$ for all $b \succ_j a$.

Differential privacy (Dwork et al. 2006) is a theoretical guarantee of privacy that requires a function to return similar outputs while receiving similar inputs.

Definition 1 (**Differential privacy).** *A function r with domain $\mathcal{D}$ is ϵ-differentially private (ϵ-DP for short) if for all $O \subseteq \mathrm{Range}(r)$ and $P, P' \in \mathcal{D}$ differing on only one record,*

$$\mathbb{P}[r(P) \in O] \leqslant \mathrm{e}^{\epsilon} \cdot \mathbb{P}[r(P') \in O].$$

In other words, a function r is ϵ-DP, if the ratio between the probabilities for the outputs of any pair of neighboring datasets to be in any given set O must be upper bounded by e^{ϵ}. In the context of social choice, r is a voting rule and

$$\mathcal{D} = \mathcal{L}(A)^* = \mathcal{L}(A) \cup \mathcal{L}(A)^2 \cup \cdots,$$

and P, P' are two profiles differing on only one voter's vote. Under this requirement, the winner of any DP voting rule will not be significantly influenced by any single voter's vote. As a result, any individual's vote will not be revealed by announcing the winner of the voting process.

Notice that Definition 1 does not require the e^{ϵ} upper bound to be tight. The tight upper bound is captured by *exact DP*, formally defined as follows.

Definition 2 (**Exact differential privacy (Dwork 2006)).** *A voting rule r is exact DP (ϵ-eDP for short) if it is ϵ-DP and there does not exist $\epsilon' < \epsilon$ such that r is ϵ'-DP.*

For both DP and eDP, the privacy budget ϵ usually is decided according to the users' demand. For example, iOS 11 requires $\epsilon \leq 43$ and iOS 10 requires $\epsilon \leq 14$ (Orr 2017)[2]. In the next section, we provide upper and lower bounds for the required noise level for any user-defined privacy budget.

3 Differentially Private Condorcet Methods

In this section, we propose a novel class of randomized voting rules. We apply three randomization mechanisms and obtain three classes of voting rules. By analyzing the worst cases, we prove that all of the three rules are differentially private, and our bounds of privacy budget are tight.

[1] In fact, absolute monotonicity and SD-strategyproof are equivalent to the nonperverseness and the strategyproofness in (Gibbard 1977), respectively.

[2] iOS has may have stronger privacy requirement for some specific data types (e.g., $\epsilon \leq 8$ for Safari Auto-play intent detection data) (Apple Inc. 2017).

Mechanism 1: Randomized Condorcet Method

Input: Profile P, Parameter λ, Randomization `Rand`
Output: Winning alternative

Function `Select_Rand(`S, λ`)`**:**
  Get randomized unweighted graph $U^{\mathtt{Rand}}_{\lambda,P}$ with randomized mechanism `Rand`;
  if *There exists Condorcet winner a for $U^{Rand}_{\lambda,P}$* **then**
    return a;
  else
    `Select_Rand(`S,λ`)`;

Function `CM_Rand(`P, λ`)`**:**
  Compute $S_P[a, b]$ for all $a, b \in A$;
  `Select_Rand(`S_P, λ`)`;

As mentioned in Section 2, the existence of Condorcet winner is completely determined by the UMG. In our mechanism, denoted by $\mathsf{CM}^{\mathtt{Rand}}_{\lambda}$, a randomization mechanism `Rand` generates a noisy UMG for the given profile, and the voting rule outputs the Condorcet winner. If the Condorcet winner does not exist, the mechanism will generate another UMG, until the Condorcet winner exists, as shown in Mechanism 1.

Remark. Notice that for each pair of alternatives $a, b \in A$, $U^{\mathtt{Rand}}_{\lambda,P}[a, b]$ and $U^{\mathtt{Rand}}_{\lambda,P}[b, a]$ are determined simultaneously, i.e., $U^{\mathtt{Rand}}_{\lambda,P}[a, b] = 1$, if and only if $U^{\mathtt{Rand}}_{\lambda,P}[b, a] = -1$. Thus, any noisy UMG $U^{\mathtt{Rand}}_{\lambda,P}$ produced in Mechanism 1 claims at most one Condorcet winner. In other words, our mechanism is a well-defined map from $\mathcal{L}(A)^*$ to $\Delta(A)$.

In the randomization process, we adopt three different methods, which are defined as follows.

Definition 3. *Given $\lambda > 0$, the three randomization mechanisms are*

- *Laplace mechanism: Given profile P, for any $a_i, a_j \in A$ that $i < j$, let $\hat{w}_P[a_i, a_j] = w_P[a_i, a_j] + X_{ij}$ for all $a_i, a_j \in A$ and $\hat{w}_P[a_j, a_i] = -\hat{w}_P[a_i, a_j]$, where $X_{ij} \overset{i.i.d}{\sim} \mathrm{Lap}(1/\lambda)$[3]. Under such a mechanism, the noisy UMG is*

$$U^{\mathrm{LAP}}_{\lambda,P}[a, b] = \mathrm{Sgn}(\hat{w}_P[a, b]).$$

- *Exponential mechanism: For profile P,*

$$\mathbb{P}[U^{\mathrm{EXP}}_{\lambda,P}[a, b] = 1] \propto \mathrm{e}^{\lambda \cdot S_P[a,b]/2},$$
$$\mathbb{P}[U^{\mathrm{EXP}}_{\lambda,P}[a, b] = -1] \propto \mathrm{e}^{\lambda \cdot S_P[b,a]/2}.$$

- *Randomized response: For the majority margin w_P of a given profile P, if $w_P[a, b] \neq 0$,*

$$U^{\mathrm{RR}}_{\lambda,P}[a, b] = \begin{cases} \mathrm{Sgn}(w_P[a, b]), & \text{w.p. } \frac{\mathrm{e}^{\lambda}}{1+\mathrm{e}^{\lambda}}, \\ -\mathrm{Sgn}(w_P[a, b]), & \text{w.p. } \frac{1}{1+\mathrm{e}^{\lambda}}. \end{cases}$$

[3] The Laplace distribution with scale parameter $1/\lambda$, of which the probability density function (PDF) is $f_{\lambda}(x) = \frac{\lambda}{2}\mathrm{e}^{-\lambda|x|}$.

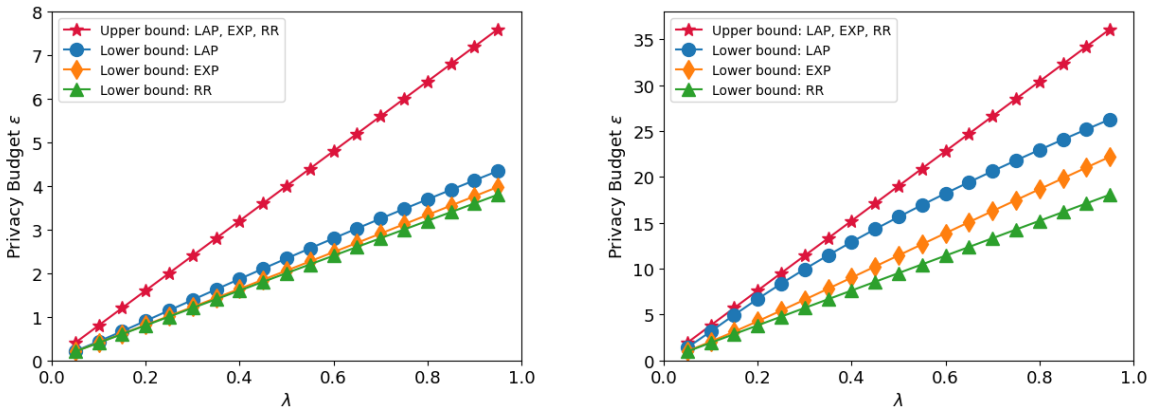

Figure 1: The lower and upper bounds of privacy budget (left: $m = 5$, right: $m = 20$).

If $w_P[a, b] = 0$*, then*

$$\mathbb{P}[U^{\mathrm{RR}}_{\lambda,P}[a, b] = 1] = \mathbb{P}[U^{\mathrm{RR}}_{\lambda,P}[a, b] = -1] = 1/2.$$

The three randomization mechanisms above are denoted by LAP, EXP, and RR, respectively. For each $\mathtt{Rand} \in \{\mathrm{LAP}, \mathrm{EXP}, \mathrm{RR}\}$, the Condorcet winner may not exist for the noisy UMG $U^{\mathrm{Rand}}_{\lambda,P}$. Thus, our mechanism may need to perform the randomization for several times. In fact, for any given profile P, the expected times of randomization is $\exp(\Theta(m))$ (see Appendix A.1). However, such a mechanism with high time complexity can be sampled efficiently, as shown in the following lemma.

Lemma 1. *For any* $\mathtt{Rand} \in \{\mathrm{LAP}, \mathrm{EXP}, \mathrm{RR}\}$ *and* $\lambda > 0$*,* $\mathsf{CM}^{\mathtt{Rand}}_{\lambda}$ *can be sampled as follows:*

- *For any* $P \in \mathcal{L}(A)^*$*,* $\mathsf{CM}^{\mathrm{LAP}}_{\lambda}(P)$ *is a probability distribution in* $\Delta(A)$*, such that for any* $a \in A$*,*

$$\mathbb{P}[\mathsf{CM}^{\mathrm{LAP}}_{\lambda}(P) = a] \propto \prod_{b \neq a} F_\lambda(w_P[a, b]),$$

  where $F_\lambda(x) = \int_{-\infty}^{x} f_\lambda(t)\mathrm{d}t$ *is the cumulative distribution function (CDF) of* $\mathrm{Lap}(1/\lambda)$*.*
- *For any* $P \in \mathcal{L}(A)^*$*,* $\mathsf{CM}^{\mathrm{EXP}}_{\lambda}(P)$ *is a probability distribution in* $\Delta(A)$*, such that for any* $a \in A$*,*

$$\mathbb{P}[\mathsf{CM}^{\mathrm{EXP}}_{\lambda}(P) = a] \propto \prod_{b \neq a} \frac{1}{1 + \mathrm{e}^{-\lambda \cdot w_P[a,b]/2}}.$$

- *For any* $P \in \mathcal{L}(A)^*$*,* $\mathsf{CM}^{\mathrm{RR}}_{\lambda}(P)$ *is a probability distribution in* $\Delta(A)$*, such that for any* $a \in A$*,*

$$\mathbb{P}[\mathsf{CM}^{\mathrm{RR}}_{\lambda}(P) = a] \propto \frac{\mathrm{e}^{\lambda \cdot |B(a)|}}{(1 + \mathrm{e}^{\lambda})^{m-1}},$$

  where $B(a) = \{b \in A : S_P[a, b] > S_P[b, a]\}$*.*

Since there are totally m alternatives, and the value of $\mathbb{P}[\mathsf{CM}^{\mathtt{Rand}}_{\lambda}(P) = a]$ for each $a \in A$ in Lemma 1 can be computed in $O(m)$ time. Therefore, $\mathsf{CM}^{\mathtt{Rand}}_{\lambda}$ can be sampled in $O(m^2)$ time.

Now, we are ready to show the DP bounds of our rules. For simplicity , we use $G_\lambda(x)$ to denote $\mathbb{P}[U^{\mathtt{Rand}}_{\lambda,P}[a, b] = 1]$, where $w_P[a, b] = x$. For example, when $\mathtt{Rand} = \mathrm{LAP}$, $G_\lambda(x) = F_\lambda(x)$; when $\mathtt{Rand} = \mathrm{EXP}$, $G_\lambda(x) = \frac{1}{1+\mathrm{e}^{-\lambda x/2}}$.

Theorem 1. *Given* $\lambda > 0$ *and* $\mathtt{Rand}$*, suppose that* $\mathsf{CM}^{\mathtt{Rand}}_{\lambda}$ *satisfies* ϵ*-eDP. When* $\mathtt{Rand} \in \{\mathrm{LAP}, \mathrm{EXP}\}$

$$\ln\left(\frac{G^{m-1}_\lambda(2) - G^{m-1}_\lambda(-2)}{G_\lambda(2) - G_\lambda(-2)} \cdot \frac{2^{m-2}}{m-1}\right) + (m-1)\lambda \leqslant \epsilon \leqslant 2(m-1)\lambda.$$

When $\mathtt{Rand} = \mathrm{RR}$*,* $(m-1)\lambda \leqslant \epsilon \leqslant 2(m-1)\lambda$*.*

Proof Sketch. For the upper bound, w.l.o.g., we make comparison between the winning probabilities of a_1 under neighboring profile P and P'. According to Lemma 1, we have

$$\mathbb{P}[\mathsf{CM}^{\mathtt{Rand}}_{\lambda}(P) = a_1] = \frac{\prod_{i=2}^{m-1} \mathbb{P}[U^{\mathtt{Rand}}_{\lambda,P}[a_1, a_i] = 1]}{\sum_{j=1}^{m-1} \prod_{i \neq j} \mathbb{P}[U^{\mathtt{Rand}}_{\lambda,P}[a_j, a_i] = 1]}.$$

Then we can prove that for any $a \in A$ and any $\mathtt{Rand}$,

$$\prod_{b \neq a} \mathbb{P}[U^{\mathtt{Rand}}_{\lambda,P}[a, b] = 1] \leqslant \mathrm{e}^{(m-1)\lambda} \cdot \prod_{b \neq a} \mathbb{P}[U^{\mathtt{Rand}}_{\lambda,P'}[a, b] = 1].$$

Further, we have

$$\frac{\sum_{a \in A} \prod_{b \neq a} \mathbb{P}[U^{\mathtt{Rand}}_{\lambda,P}[a, b] = 1]}{\sum_{a \in A} \prod_{b \neq a} \mathbb{P}[U^{\mathtt{Rand}}_{\lambda,P'}[a, b] = 1]} \leqslant \mathrm{e}^{(m-1)\lambda}.$$

As a consequence, for any $\mathtt{Rand} \in \{\mathrm{LAP}, \mathrm{EXP}, \mathrm{RR}\}$,

$$\frac{\mathbb{P}[\mathsf{CM}^{\mathtt{Rand}}_{\lambda}(P) = a_1]}{\mathbb{P}[\mathsf{CM}^{\mathtt{Rand}}_{\lambda}(P') = a_1]} \leqslant \mathrm{e}^{2(m-1)\lambda},$$

which completes the proof of upper bound.

For the lower bound, consider the profile P with $n = 2k$:

- k voters: $a_1 \succ a_2 \succ \cdots \succ a_m$;
- $k-1$ voters: $a_{m-1} \succ a_{m-2} \succ \cdots a_1 \succ a_m$;
- 1 voter: $a_m \succ a_{m-1} \succ \cdots \succ a_1$.

And another profile P':

- $k+1$ voters: $a_1 \succ' a_2 \succ' \cdots \succ' a_m$;
- $k-1$ voters: $a_{m-1} \succ' a_{m-2} \succ' \cdots a_1 \succ' a_m$.

Then the lower bound is $\frac{\mathbb{P}[\mathsf{CM}_\lambda^{\mathtt{Rand}}(P)=a_m]}{\mathbb{P}[\mathsf{CM}_\lambda^{\mathtt{Rand}}(P')=a_m]}$. □

With Theorem 1, we can get the upper and lower bounds of exact privacy budget ϵ for any given λ. The relations between the lower and upper bounds when $m = 5$ and $m = 20$ are shown in Figure 1.

4 Axioms-Privacy Tradeoff

In this section, we analyze our voting rules with axioms mentioned in Section 2. We show that our rules do not satisfy Condorcet criterion and Pareto efficiency. To address these challenges, we explore probabilistic variants of them. Then we discuss the satisfaction to absolute monotonicity, SD-strategyproofness, and lexi-participation.

To begin with, we analyze our voting rules with Condorcet criterion. But unfortunately, $\mathsf{CM}_\lambda^{\mathtt{Rand}}$ does not satisfy Condorcet criterion with any $\mathtt{Rand} \in \{\mathrm{LAP}, \mathrm{EXP}, \mathrm{RR}\}$ and $\lambda \in \mathbb{R}_+$, though it is designed based on the Condorcet method. Intuitively, for any $P \in \mathcal{L}(A)^n$ and any pair of alternatives $a, b \in A$, $\mathbb{P}[U_{\lambda,P}^{\mathtt{Rand}}[a, b] = 1] < 1$. Then

$$\mathbb{P}[\mathsf{CM}_\lambda^{\mathtt{Rand}}(P) = a] \leqslant \prod_{b \in A\backslash\{a\}} \mathbb{P}[U_{\lambda,P}^{\mathtt{Rand}}[a, b] = 1] < 1,$$

even when a is the Condorcet winner. To deal with this, we propose a probabilistic variant of Condorcet criterion, which is shown in the following definition.

Definition 4 (Probabilistic Condorcet criterion). *A randomized voting rule r satisfies probabilistic Condorcet criterion (p-Condorcet) if for every profile P that $\mathrm{CW}(P)$ exists and all $a \in A\backslash\{\mathrm{CW}(P)\}$,*

$$\mathbb{P}[r(P) = \mathrm{CW}(P)] \geqslant \mathbb{P}[r(P) = a].$$

At a high level, Definition 4 is a relaxation of the Condorcet criterion, since it does not always require the voting rule r to select the Condorcet winner. Further, the following theorem holds.

Theorem 2. *For any $\lambda > 0$, $\mathsf{CM}_\lambda^{\mathrm{RR}}$ satisfies p-Condorcet.*

In other words, $\mathsf{CM}_\lambda^{\mathrm{RR}}$ satisfies a weak version of Condorcet criterion for any $\lambda > 0$. However, the results for $\mathsf{CM}_\lambda^{\mathrm{EXP}}$ and $\mathsf{CM}_\lambda^{\mathrm{LAP}}$ are relatively negative.

Proposition 1. *$\mathsf{CM}_{0.5}^{\mathrm{EXP}}$ and $\mathsf{CM}_{0.5}^{\mathrm{LAP}}$ do not satisfy p-Condorcet.*

To measure how much $\mathsf{CM}_\lambda^{\mathrm{EXP}}$ and $\mathsf{CM}_\lambda^{\mathrm{LAP}}$ deviate from p-Condorcet, we further extend the axiom.

Definition 5 (α-Probabilistic Condorcet criterion). *A randomized voting rule r satisfies α-probabilistic Condorcet criterion (α-p-Condorcet) if for every profile P that $\mathrm{CW}(P)$ exists and for all $a \in A\backslash\{\mathrm{CW}(P)\}$,*

$$\mathbb{P}[r(P) = \mathrm{CW}(P)] \geqslant \alpha \cdot \mathbb{P}[r(P) = a].$$

Note that a larger α is more desirable, as α-p-Condorcet is almost equivalent to the standard Condorcet criterion when $\alpha \to \infty$. Especially, it reduces to p-Condorcet when $\alpha = 1$.

For $\mathsf{CM}_\lambda^{\mathrm{EXP}}$ and $\mathsf{CM}_\lambda^{\mathrm{LAP}}$, the following theorem holds.

Theorem 3. *For any $\lambda > 0$,*

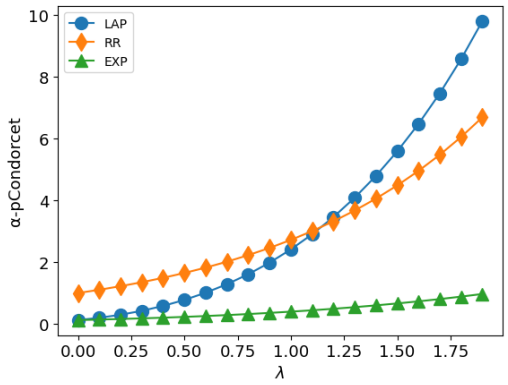

Figure 2: The satisfaction of p-Condorcet with different λ.

- $\mathsf{CM}_\lambda^{\mathrm{EXP}}$ *satisfies* $\frac{1+\mathrm{e}^{\lambda/2}}{\left(1+\mathrm{e}^{-\lambda/2}\right)^{m-1}}$*-p-Condorcet;*
- $\mathsf{CM}_\lambda^{\mathrm{LAP}}$ *satisfies* $2\mathrm{e}^{\lambda}\left(1-\frac{\mathrm{e}^{-\lambda}}{2}\right)^{m-1}$*-p-Condorcet;*
- $\mathsf{CM}_\lambda^{\mathrm{RR}}$ *satisfies* e^{λ}*-p-Condorcet.*

Proof Sketch. Since there is only one profile P, the normalizations in Lemma 1 are not necessarily considered anymore. Here we take $\mathsf{CM}_\lambda^{\mathrm{EXP}}$ as an example. Since $\mathrm{CW}(P) = a$, we have $w_P[a, c] \geqslant 1$, for any $c \in A\backslash\{a\}$, which indicates that $\mathbb{P}[\mathsf{CM}_\lambda^{\mathrm{EXP}}(P) = a]$ has a lower bound. Similarly, $\mathbb{P}[\mathsf{CM}_\lambda^{\mathrm{EXP}}(P) = b]$ has an upper bound for any $b \in A\backslash\{a\}$. Then the lower bound of $\frac{\mathbb{P}[\mathsf{CM}_\lambda^{\mathrm{EXP}}(P)=a]}{\mathbb{P}[\mathsf{CM}_\lambda^{\mathrm{EXP}}(P)=b]}$ can be derived. □

With Theorem 3, we can obtain a more general version of Proposition 1, which is shown as follows.

Proposition 2. *Given $\lambda > 0$, $\mathsf{CM}_\lambda^{\mathrm{EXP}}$ satisfies p-Condorcet when $\frac{\ln(\mathrm{e}^{\lambda/2}+1)}{\ln(\mathrm{e}^{\lambda/2}+1)-\lambda/2}+1 \geqslant m$; $\mathsf{CM}_\lambda^{\mathrm{LAP}}$ satisfies p-Condorcet when $\frac{\lambda+\ln 2}{\ln 2-\ln(2-\mathrm{e}^{-\lambda})}+1 \geqslant m$.*

Notice that both of the LHS of the two inequalities in Proposition 2 are increasing functions of λ which diverge when $\lambda \to \infty$. Thus, for any m, there must exist some λ satisfying the inequalities.

Since the upper and lower bounds of the privacy budget can completely be determined by λ, we use λ to denote the privacy level. Also, we use the parameter α in Definition 4 to denote the level of satisfaction to p-Condorcet, then the tradeoff curves when $m = 5$ are shown in Figure 2.

Similar to the Condorcet criterion, our new class of voting rules do not satisfy Pareto efficiency either. Suppose there is an alternative $b \in A$, which is Pareto dominated by $a \in A$ in profile P, i.e., $a \succ_j b$ for all $j \in N$. Then for any λ and $\mathtt{Rand} \in \{\mathrm{LAP}, \mathrm{EXP}, \mathrm{RR}\}$, we have

$$\mathbb{P}[\mathsf{CM}_\lambda^{\mathtt{Rand}}(P) = b] \leqslant \prod_{c \neq b} \mathbb{P}[U_{\lambda,P}^{\mathtt{Rand}}[b, c] = 1].$$

According to the Definition of LAP, EXP, and RR, $\mathbb{P}[U_{\lambda,P}^{\mathtt{Rand}}(b, c)] > 0$ for all $c \in A$, which indicates that $\mathsf{CM}_\lambda^{\mathtt{Rand}}$ does not satisfy Pareto efficiency. However, a still dominates b in another way. Formally, we have the following definition.

Definition 6 (**Probabilistic Pareto efficiency**). *A randomized voting rule r satisfies probabilistic Pareto efficiency (p-Pareto) if for each pair of alternatives $a, b \in A$ that $a \succ_k b$ holds for all $k \in N$,*

$$\mathbb{P}[r(P) = a] \geqslant \mathbb{P}[r(P) = b].$$

This definition is a relaxation of Pareto efficiency. For our voting rules, the following theorem holds.

Theorem 4. *For any $\lambda > 0$, $\mathsf{CM}_\lambda^{\text{RR}}$, $\mathsf{CM}_\lambda^{\text{EXP}}$, and $\mathsf{CM}_\lambda^{\text{LAP}}$ satisfy p-Pareto.*

Proof Sketch. Let $a, b \in A$ be the pair of alternatives that $a \succ_j b$ holds for any $j \in N$. Then for any $j \in N$ and any $c \in A \backslash \{a, b\}$ such that $b \succ_j c$, we have $a \succ_j c$. Further,

$$\begin{aligned} S[a, c] &= |\{j \in N : a \succ_j c\}| \\ &\geqslant |\{j \in N : b \succ_j c\}| = S[b, c]. \end{aligned}$$

Hence, $w_P[a, c] \geqslant w_P[b, c]$, for all $c \in A \backslash \{a, b\}$. Then the theorem holds due to the monotonicity of CDFs. □

Unlike Condorcet and Pareto, the definition of monotonicity, strategyproofness, and participation are related to two distinct profiles. For monotonicity, we use the notion of absolute monotonicity in (Brandl, Brandt, and Stricker 2018). Intuitively, in Mechanism 1, for any $a \in A$, whenever a voter i switches her preference $\succ_i$ to $\succ_i'$ by lifting a simply, a will be more likely to defeat any $b \in A \backslash \{a\}$ in the one-on-one comparisons. As a consequence, a will be more likely to be the winning alternative in our $\mathsf{CM}_\lambda^{\texttt{Rand}}$. Formally, we have the following theorem.

Theorem 5. *For any $\lambda > 0$, $\mathsf{CM}_\lambda^{\text{RR}}$, $\mathsf{CM}_\lambda^{\text{EXP}}$, and $\mathsf{CM}_\lambda^{\text{LAP}}$ satisfy a-monotonicity.*

Proof Sketch. Let P and P be two profiles in $\mathcal{L}(A)^n$, such that $P_{-j} = P'_{-j}$ and $\succ_j'$ is a pushup of $a \in A$ in $\succ_j$. Then we only need to prove that

$$\prod_{b \neq a} \mathbb{P}[U_{\lambda,P}^{\texttt{Rand}}[a, b] = 1] \geqslant \prod_{b \neq a} \mathbb{P}[U_{\lambda,P'}^{\texttt{Rand}}[a, b] = 1],$$

which is true due to the monotonicity of CDFs. □

Next, we discuss the strategyproofness of $\mathsf{CM}_\lambda^{\texttt{Rand}}$. We adopt the notion of SD-strategyproofness (Aziz, Brandt, and Brill 2013), which implies the absolute monotonicity. However, the results for our rules are not so positive.

Proposition 3. *$\mathsf{CM}_1^{\text{RR}}$, $\mathsf{CM}_1^{\text{EXP}}$, and $\mathsf{CM}_1^{\text{LAP}}$ do not satisfy SD-strategyproofness.*

Similar to Definition 5, we extend the notion of SD-strategyproofness.

Definition 7 (**α-SD-Strategyproofness**). *A voting rule r satisfies α-SD-strategyproofness (α-SD-SP for short) if for all P, P' and $j \in N$, such that $P_{-j} = P'_{-j}$ and $\succ_j \neq \succ_j'$,*

$$\sum_{b \succ_j a} \mathbb{P}[r(P) = b] \geqslant \alpha \cdot \sum_{b \succ_j' a} \mathbb{P}[r(P') = b], \text{ for all } a \in A.$$

Especially, α-SD-strategyproofness reduces to the standard SD-strategyproofness when $\alpha = 1$. For our rules, the following theorem holds.

Theorem 6. *For any $\lambda > 0$, $\mathsf{CM}_\lambda^{\text{RR}}$, $\mathsf{CM}_\lambda^{\text{LAP}}$ and $\mathsf{CM}_\lambda^{\text{EXP}}$ satisfy $\mathrm{e}^{(2-2m)\lambda}$-SD-strategyproofness.*

Proof Sketch. W.l.o.g., for any $\texttt{Rand} \in \{\text{LAP}, \text{EXP}, \text{RR}\}$ and any profiles $P = \{\succ_1, \succ_2, \ldots, \succ_n\}, P' = \{\succ_1', \succ_2, \ldots, \succ_n\}$ and $a \in A$, we have

$$\mathbb{P}[\mathsf{CM}_\lambda^{\texttt{Rand}}(P) = a] \geqslant \mathrm{e}^{(2-2m)\lambda} \cdot \mathbb{P}[\mathsf{CM}_\lambda^{\texttt{Rand}}(P') = a].$$

Therefore, for any $a \in A$,

$$\frac{\sum_{b \succ_1 a} \mathbb{P}[\mathsf{CM}_\lambda^{\texttt{Rand}}(P) = b]}{\sum_{b \succ_1' a} \mathbb{P}[\mathsf{CM}_\lambda^{\texttt{Rand}}(P') = b]} \geqslant \mathrm{e}^{(2-2m)\lambda},$$

which completes the proof. □

Finally, we discuss the participation of our voting rules. We use the notion of lexi-participation, which requires that a participating agent is always no worse off under lexicographical order. In our rules, each participating voter j can benefit herself, since the majority margin $w[a, b]$ for any $a \succ_j b$ will increase due to her vote. Formally, the following theorem holds.

Theorem 7. *For any $\lambda > 0$, $\mathsf{CM}_\lambda^{\text{LAP}}$, $\mathsf{CM}_\lambda^{\text{EXP}}$, and $\mathsf{CM}_\lambda^{\text{RR}}$ satisfy lexi-participation.*

Proof Sketch. On the one hand, we can prove that the winning probability of the top-ranked alternative of the voter i will not decrease after she submits her vote. On the other hand, if there exists any alternative a that

$$\mathbb{P}[\mathsf{CM}_\lambda^{\texttt{Rand}}(P) = a] < \mathbb{P}[\mathsf{CM}_\lambda^{\texttt{Rand}}(P \backslash \{\succ i\}) = a],$$

there will exists another alternative b that $b \succ_i a$, and

$$\mathbb{P}[\mathsf{CM}_\lambda^{\texttt{Rand}}(P) = b] < \mathbb{P}[\mathsf{CM}_\lambda^{\texttt{Rand}}(P \backslash \{\succ i\}) = b],$$

which completes the proof. □

Theorem 7 shows that for any $\texttt{Rand} \in \{\text{LAP}, \text{EXP}, \text{RR}\}$, $\mathsf{CM}_\lambda^{\texttt{Rand}}$ will not harm any participating voter under lexicographical order. Further more, $\mathsf{CM}_\lambda^{\text{LAP}}$ and $\mathsf{CM}_\lambda^{\text{EXP}}$ satisfy a stronger notion, which is defined as follows.

Definition 8 (Strong lexi-participation). *A voting rule r satisfies strong lexi-participation if for all P, P' that $P' = P \backslash \{\succ_j\}$, there exists $a \in A$, such that $\mathbb{P}[r(P) = a] > \mathbb{P}[r(P') = a]$ and $\mathbb{P}[r(P) = b] = \mathbb{P}[r(P') = b]$ for all $b \succ_j a$.*

Intuitively, strong lexi-participation guarantees that each voter can benefit from her vote, while lexi-participation only ensures that each voter will not be harmed by her vote. For $\mathsf{CM}_\lambda^{\text{LAP}}$ and $\mathsf{CM}_\lambda^{\text{EXP}}$, we have the following theorem.

Theorem 8. *For any $\lambda > 0$, $\mathsf{CM}_\lambda^{\text{LAP}}$ and $\mathsf{CM}_\lambda^{\text{EXP}}$ satisfy strong lexi-participation.*

Proof Sketch. When $\texttt{Rand} = \{\text{LAP}, \text{EXP}\}$, for any $\lambda > 0$, any pair of $a, b \in A$, $\mathbb{P}[U_{\lambda,P}^{\texttt{Rand}}[a, b] = 1]$ are strictly increasing functions of $w_P[a, b]$. Thus, the top-ranked alternative of a voter strictly increases when she submits her vote, which completes the proof. □

However, $\mathsf{CM}_\lambda^{\text{RR}}$ does not satisfy strong lexi-participation, since only one vote may be not able to increase the winning probability of any alternative. For example, consider

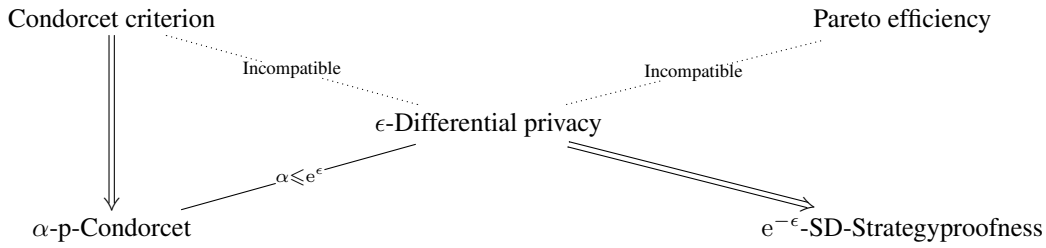

Figure 3: Relations between ϵ-DP and other axioms, where $X \Rightarrow Y$ indicates that X implies Y, a solid line between X and Y indicates that X, Y are compatible with some condition, and a dash line between X and Y means that X, Y are incompatible.

the profile P, where the votes of all n ($n \geqslant 3$) voters are exactly the same, $a_1 \succ a_2 \succ \cdots \succ a_m$. Then for any $i_1 < i_2$, we have $S_P[a_{i_1}, a_{i_2}] = n$ and $S_P[a_{i_2}, a_{i_1}] = 0$. For any P' that $P' = P \backslash \{\succ_j\}$, we have $S_{P'}[a_{i_1}, a_{i_2}] = n - 1$ and $S_{P'}[a_{i_2}, a_{i_1}] = 0$ for any $i_1 < i_2$. Then it follows that $U^{\mathrm{RR}}_{\lambda,P}[a_{i_1}, a_{i_2}] = U^{\mathrm{RR}}_{\lambda,P'}[a_{i_1}, a_{i_2}]$, for all $i_1, i_2 \in \{1, 2, \ldots, m\}$. As a result, we have

$$\mathbb{P}[\mathsf{CM}^{\mathrm{RR}}_{\lambda}(P) = a_i] = \mathbb{P}[\mathsf{CM}^{\mathrm{RR}}_{\lambda}(P') = a_i], \text{ for all } a_i \in A,$$

which indicates that $\mathsf{CM}^{\mathrm{RR}}_{\lambda}$ does not satisfy strong lexiparticipation.

5 Differential Privacy as a Voting Axiom

In Section 4, we explored the tradeoffs between privacy and some voting axioms. In this section, differential privacy is regarded as an axiomatic property of voting rules. The relations between DP and some of the voting axioms are discussed. Our results are summarized in Figure 3.

As proved previously, for any $\mathtt{Rand} \in \{\mathrm{LAP}, \mathrm{EXP}, \mathrm{RR}\}$, $\mathsf{CM}^{\mathtt{Rand}}_{\lambda}$ does not satisfy Condorcet criterion under DP. Furthermore, we can prove that they are incompatible.

Proposition 4. *There is no voting rule r satisfying Condorcet criterion and ϵ-DP for any $\epsilon > 0$.*

Proof Sketch. Suppose there is a voting rule r satisfying Condorcet criterion and ϵ-DP, where $\epsilon > 0$. Then the Condorcet criterion ensures that there exist some profile P and alternative $a \in A$ that $\mathbb{P}[r(P) = a] = 0$. However, all profiles in $\mathcal{L}(A)^n$ are somehow connected by the DP-inequality. As a result, for such a profile P, we have $\mathbb{P}[r(P) = b] = 0$ for each $b \in A$, which leads to a contradiction. □

Similarly, Pareto efficiency is also incompatible with DP, which indicates that the stronger notions of efficiency, e.g., PC-efficiency and SD-efficiency (Brandt 2017) are all incompatible with DP. Formally, we have the following result.

Proposition 5. *There is no voting rule r satisfying Pareto efficiency and ϵ-DP for any $\epsilon > 0$.*

To show the limitation of the tradeoff between Condorcet criterion and DP by any mechanism, we measure the incompatibility between Condorcet criterion and DP using the notion of α-p-Condorcet. The result is shown as follows.

Proposition 6. *There is no voting rule satisfying ϵ-DP and α-p-Condorcet with $\alpha > e^{\epsilon}$.*

Proof. Let P, P' be profiles that $\mathrm{CW}(P) = a$,$\mathrm{CW}(P') = b$, $P_{-j} = P'_{-j}$, and $\succ_j \neq \succ'_j$. Then

$$\mathbb{P}[f(P) = a] \geqslant \alpha \cdot \mathbb{P}[f(P) = b] \geqslant \alpha \cdot e^{\epsilon} \cdot \mathbb{P}[f(P') = b]$$
$$\geqslant \alpha^2 \cdot e^{\epsilon} \cdot \mathbb{P}[f(P') = a] \geqslant \alpha^2 \cdot e^{-2k\epsilon} \cdot \mathbb{P}[f(P) = a].$$

Thus, $\alpha^2 e^{-2\epsilon} \leqslant 1$, i.e., $\alpha \leqslant e^{\epsilon}$. □

The SD-strategyproofness is compatible with DP, as the trivial voting rule, i.e., $\mathbb{P}[r(P) = a] = 1/m$, for all $a \in A$, satisfies SD-strategyproofness and 0-DP. In fact, DP admits a lower bound of satisfaction to strategyproofness. To be more precise, we use the notion of α-SD-strategyproofness.

Then the following proposition holds.

Proposition 7. *Any voting rule satisfying ϵ-DP satisfies $e^{-\epsilon}$-SD-strategyproofness.*

Proof. Suppose that r is a voting rule satisfying ϵ-DP and P, P' are profiles differing on only one voter's ballot. Then, for any $j \in N$ and any $a \in A$, DP indicates that

$$\mathbb{P}[r(P) = a] \leqslant e^{-\epsilon} \cdot \mathbb{P}[r(P') = a].$$

Then we have

$$\sum_{b \succ_i a} \mathbb{P}[r(P) = b] \leqslant e^{-\epsilon} \cdot \sum_{b \succ_i a} \mathbb{P}[r(P') = b],$$

which completes the proof. □

As is shown in Proposition 7, the satisfaction to strategyproofness increases when ϵ decreases. This is quite intuitive, since there is little motivation for an adversary to manipulate the voting process when the outcomes of neighboring datasets are very similar.

6 Conclusion and Future Work

In the paper, we proposed three classes of differentially private Condorcet methods and explored their accuracy-privacy tradeoff and axioms-privacy tradeoff. Further, we investivated the relations between DP and other axioms. For future work, we plan to explore more axioms for randomized voting rules and design new voting rules performing better on satisfaction to the axioms. The design and analysis of DP mechanisms for other social choice problems, such as multi-winner elections, fair division, and participatory budgeting, are also promising directions for future work.

Acknowledgments

ZL acknowledges the National Key R&D Program of China under Grant 2020YFB1005702 for support. AL acknowledges the RPI-IBM AI Horizons scholarship for support. LX acknowledges NSF #1453542 and a Google Research Award for support. YC acknowledges NSFC under Grants 62172016 and 61932001 for support. HW acknowledges NSFC under Grant 61972005 for support.

References

Apple Inc. 2017. Differential Privacy Technical Overview. https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf. Accessed: 2022-08-16.

Aziz, H.; Brandl, F.; and Brandt, F. 2014. On the incompatibility of efficiency and strategyproofness in randomized social choice. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI-14)*, 545–551.

Aziz, H.; Brandl, F.; and Brandt, F. 2015. Universal Pareto dominance and welfare for plausible utility functions. *Journal of Mathematical Economics*, 60: 123–133.

Aziz, H.; Brandt, F.; and Brill, M. 2013. On the tradeoff between economic efficiency and strategy proofness in randomized social choice. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-13)*, 455–462. IFAAMAS.

Bassily, R.; Groce, A.; Katz, J.; and Smith, A. 2013. Coupled-worlds privacy: Exploiting adversarial uncertainty in statistical data privacy. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS-13)*, 439–448. IEEE Computer Society.

Benadè, G.; Gölz, P.; and Procaccia, A. D. 2019. No stratification without representation. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC-19)*, 281–314. ACM.

Birrell, E.; and Pass, R. 2011. Approximately strategy-proof voting. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*, 67–72. AAAI Press.

Brandl, F.; Brandt, F.; and Hofbauer, J. 2015. Incentives for participation and abstention in probabilistic social choice. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS-15)*, 1411–1419. IFAAMAS.

Brandl, F.; Brandt, F.; and Hofbauer, J. 2019. Welfare maximization entices participation. *Games and Economic Behavior*, 114: 308–314.

Brandl, F.; Brandt, F.; and Stricker, C. 2018. An analytical and experimental comparison of maximal lottery schemes. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18)*, 114–120. Morgan Kaufmann.

Brandt, F. 2017. Rolling the Dice: Recent Results in Probabilistic Social Choice. In Endriss, U., ed., *Trends in Computational Social Choice*, chapter 1. AI Access.

Dwork, C. 2006. Differential Privacy. In *33rd International Colloquium on Automata, Languages, and Programming (ICALP-06)*, 1–12. Springer.

Dwork, C.; Kenthapadi, K.; McSherry, F.; Mironov, I.; and Naor, M. 2006. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques (Eurocrypt-06)*, 486–503. Springer.

Flanigan, B.; Gölz, P.; Gupta, A.; Hennig, B.; and Procaccia, A. D. 2021. Fair algorithms for selecting citizens' assemblies. *Nature*, 596(7873): 548–552.

Flanigan, B.; Gölz, P.; Gupta, A.; and Procaccia, A. D. 2020. Neutralizing self-selection bias in sampling for sortition. In *Advances in Neural Information Processing Systems (NeurIPS-20)*, volume 33, 6528–6539. MIT Press.

Gibbard, A. 1977. Manipulation of schemes that mix voting with chance. *Econometrica: Journal of the Econometric Society*, 45(3): 665–681.

Gross, S.; Anshelevich, E.; and Xia, L. 2017. Vote until two of you agree: Mechanisms with small distortion and sample complexity. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI-17)*, 544–550.

Hay, M.; Elagina, L.; and Miklau, G. 2017. Differentially private rank aggregation. In *Proceedings of the 2017 SIAM International Conference on Data Mining (SDM-17)*, 669–677. SIAM.

Hsu, J.; Huang, Z.; Roth, A.; Roughgarden, T.; and Wu, Z. S. 2016. Private matchings and allocations. *SIAM Journal on Computing*, 45(6): 1953–1984.

Kannan, S.; Morgenstern, J.; Rogers, R.; and Roth, A. 2018. Private Pareto optimal exchange. *ACM Transactions on Economics and Computation*, 6(3-4): 1–25.

Kohli, N.; and Laskowski, P. 2018. Epsilon voting: Mechanism design for parameter selection in differential privacy. In *2018 IEEE Symposium on Privacy-Aware Computing (PAC-18)*, 19–30. IEEE.

Lee, D. T. 2015. Efficient, private, and eps-strategyproof elicitation of tournament voting rules. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-15)*, 2026–2032. AAAI Press.

Liu, A.; Lu, Y.; Xia, L.; and Zikas, V. 2020. How private are commonly-used voting rules? In *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI-20)*, volume 124, 629–638. PMLR.

Orr, A. 2017. Google's differential privacy may be better than Apple's. Technical report, the Mac Observer.

Pai, M. M.; and Roth, A. 2013. Privacy and mechanism design. *ACM SIGecom Exchanges*, 12(1): 8–29.

Plott, C. R. 1976. Axiomatic Social Choice Theory: An Overview and Interpretation. *American Journal of Political Science*, 20(3): 511–596.

Procaccia, A. D. 2010. Can approximation circumvent Gibbard-Satterthwaite? In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*, 836–841. AAAI Press.

Shang, S.; Wang, T.; Cuff, P.; and Kulkarni, S. R. 2014. The application of differential privacy for rank aggregation: Privacy and accuracy. In *17th International Conference on Information Fusion (FUSION-14)*, 1–7.

Torra, V. 2019. Random dictatorship for privacy-preserving social choice. *International Journal of Information Security*, 19(4): 1–9.

Walsh, T.; and Xia, L. 2012. Lot-based voting rules. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS-12)*, 603–610. IFAAMAS.

Xiao, D. 2013. Is privacy compatible with truthfulness? In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science (ITCS-13)*, 67–86. ACM.

Yan, Z.; Li, G.; and Liu, J. 2020. Private rank aggregation under local differential privacy. *International Journal of Intelligent Systems*, 35(1): 1492–1519.