
How Private Are Commonly-Used Voting Rules?*

Ao Liu¹, Yun Lu², Lirong Xia¹, Vassilis Zikas²

¹ Rensselaer Polytechnic Institute, liua6@rpi.edu, xial@cs.rpi.edu

² The University of Edinburgh, y.lu-59@sms.ed.ac.uk, vzikas@inf.ed.ac.uk

Abstract

Differential privacy has been widely applied to provide privacy guarantees by adding random noise to the function output. However, it inevitably fails in many high-stakes voting scenarios, where voting rules are required to be deterministic. In this work, we present the first framework for answering the question: “*How private are commonly-used voting rules?*” Our answers are two-fold. First, we show that deterministic voting rules provide sufficient privacy in the sense of *distributional differential privacy (DDP)*. We show that assuming the adversarial observer has uncertainty about individual votes, even publishing the histogram of votes achieves good DDP. Second, we introduce the notion of *exact* privacy to compare the privacy preserved in various commonly-studied voting rules, and obtain dichotomy theorems of exact DDP within a large subset of voting rules called *generalized scoring rules*.

1 INTRODUCTION

Differential privacy (DP) has gained much public attention recently, partly due to its use in the United States 2020 Census. Improving upon ad-hoc privacy techniques

*VZ: Research supported in part by Sunday Group, Inc., and part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via 2019-1902070008. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein. LX: acknowledges NSF #1453542 and #1716333, and ONR #N00014-17-1-2621 for support. AL: acknowledges an IBM AIHN scholarship for support.

that were broken in the previous census [Garfinkel *et al.*, 2018], formal privacy definition like DP are much more suitable for controlling the leakage of sensitive data.

Yet, sensitive data is still published today without necessarily understanding the privacy leakage it incurs. In particular, voting data has been surprisingly accessible. In the US, histograms of votes are revealed per county, and voting and registration tables are released [US Census Bureau, 2019], which include fields like sex, race, age, location, and marital status. This abundance of information has enabled politicians to buy voter profiles from data mining companies to manipulate public opinion [Verini, 2007; Bradshaw and Howard, 2018].

Unfortunately, it is not easy to achieve (differential) privacy for voting. It is insufficient to protect voter registration tables with proven privacy techniques; releasing the election outcome can also be a cause of information leakage. To see how an individual’s vote can be inferred by observing the winner of the election, we consider the following example. Suppose Alice cast a vote in an election, and then the winner is announced. Further suppose that an adversary can accurately estimate other votes from questionnaires or by machine learning from the other voters’ social media, and it turns out these other votes ended up with a tie among the candidates. In this case, the adversary can distinguish Alice’s vote even if he knows nothing about Alice, since Alice must have voted for the winner as the tie-breaker.

The strict definition of differential privacy means the mere *possibility* of the above scenario is a privacy violation. Moreover, ties do occur quite often in real life elections. For example, 9.2% of STV elections on Preflib election data [Mattei and Walsh, 2013] are tied [Wang *et al.*, 2019]. Even if we consider another formal privacy definition that accepts the uncertainty stemming from machine learning methods or low likelihood of ties as helpful in disguising votes, it is unclear how to quantitatively measure the effect of such uncertainty, and how

(or whether) privacy differs for different voting rules.

Motivated by the privacy concern in voting, we focus on the following key question in this paper.

How private are commonly-used voting rules?

The importance of answering this question is both practical and theoretical. On the practical side, minimizing the amount of information leakage from voting rules helps protect against censorship, coercion, and vote buying. On the theoretical side, privacy provides a new angle to comparing voting rules and designing new ones.

A first attempt would be to employ *differential privacy* (DP) [Dwork, 2006], measure of privacy widely-accepted and widely-applied in the cryptographic community. Mathematically, a voting rule \mathbf{M} for $n \in \mathbb{N}$ voters is a mapping $\mathbf{M} : \mathcal{U}^n \rightarrow \mathcal{R}$, where \mathcal{U} is the set of all possible votes; \mathcal{R} is the set of all possible outcomes of voting, e.g. winners or histograms of votes. \mathbf{M} is (ϵ, δ) -*differentially private* if for any pair of *preference profiles* $\vec{X} \in \mathcal{U}^n$ and $\vec{X}' \in \mathcal{U}^n$ that only differ on one vote, and any subset of outcomes $\mathcal{S} \subseteq \mathcal{R}$, the following inequality holds:

$$\Pr [\mathbf{M}(\vec{X}) \in \mathcal{S}] \leq e^\epsilon \Pr [\mathbf{M}(\vec{X}') \in \mathcal{S}] + \delta. \quad (1)$$

Smaller ϵ, δ are desirable as it means the outcome of \mathbf{M} is not affected much by one vote, and thus reveals little about an individual voter. Note in general \mathbf{M} must be randomized to satisfy Inequality (1); indeed [Shang *et al.*, 2014; Lee, 2015; Hay *et al.*, 2017] achieved DP via randomized voting.

Yet most, if not all, voting rules used in high-stakes political elections are deterministic, since randomized voting rules suffer from difficulties in verifying implementation correctness, e.g. the controversy in the 2016 Democratic primary election in Iowa [Clayworth and Noble, 2016]. Unfortunately, the randomness in Inequality (1) comes from the voting rule itself, so deterministic rules cannot achieve DP except with the trivial parameter of $\delta \geq 1$, which always holds (see Example 1 for more details).

1.1 OUR CONTRIBUTIONS

To overcome the critical limitation of DP in high-stakes voting scenarios, we study the privacy of deterministic voting rules using *distributional differential privacy* (DDP) [Bassily *et al.*, 2013], a well-accepted notion of privacy that works for deterministic functions. DDP measures the amount of individual information leakage, assuming the adversary only has uncertain information about voter preferences, for example when using a machine learning algorithm. Our result on the DDP of commonly-used voting rules carries the following encouraging message:

Main Message 1: Many commonly-used voting rules achieve good DDP in natural settings.

More precisely, we focus on a natural DDP setting where the adversary’s information is represented by a set of i.i.d. distribution’s over preference profiles, denoted by $\Delta \subseteq \Pi(\mathcal{U})$, where $\Pi(\mathcal{U})$ is the set of all probability distributions over \mathcal{U} with full support. A voting rule \mathbf{M} ’s DDP is now measured by three parameters $(\epsilon, \delta, \Delta)$. A deterministic function is DDP (Definition 2) if it satisfies an inequality similar to Inequality (1), but now the randomness is replenished by the adversary’s uncertainty about the profile \vec{X} , represented by Δ . Like DP, smaller ϵ and δ in DDP are more desirable.

With DDP, we can quantitatively measure the privacy of the histogram rule **Hist**, which outputs the frequency of each type of vote in the preference profile, in the following Theorem 1. As an immediate consequence, many common voting rules also achieve good privacy.

Theorem 1 (DDP for Hist). *Given any $\mathcal{U} = \{x_1, \dots, x_l\}$ and $\Delta \subseteq \Pi(\mathcal{U})$ with $|\Delta| < \infty$, let $p_{\min} = \min_{\pi \in \Delta, i \leq l} (\pi(x_i))$. For any $n \in \mathbb{N}$ and any $\epsilon \geq 2 \ln \left(1 + \frac{1}{p_{\min} n}\right)$, **Hist** for n voters is $(\epsilon, \delta, \Delta)$ -DDP where $\delta = \exp(-\Omega(np_{\min}[\min(2 \ln(2), \epsilon)]^2))$.*

Theorem 1 states that **Hist** is private with good parameters, as even a small ϵ results in δ that is considered *negligible* in cryptography literature. The winner of many commonly-used voting rules depends only on the outcome of **Hist**, and thus contain (often strictly) less information than **Hist**. Thus, they achieve *at least as good* privacy w.r.t. DDP as simply outputting the histogram.

Next, we highlight that DDP (as well as DP and its variants) parameters only describe loose bounds on privacy—by definition, if a voting rule satisfies $(\epsilon, \delta, \Delta)$ -DDP, it also satisfies $(\epsilon + 0.1, \delta + 0.1, \Delta)$ -DDP. To compare the privacy-preserving capability of voting rules, we introduce the notion of *exact distributional differential privacy* (eDDP), whose parameters describe tight bounds on ϵ and δ . We focus on the $\epsilon = 0$ case as a first step to compare various voting rules with their eDDP in the δ parameter. Our results on the eDDP of commonly-used voting rules carry the following message:

Main Message 2: For many combinations of commonly-used voting rules and Δ , the $(0, \delta, \Delta)$ -eDDP exhibits a dichotomy between $\delta = \Theta(\sqrt{1/n})$ and $\delta = \exp(-\Omega(n))$.

More precisely, we prove the following dichotomy theorem for two candidates $\{a, b\}$ and α -biased majority rules with $\alpha \in (0, 1)$, which chooses a as the winner

iff at least αn out of n votes prefer a .

Theorem 2 (Dichotomy in Exact DDP for α -Majority Rules over Two Candidates, Informal) *Fix two candidates $\{a, b\}$ and $\Delta \subseteq \Pi(\{a, b\})$ with $|\Delta| < \infty$. For any $\alpha \in (0, 1)$, the α -biased majority rule is $(0, \delta, \Delta)$ -eDDP for all n , where δ is either $\Theta(\sqrt{1/n})$, when Δ contains a distribution π with $\pi(a) = \alpha$, or exponentially small otherwise.*

For more than two candidates, we prove the following dichotomy theorem for a large family of voting rules and $\Delta \subseteq \Pi(\mathcal{U})$.

Theorem 3 (Dichotomy in Exact DDP of A Large Class of Voting Rules and Δ , Informal) *For any fixed number of candidates, and any voting rule in a large family, the $(0, \delta, \Delta)$ -eDDP is $\delta = \Theta(\sqrt{1/n})$, when Δ contains the uniform distribution, or $\delta = \exp(-\Omega(n))$, when Δ is finite and does not contain any unstable distributions.*

Intuitively, a distribution π is *unstable* under a voting rule \mathbf{M} if adding small perturbations can cause a different candidate to win (Definition 7). Instead of conducting case-by-case studies of eDDP for commonly-used voting rules, we prove Theorem 3 for a large family of voting rules called *generalized scoring rules* [Xia and Conitzer, 2008] that further satisfy *monotonicity*, *local stability*, and *canceling-out*. We show that many commonly-used voting rules satisfy these conditions (Section 5). We also compute and compare the concrete δ values for small elections (Table 1, Section 6 and Appendix E).

1.2 RELATED WORK

Differential privacy [Dwork, 2006] was recently used to add privacy to voting. Shang *et al.* (2014) applied Gaussian noise to the histogram of linear orders, while Hay *et al.* (2017) used Laplace and Exponential mechanisms applied to specific voting rules. Lee (2015) also developed a method of random selection of votes to achieve differential privacy. One interesting aspect of adding noise to the output that was observed in [Birrell and Pass, 2011; Lee, 2015] is that it enables an approximate strategy-proofness; the idea here is that the added noise dilutes the effect of any individual deviation, thereby making strategies which would slightly perturb the outcome irrelevant. We remark that if one wishes to achieve DP for a large number of voting rules, well-known DP mechanisms (like adding Laplace noise [Dwork *et al.*, 2006]) can be applied to rules in GSR in a straightforward way, by adding noise to each component of the score vector and outputting the winner based on the noised score vector. Our work is different because we focus on exact pri-

vacy of deterministic voting rules.

In our work, we compare deterministic functions by their exact privacy. In differential privacy literature where functions must be randomized, their accuracy, or utility, is used to compare them. A number of works have defined utility as a metric which describes the comparative desirability of ϵ -DP mechanisms. In [McSherry and Talwar, 2007], utility is an arbitrary user-defined function, used in the exponential mechanism. The works of [Blum *et al.*, 2008; Hardt and Talwar, 2010; Bassily and Smith, 2015] define utility in terms of error, where the closer (by some metric) the output of the function, which uses this mechanism to apply noise, is from the desired (deterministic) query's, the higher the utility; the definition of [Ghosh *et al.*, 2009] in addition allows the user to define as a parameter, the prior distribution on the query output. In contrast, our results imply that in the context of distributional differential privacy, voting rules achieve a well-accepted notion of privacy while preserving perfect accuracy, or utility.

1.3 DISCUSSIONS

While DP has been widely applied to measure privacy and has been applied to voting, as we discussed in the Introduction, it fails for deterministic functions such as voting rules in high-stakes elections. This critical limitation motivates our study. To the best of our knowledge, we are the first to illustrate how to measure privacy in high-stakes voting using (e)DDP in a natural setting. We will see that the problem, though challenging, can be solved by our novel *trails* technique. Below we explicitly discuss our conceptual and technical contributions and closely related works. More comprehensive discussions of related work can be found in Appendix A.

Conceptual contributions. Our first conceptual contribution is the application of DDP to deterministic voting rules. As discussed earlier, while previous works add random noise to achieve DP, to the best of our knowledge, no previous studies were done for deterministic voting rules. We note that the *truncated* histogram result of [Bassily *et al.*, 2013] does not suffice, since in general, votes are not removed in an election. Moreover, we prove our results in a simpler definition than DDP; the equivalence of this definition and DDP is proven in Appendix B.1. Our second conceptual contribution is the introduction of *exact DDP*, addressing the issue that parameters of DDP (and other relaxations of DP [Bassily *et al.*, 2013; Groce, 2014; Kasiviswanathan and Smith, 2008; Hall *et al.*, 2012; Duan, 2009; Bhaskar *et al.*, 2011]) describe only upper bounds on privacy. We are not aware of other works that explicitly propose to characterize tight bounds on the pri-

M	Borda	STV	Maximin	Plurality	2-approval
$\delta(n)$	$\frac{1}{\sqrt{1.347n + 0.5263}}$	$\frac{1}{\sqrt{1.495n + 0.02669}}$	$\frac{1}{\sqrt{1.553n + 4.433}}$	$\frac{1}{\sqrt{1.717n - 0.09225}}$	$\frac{1}{\sqrt{1.786n + 0.3536}}$

Table 1: δ values in $(0, \delta, \Delta)$ -eDDP for some commonly-used voting rules under the i.i.d. uniform distribution, $m = 3$ and $n \leq 50$. From left to right, we rank rules from least to most private.

vacy parameters ϵ and δ .

Technical contributions. Our first theorem (Theorem 1) is quite positive, showing the privacy of outputting histograms. Theorem 2 and 3 characterize eDDP in terms of δ values by fixing $\epsilon = 0$. We do so for the two reasons: (1) it is the common convention to compute δ based on a fixed ϵ for DP or DDP; (2) $\epsilon = 0$ is the most informative choice, since Theorem 1 shows that even for small non-zero ϵ , any difference we can observe in the δ of two voting rules is exponentially small—considered negligible in cryptography literature. While our theorems appear similar and related to the dichotomy theorems on the probability of ties in voting [Xia and Conitzer, 2008; Xia, 2015], the definition and mathematical analysis are quite different, and previous techniques do not work for all cases; see more discussions in the proof sketch for Theorem 3. To address the challenge, we developed the *trails* technique, which significantly simplifies calculations.

Generality of our setting. As the first work towards answering our key question, we assume the adversary’s beliefs are modeled by a set of i.i.d. distributions over the votes. A special case is the i.i.d. uniform distribution, which is known as the *impartial culture* assumption in social choice [Georges-Théodule, 1952]. Extending to general (ϵ, δ) , and non-i.i.d. distributions is an important and challenging future direction. Lastly, though our definitions and results are presented in the context of voting for the sake of presentation, they can easily be extended to general applications.

2 PRELIMINARIES

Let $\mathcal{C} = \{c_1, \dots, c_m\}$ be a set of $m \geq 2$ candidates, and $\mathcal{L}(\mathcal{C})$ denote the set of all *linear orders* over \mathcal{C} , that is, the set of all antisymmetric, transitive, and total binary relations. Let \mathcal{U} denote the set of all possible votes. Given $n \in \mathbb{N}$, we let $\vec{X} = (X_1, \dots, X_n) \in \mathcal{U}^n$ denote a collection of n votes called a *preference profile*. Let \mathcal{R} denote the set of outcomes of voting. A (deterministic) voting rule for n voters is a mapping $\mathbf{M} : \mathcal{U}^n \rightarrow \mathcal{R}$.

For example, in the *plurality* rule, $\mathcal{U} = \mathcal{R} = \mathcal{C}$; each voter votes for one favorite candidate, and the winner is the candidate with the most votes. In the *Borda* rule,

$\mathcal{U} = \mathcal{L}(\mathcal{C})$ and $\mathcal{R} = \mathcal{C}$; each voter cast a linear order X over \mathcal{C} , denoted by $c_{i_1} \succ c_{i_2} \succ \dots \succ c_{i_m}$, where $a \succ b$ means that a is preferred over b ; each candidate c gets $m - i$ points in each vote, where i is the rank of c in the vote; the winner is the candidate with the highest total points. A tie-breaking mechanism is used when there are ties in plurality and Borda.

Definition 1 (The histogram rule). *Let $\mathcal{U} = \{x_1, \dots, x_l\}$. For any $n \in \mathbb{N}$, the histogram function, denoted by $\mathbf{Hist} : \mathcal{U}^n \rightarrow \mathbb{N}^l$, takes as input a preference profile $\vec{X} = (X_1, \dots, X_n) \in \mathcal{U}^n$ and outputs a l -dimensional integer vector whose i th component is $|\{j : X_j = x_i, j \in \{1, \dots, n\}\}|$.*

For example, when applied to the setting of the plurality rule, $l = m$ and \mathbf{Hist} outputs the number of votes each candidate receives. When applied to the setting of the Borda rule, $l = m!$ and \mathbf{Hist} outputs the number of occurrences of each linear order.

3 DISTRIBUTIONAL DIFFERENTIAL PRIVACY FOR VOTING

As we discussed, DP is not a suitable notion to analyze nontrivial deterministic voting rules as shown in the following example, which motivates our use of distributional differential privacy (DDP) [Bassily *et al.*, 2013].

Example 1 (DP fails for deterministic voting rules). *Consider the plurality rule for two candidates $\{a, b\}$ and three voters ($n = 3$). We have $\mathcal{U} = \mathcal{R} = \{a, b\}$. In Inequality (1), let $\vec{X} = (a, a, b)$, $\vec{X}' = (b, a, b)$, and $\mathcal{S} = \{a\}$. Then, (1) becomes $1 \leq e^\epsilon \times 0 + \delta$, which means that $\delta \geq 1$.*

At a high level, the DDP of a (deterministic or randomized) function is characterized by three parameters $(\epsilon, \delta, \Delta)$, where ϵ and δ are privacy parameters similar to DP, and Δ is a set describing the adversary’s knowledge about the preference profile. We consider adversaries that can be modeled as $\Delta \subseteq \Pi(\mathcal{U})$, which encodes each of the adversary’s possible uncertainties as a distribution where each vote is i.i.d..

Example 2 (Adversary’s information Δ). *Suppose $\mathcal{U} = \mathcal{R} = \mathcal{C} = \{a, b\}$, and the n votes could be i.i.d. generated from either $\pi_{0.2}$ or $\pi_{0.7}$. Here, for any $\gamma \in [0, 1]$,*

$\pi_\gamma(a) = \gamma$. Then, the adversary’s information is represented by $\Delta = \{\pi_{0.2}, \pi_{0.7}\}$. Say we prove that some voting rule is $(\epsilon = 0.5, \delta = 0.1, \Delta)$ -DDP for the above Δ . Intuitively, this means that the voting rule has privacy $\epsilon = 0.5, \delta = 0.1$, given the adversary’s knowledge can be modeled by any distribution in Δ . We remark that this privacy holds without the need to add noise to the outcome of the election, contrasting with DP.

To simplify presentation, below we will introduce the definition of DDP studied in this paper. In our setting of this paper, our simpler definition is equivalent to the original DDP. More details can be found in Appendix B.1.

Definition 2 (DDP studied in this paper). For any $\Delta \subseteq \Pi(\mathcal{U})$, $\epsilon > 0$, and $\delta > 0$, a voting rule $\mathbf{M} : \mathcal{U}^n \rightarrow \mathcal{R}$ is $(\epsilon, \delta, \Delta)$ -DDP if for every $\pi \in \Delta$, $i \leq n$, $x, x' \in \mathcal{U}$, and $S \subseteq \mathcal{R}$, the following inequality holds.

$$\begin{aligned} & \Pr_{\vec{X} \sim \pi}(\mathbf{M}(\vec{X}) \in S | X_i = x) \\ & \leq e^\epsilon \Pr_{\vec{X} \sim \pi}(\mathbf{M}(\vec{X}) \in S | X_i = x') + \delta, \end{aligned} \quad (2)$$

where $\vec{X} = (X_1, \dots, X_n)$ is a preference profile where each vote is i.i.d. generated from π .

For deterministic \mathbf{M} , the randomness in Inequality (2) comes from the adversary’s incomplete information, captured by Δ . We show that **Hist** satisfies good DDP.

Theorem 1 (DDP of **Hist**, proof in Appendix B.2). Given any $\mathcal{U} = \{x_1, \dots, x_l\}$ and $\Delta \subseteq \Pi(\mathcal{U})$ with $|\Delta| < \infty$, let $p_{\min} = \min_{\pi \in \Delta, i < l}(\pi(x_i))$. For any $n \in \mathbb{N}$ and any $\epsilon \geq 2 \ln(1 + \frac{1}{p_{\min}^n})$, **Hist** for n voters is $(\epsilon, \delta, \Delta)$ -DDP where $\delta = \exp(-\Omega(np_{\min}[\min(2 \ln(2), \epsilon)]^2))$.

As corollary, these privacy parameters of **Hist** automatically apply to all functions that only depend on the output of **Hist**, i.e. most voting rules, or outputting the histogram in addition to the winner as in US presidential elections. This follows immediately from a property of DDP called *immunity to post processing* (see Lemma 3 in Appendix B.2). We note the result is similar to that of [Bassily *et al.*, 2013], but they assume lower-frequency items in the histogram are truncated (which is not the case in general when election results are posted) and describe a less precise δ .

4 EXACT PRIVACY OF VOTING RULES: TWO-CANDIDATE CASE

In this section, we first present the definition of *exact distributional differential privacy* (exact DDP or eDDP), then characterize $(0, \delta, \Delta)$ -eDDP for two candidates under any α -biased majority rule. The proof of this theorem will serve as a toy application of our *trails technique*, useful for proving our main result Theorem 3.

Intuitively, a function has *exact privacy* with parameters ϵ and δ if the function cannot satisfy the privacy definition with strictly better parameters. We remark that this definition can easily be altered to define (ϵ, δ) -*exact differential privacy* (eDP) by omitting Δ .

Definition 3 (Exact Distributional Differential Privacy (eDDP)). A voting rule \mathbf{M} is $(\epsilon, \delta, \Delta)$ -Exact Distributional Differential Privacy (eDDP) if it is $(\epsilon, \delta, \Delta)$ -DDP and there does not exist $(\epsilon' \leq \epsilon, \delta' < \delta)$ nor $(\epsilon' < \epsilon, \delta' \leq \delta)$ such that \mathbf{M} is $(\epsilon', \delta', \Delta)$ -DDP.

The α -biased majority rule, denoted by \mathbf{M}_α , over two candidates (a, b) outputs a as the winner if at least α fraction of votes prefer a over b . An example of this type of voting rule is *supermajority*, used in government decisions around the world.

Theorem 2 (Exact DDP for Majority Rules, full proof in Appendix C.2). Fix two candidates $\{a, b\}$ and $\Delta \subseteq \Pi(\{a, b\})$ with $|\Delta| < \infty$. For any $\alpha \in (0, 1)$, the α -biased majority rule is $(0, \delta, \Delta)$ -eDDP for all n , where

$$\delta = \max_{p=\pi(a): \pi \in \Delta} \Theta \left(\sqrt{\frac{1}{n}} \left[\left(\frac{p}{\alpha} \right)^\alpha \left(\frac{1-p}{1-\alpha} \right)^{1-\alpha} \right]^n \right).$$

In particular, $\delta = \Theta(\sqrt{1/n})$ if there exists $\pi \in \Delta$ with $\pi(a) = \alpha$; otherwise $\delta = \exp(-\Omega(n))$.

In the following subsections, we will present our *trails* technique for analyzing DDP in voting, followed by a proof sketch of Theorem 2 using the trails technique.

4.1 OUR TOOL TO ANALYZE PRIVACY: TRAILS TECHNIQUE

Let us describe the trails technique using a simple, toy example: suppose there are two candidates $\{a, b\}$, and $n = 5$ votes. Let \mathbf{M} be the majority rule where ties are broken in favor of a , i.e. $\alpha = 0.5$. We want to compute $(0, \delta, \Delta)$ -eDDP of \mathbf{M} for any $\Delta \subseteq \Pi(\{a, b\})$. In light of Definitions 2 and 3, we have:

$$\begin{aligned} \delta = \max_{S, x, x', i, \pi \in \Delta} & \left[\Pr_{\vec{X} \sim \pi}(\mathbf{M}(\vec{X}) \in S | X_i = x) \right. \\ & \left. - \Pr_{\vec{X} \sim \pi}(\mathbf{M}(\vec{X}) \in S | X_i = x') \right]. \end{aligned} \quad (3)$$

Now, the majority rule is *anonymous*, that is, the identity of the voter is irrelevant and it chooses the winner only based on the histogram of votes. We can thus write $\mathbf{M} = \mathbf{f} \circ \mathbf{Hist}$, where $t = (t_a, t_b)$ and $\mathbf{f}(t)$ outputs a if $t_a \geq t_b$ and outputs b otherwise. Then, Equation (3) can be rewritten with probabilities over histograms, which is

easier to compute (below, $\vec{X} \sim \pi$ is implicit).

$$\begin{aligned} \delta &= \max_{\mathcal{S}, x, x', i, \pi \in \Delta} \left[\Pr(\mathbf{f}(\mathbf{Hist}(\vec{X})) \in \mathcal{S} | X_i = x) \right. \\ &\quad \left. - \Pr(\mathbf{f}(\mathbf{Hist}(\vec{X})) \in \mathcal{S} | X_i = x') \right] \\ &= \max_{\mathcal{S}, x, x', i, \pi \in \Delta} \left[\sum_{t: \mathbf{f}(t) \in \mathcal{S}} \Pr(\mathbf{Hist}(\vec{X}) = t | X_i = x) \right. \\ &\quad \left. - \sum_{t: \mathbf{f}(t) \in \mathcal{S}} \Pr(\mathbf{Hist}(\vec{X}) = t | X_i = x') \right]. \end{aligned} \quad (4)$$

For example, if $\mathcal{S} = \{a\}$, then $\mathbb{T} \equiv \{t: \mathbf{f}(t) \in \mathcal{S}\} = \{(5, 0), (4, 1), (3, 2)\}$ is an example of what we call a *trail*. Intuitively, a trail \mathbb{T} is a set of histograms *consecutive* in the sense that, starting from some t , we can list exactly the elements of \mathbb{T} by iteratively subtracting 1 from and adding 1 to two components of t , respectively. We see that \mathbb{T} can be listed in such a way, starting from entry $\text{Enter}(\mathbb{T}) = (5, 0)$ and ending at exit $\text{Exit}(\mathbb{T}) = (3, 2)$, by iteratively subtracting from the first component and adding to the second component of $(5, 0)$ (we say the *direction* of \mathbb{T} is $(1, 2)$). See Figure 1.

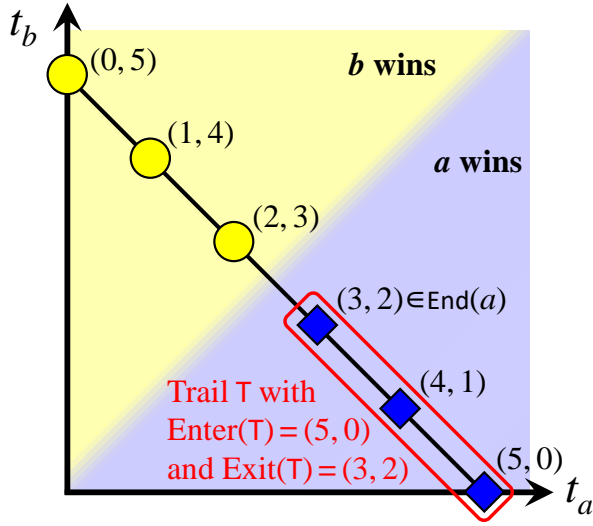


Figure 1: A trail for two candidates. A graph of number of votes for candidate a ($= t_a$) versus votes for candidate b ($= t_b$). Each point in the line is a histogram where the total number of votes is $n = 5$. The set $\{(5, 0), (4, 1), (3, 2)\}$ forms a trail. We denote by $\text{End}(a)$ (used in the proof of Theorem 3) the set of histograms which are exits of trails where a is the winner. In this example $\text{End}(a) = \{(3, 2)\}$.

We now give intuition for our key Lemma 1 presented below using this example. Suppose in Equation (4) the maximizing \mathcal{S} is $\{a\}$ (so that $\{t: \mathbf{f}(t) \in \mathcal{S}\} = \mathbb{T}$), $x = a$,

and $x' = b$. Then, for any i , and any $\pi \in \Delta$:

$$\begin{aligned} \delta &= \sum_{t \in \{(5,0), (4,1), (3,2)\}} \Pr(\mathbf{Hist}(\vec{X}) = t | X_i = a) \\ &\quad - \sum_{t \in \{(5,0), (4,1), (3,2)\}} \Pr(\mathbf{Hist}(\vec{X}) = t | X_i = b). \end{aligned}$$

The core of Lemma 1 is the observation that when votes are independent (e.g. when $\Delta \subseteq \Pi(\{a, b\})$), then for all $t = (t_a, t_b)$ such that $t_a > 0$, the following holds

$$\begin{aligned} &\Pr(\mathbf{Hist}(\vec{X}) = (t_a, t_b) | X_i = a) \\ &= \Pr(\mathbf{Hist}(\vec{X}) = (t_a - 1, t_b + 1) | X_i = b). \end{aligned}$$

In light of this, $\Pr(\mathbf{Hist}(\vec{X}) = (5, 0) | X_i = a)$ cancels out with $\Pr(\mathbf{Hist}(\vec{X}) = (4, 1) | X_i = b)$, and $\Pr(\mathbf{Hist}(\vec{X}) = (4, 1) | X_i = a)$ cancels out with $\Pr(\mathbf{Hist}(\vec{X}) = (3, 2) | X_i = b)$. This leaves

$$\begin{aligned} \delta &= \Pr(\mathbf{Hist}(\vec{X}) = (3, 2) = \text{Exit}(\mathbb{T}) | X_i = a) \\ &\quad - \Pr(\mathbf{Hist}(\vec{X}) = (5, 0) = \text{Enter}(\mathbb{T}) | X_i = b). \end{aligned}$$

We note that here $\Pr(\mathbf{Hist}(X) = \text{Enter}(\mathbb{T}) | X_i = b) = 0$, but this does not hold generally for all trails for $m \geq 2$. This calculation can be extended to the more general Lemma 1 below. Before that, let us formally define trails. For any histogram $t = (t_1, \dots, t_l) \in \mathbb{N}^l$, any $z \in \mathbb{Z}$ and $j \leq l$, we let $(t_1, \dots, t_l) + zx_j$ denote the histogram $(t_1, \dots, t_j + z, \dots, t_l)$.

Definition 4 (Trails). Given a pair of indices (j, k) where $j \neq k$, a histogram t , and a length q , we define the trail $\mathbb{T}_{t, x_j, x_k, q} = \{t - zx_j + zx_k : 0 \leq z \leq q\}$, where (j, k) is called the *direction* of the trail, t is then the *entry* of this trail, also denoted by $\text{Enter}(\mathbb{T}_{t, x_j, x_k, q})$, and $t - qx_j + qx_k$ is called the *exit* of the trail, denoted by $\text{Exit}(\mathbb{T}_{t, x_j, x_k, q})$.

Alternatively, a trail \mathbb{T} can be defined by just its entry and exit.

Lemma 1. Let \mathbb{T} be a trail with direction (j, k) , and let $\pi \in \Pi(\mathcal{U})$. For any $i, x_j, x_k \in \mathcal{U}$, we have:

$$\begin{aligned} &\Pr_{\vec{X} \sim \pi}(\mathbf{Hist}(\vec{X}) \in \mathbb{T} | X_i = x_j) \\ &\quad - \Pr_{\vec{X} \sim \pi}(\mathbf{Hist}(\vec{X}) \in \mathbb{T} | X_i = x_k) \\ &= \Pr_{\vec{X} \sim \pi}(\mathbf{Hist}(\vec{X}) = \text{Exit}(\mathbb{T}) | X_i = x_j) \\ &\quad - \Pr_{\vec{X} \sim \pi}(\mathbf{Hist}(\vec{X}) = \text{Enter}(\mathbb{T}) | X_i = x_k). \end{aligned}$$

Proof. Fix distribution π over n votes, where each vote is independently distributed. For $\vec{X} \sim \pi$, denote X_{-i} as the random variable \vec{X} but without the i th vote. The equality in the lemma comes from the simple observation

that when votes are independently distributed, for any histogram $t \in \mathbb{N}^l$ and any $j \in [l]$

$$\Pr_{\vec{X} \sim \pi} (\mathbf{Hist}(\vec{X}) = t | X_i = x_j) = \Pr_{\vec{X} \sim \pi} (\mathbf{Hist}(X_{-i}) = t - x_j)$$

(Below, $\vec{X} \sim \pi$ is implicit). Let q be the length of the trail. For any $0 \leq z < q$, let $t_z = \text{Enter}(\mathbb{T}) - zx_j + zx_k$. Then,

$$\begin{aligned} & \Pr(\mathbf{Hist}(\vec{X}) = t_z | X_i = x_j) \\ &= \Pr(\mathbf{Hist}(X_{-i}) = t_z - x_j) \\ &= \Pr(\mathbf{Hist}(\vec{X}) = t_z - x_j + x_k | X_i = x_k) \\ &= \Pr(\mathbf{Hist}(\vec{X}) = t_{z+1} | X_i = x_k). \end{aligned}$$

In other words,

$$\begin{aligned} & \Pr(\mathbf{Hist}(\vec{X}) \in \mathbb{T} | X_i = x_j) \\ & - \Pr(\mathbf{Hist}(\vec{X}) \in \mathbb{T} | X_i = x_k) \\ &= \Pr(\mathbf{Hist}(\vec{X}) = t_q | X_i = x_j) \\ & - \Pr(\mathbf{Hist}(\vec{X}) = t_0 | X_i = x_k) \\ & + \sum_{0 \leq z < q} \left(\Pr(\mathbf{Hist}(\vec{X}) = t_z | X_i = x_j) \right. \\ & \quad \left. - \Pr(\mathbf{Hist}(\vec{X}) = t_{z+1} | X_i = x_k) \right) \\ &= \Pr(\mathbf{Hist}(\vec{X}) = t_q | X_i = x_j) \\ & - \Pr(\mathbf{Hist}(\vec{X}) = t_0 | X_i = x_k) \end{aligned}$$

(Every term in the summation of differences cancels out.)

$$\begin{aligned} &= \Pr(\mathbf{Hist}(\vec{X}) = \text{Exit}(\mathbb{T}) | X_i = x_j) \\ & - \Pr(\mathbf{Hist}(\vec{X}) = \text{Enter}(\mathbb{T}) | X_i = x_k) \end{aligned}$$

□

Remark. In this subsection's example, no matter the \mathcal{S} , the set $\{t: \mathbf{f}(t) \in \mathcal{S}\}$ forms one single trail, but this does not hold in general. Instead, to prove our main theorem we will partition this set into multiple trails, and apply Lemma 1 to simplify probabilities over each trail.

4.2 A SIMPLE APPLICATION OF TRAILS TECHNIQUE: PROOF OF THEOREM 2

Proof. [Proof sketch for Theorem 2, see Appendix C.2 for the full proof]. For any $\pi \in \Delta$, let $p = \pi(a)$. Let trails $\mathbb{T}_a = \{t: t = (k, n-k), k \geq \alpha n\}$ and $\mathbb{T}_b = \{t: t = (k, n-k), k < \alpha n\}$. It follows that any histogram in \mathbb{T}_a results in a being the winner, and any in \mathbb{T}_b results in b as the winner. Also, Equation (4) implies we should *not* consider $\mathcal{S} = \{a, b\}$ nor $\mathcal{S} = \emptyset$ as otherwise $\delta = 0$ (the default lower bound on δ). Thus, we only consider $\mathcal{S} = \{a\}$ (when winner is a , corresponding to trail \mathbb{T}_a) or $\mathcal{S} = \{b\}$ (trail \mathbb{T}_b). Then Equation (4) becomes (we disregard the value of i since votes are

i.i.d.):

$$\begin{aligned} \delta &= \max_{j \in \{a, b\}, x, x'} \left[\Pr_{\vec{X} \sim \pi} (\mathbf{Hist}(\vec{X}) \in \mathbb{T}_j | X_i = x) \right. \\ & \quad \left. - \Pr_{\vec{X} \sim \pi} (\mathbf{Hist}(\vec{X}) \in \mathbb{T}_j | X_i = x') \right] \quad (\text{Equation (4)}) \\ &= \max_{j \in \{a, b\}, x, x'} \left[\Pr(\mathbf{Hist}(\vec{X}) = \text{Exit}(\mathbb{T}_j) | X_i = x) \right. \\ & \quad \left. - \Pr(\mathbf{Hist}(\vec{X}) = \text{Enter}(\mathbb{T}_j) | X_i = x') \right]. \quad (\text{Lemma 1}) \end{aligned}$$

We first discuss $\mathcal{S} = \{a\}$ whose corresponding trail \mathbb{T}_a starts at $\text{Enter}(\mathbb{T}_a) = (n, 0)$ and exits at $\text{Exit}(\mathbb{T}_a) = (\lceil \alpha n \rceil, \lfloor (1 - \alpha)n \rfloor)$. Here, $x = a$ and $x' = b$ maximize δ . Then,

$$\begin{aligned} & \Pr(\mathbf{Hist}(\vec{X}) = \text{Enter}(\mathbb{T}_a) | X_i = b) \\ &= \Pr(\mathbf{Hist}(\vec{X}) = (n, 0) | X_i = b) = 0, \end{aligned}$$

and

$$\begin{aligned} & \Pr(\mathbf{Hist}(\vec{X}) = \text{Exit}(\mathbb{T}_a) | X_i = a) \\ &= \Theta \left(\sqrt{\frac{1}{n}} \left[\left(\frac{p}{\alpha} \right)^\alpha \left(\frac{1-p}{1-\alpha} \right)^{1-\alpha} \right]^n \right). \end{aligned}$$

The case for $\mathcal{S} = \{b\}$ is similar and Theorem 2 follows by maximizing δ over $\pi \in \Delta$. □

5 EXACT PRIVACY OF VOTING RULES: GENERAL CASE

The main result of this section, Theorem 3, characterizes $(0, \delta, \Delta)$ -exact DDP of *generalized scoring rules (GSR)* for arbitrary number of candidates, defined below. The main message is that the characterization holds for commonly-used voting rules (Corollary 1). Therefore, to get the main message, a reader can skip the technical descriptions and definitions below to Corollary 1.

Definition 5 (Generalized Scoring Rules (GSR) [Xia and Conitzer, 2008]). *A Generalized Scoring Rule (GSR) is defined by a number $K \in \mathbb{N}$ and two functions $\mathbf{f}: \mathcal{L}(\mathcal{C}) \rightarrow \mathbb{R}^K$ and \mathbf{g} , which maps weak orders over the set $\{1, \dots, K\}$ to \mathcal{C} . Given a vote $V \in \mathcal{L}(\mathcal{C})$, $\mathbf{f}(V)$ is the generalized score vector of V . Given a profile P , we call $\mathbf{f}(P) = \sum_{V \in P} \mathbf{f}(V)$ the score. Then, the winner is given by $\mathbf{g}(\text{Ord}(\mathbf{f}(P)))$, where Ord outputs the weak order of the K components in $\mathbf{f}(P)$.*

We say that a rule is a GSR if it can be described by some \mathbf{f}, \mathbf{g} as above. Most popular voting rules (i.e., Borda, Plurality, k -approval and ranked pairs) are GSRs. See Example 3 and Example 4 for \mathbf{f}, \mathbf{g} for plurality rule and majority rule. The domain of GSRs can be naturally extended to *weighted* profiles, where each type of vote is weighted by a real number, due to the linearity of \mathbf{f} .

Example 3. *The simplest example of a GSR is plurality. This is the voting rule where each voter chooses exactly one candidate, and the candidate with the most votes is the winner. Here, K is equal to the number of candidates m . Suppose V is a vote (linear order over candidates) where the top candidate is x_i . The function \mathbf{f} would map V to a vector $\mathbf{f}(V) = (0, \dots, 0, 1, 0, \dots, 0)$ where the 1 is at position i in the vector. Then, $\mathbf{f}(P)$ is exactly the histogram counting the number of times each candidate is ranked at the top of a vote. Finally, the function \mathbf{g} chooses the winner.*

We now define a set of properties of GSRs to present our characterization of eDDP in Theorem 3.

Definition 6 (Canceling-out, Monotonicity, and Local stability). *A voting rule \mathbf{M} satisfies canceling-out if for any profile \vec{X} , adding a copy of every ranking does not change the winner. That is, $\mathbf{M}(\vec{X}) = \mathbf{M}(\vec{X} \cup \mathcal{L}(C))$.*

A voting rule satisfies monotonicity one cannot prevent a candidate from winning by raising its ranking in a vote while maintaining the order of other candidates.

A voting rule \mathbf{M} satisfies local stability if there exist locally stable profile. A profile \vec{X}^ is locally stable (to \mathbf{M}), if there exists a candidate a , a ranking W , and another ranking V that is obtained from W by raising the position of a without changing the order of other candidates, such that for any \vec{X}' in the γ neighborhood of \vec{X}^* in terms of L_∞ norm, we have (1) $\mathbf{M}(\vec{X}') \neq a$, and (2) the winner is a when all W votes in \vec{X}' becomes V votes.*

Definition 7 (Unstable distributions). *Given a GSR \mathbf{M} , a distribution π over \mathcal{U} is unstable, if for any $\epsilon > 0$, there exists \vec{p} and \vec{q} with $\|\vec{p}\|_2 = \|\vec{q}\|_2 < \epsilon$, such that $\mathbf{M}(\pi + \vec{q}) \neq \mathbf{M}(\pi + \vec{p})^1$, where $\|\cdot\|_2$ is the ℓ_2 -norm.*

Theorem 3 (Dichotomy of Exact DDP for GSR, full proof in Appendix D.1). *Fix $m \geq 2$ and $\Delta \subseteq \Pi(\mathcal{L}(C))$ with $|\Delta| < \infty$. For any n , any GSR \mathbf{M} that satisfies monotonicity, local stability, and canceling-out is $(0, \delta, \Delta)$ -DDP, where δ is*

- $\Theta(\sqrt{1/n})$, if Δ contains the uniform distribution over $\mathcal{L}(C)$, or
- $\exp(-\Omega(n))$, if Δ does not contain any unstable distribution.

Proof sketch for Theorem 3. (See Appendix D.1 for the full proof) We first prove the $\delta = \exp[-\Omega(n)]$ case. Recalling the proof of Theorem 2, we know that δ is closely related to the probability of $\text{End}(a)$ for some $a \in \mathcal{C}$. It turns out that this is also the case for any GSR \mathbf{M}

¹We slightly abuse notation— $\mathbf{M}(\pi)$ denotes the output of \mathbf{M} when the voters cast fractional votes according to π .

that also satisfies monotonicity. Applying our trails technique, we have

$$\delta \leq \max_a \sum_{P \in \text{End}(a)} \Pr(P - V),$$

where V is a vote s.t. there exists vote W with $\mathbf{M}(P - V + W) \neq a$. Thus, we know δ is upper bounded by the probability of all profiles $(P - V)$ “close” to a tie of voting rule r . For any unstable distribution π , we can prove that the center of π is reasonably “far” away from any profile in $\text{End}(a)$ (or “far” away from any ties). Then, the exponential upper bound follows after Chernoff bound and union bound. The proof for this part is similar to the analysis of probabilities of tied profiles as in [Xia and Conitzer, 2008].

We now move on to the $\delta = \Theta(\sqrt{1/n})$ case. The upper bound $O(\sqrt{1/n})$ also derived from the trails technique’s result: $\delta \leq \max_a \sum_{P \in \text{End}(a)} \Pr(P - V)$. General framework of our proof is similar with the $\delta = \exp[-\Omega(n)]$ case. Since adding any vote to a uniform profile results in a new winner, we know the uniform distribution of preferences is always an unstable distribution when requirements in Theorem 3 are met. Thus, we can prove that the center of the profiles’ distribution (multinomial distribution in $m!$ -dimensional space) is “close” to a tie. Then, we apply Stirling’s formula to each trails and give an upper bounds to $\Pr(P - V)$ for profiles $P \in \text{End}(a)$.

For the lower bound $\Omega(\sqrt{1/n})$, canceling-out and locally stability are used to construct a “good” subset of profiles. At a high level, canceling-out ensures that the constructed subset is large enough, and locally stability ensures the trails constructed from the selected subset is long enough. Our subset is contracted by certain profiles with $O(\sqrt{n})$ distance² from the center of profile distribution in the direction of local stable profile. Giving a lower bound to the $\Pr(P - V)$ for any profile P in our selected subset is the most non-trivial part of this proof and is quite different from the proof in [Xia and Conitzer, 2008]. Unlike the profiles P in our selected subset of profiles, $P - V$ do not necessarily concentrated in a specific region in the space of profiles. Here, we use a non-i.i.d. version of Lindeberg-Levy central limit theorem [Greene, 2003] to analyze the multinomial distribution of $m!$ kinds of votes. \square

Next, we use a simple example of majority rule to show the results in Theorem 3 matches the 2-candidate results in Section 4. In the following example, we also provide the intuitions on how to describe voting rules in the language of GSR.

²we use ℓ_2 distance in the $m!$ -dimensional space of profile.

Example 4 (Example of Definition 5 and Theorem 3). Let $\mathcal{U} = \mathcal{R} = \mathcal{C} = \{c_1, c_2\}$, $V = [c_1 \succ c_2]$, and $W = [c_2 \succ c_1]$. For the majority rule with $\alpha = 0.5$, we have $\mathbf{f}(V) = (1, 0)$ and $\mathbf{f}(W) = (0, 1)$. Then, the winner is chosen according to \mathbf{g} corresponding to the largest component in $\mathbf{f}(P)$. Recalling our definition of unstable distribution, we know $(\frac{1}{2}, \frac{1}{2})$ is the only unstable distribution for 2-candidate majority rule. This is the intuitive reason behind $\delta = \Theta(\sqrt{1/n})$ when $\pi = (\frac{1}{2}, \frac{1}{2})$ for both Theorem 3 and Theorem 2 (when $\alpha = 0.5$). For any other $\pi \neq (\frac{1}{2}, \frac{1}{2})$, these two theorems result in $\delta = \exp[-\Omega(n)]$. We note that while Theorem 3 covers more voting rules, Theorem 2 is a more fine-grained result for two candidates.

Corollary 1. *Plurality, veto, k-approval, Borda, maximin, Copeland, Bucklin, Ranked Pairs, Schulze (see e.g. [Xia and Conitzer, 2008]) are $(0, \Theta(1/\sqrt{n}), \Delta)$ -eDDP when Δ contains the uniform distribution.*

Proof. As shown in Definition 6, *cancelling-out* and *monotonicity* are very natural properties of most voting rules. These two properties can be easily checked according to the definitions of voting rules discussed in Corollary 1. In the next proposition, we prove a more generalized version of Corollary 1 for *local stability*, which indicate a large subset of the voting rules can satisfy all properties required by Theorem 3.

Proposition 1. *All positional scoring rules and all Condorcet consistent and monotonic rules satisfy the property of local stability.*

Proof. Let s_i to denote the score of the i -th candidate ($f(P)$ in definition 5). Suppose $s_1 = \dots = s_l > s_{l+1}$. We let $V = [a \succ c_1 \succ c_{l-1} \succ b \succ \text{others}]$ and $W = [c_1 \succ c_{l-1} \succ b \succ a \succ \text{others}]$. Let M be the permutation $c_1 \rightarrow c_2 \rightarrow \dots c_{m-2} \rightarrow c_1$. Let $V_1 = [a \succ b \succ \text{others}]$ and $V_2 = [b \succ a \succ \text{others}]$. Let $P' = \bigcup_{i=1}^{m-2} M^i(V_1) \cup M^i(V_2)$. Let $P^* = 2P' \cup \{V, W\}$. It follows that a and b are the only two candidates tied in the first place in P^* . Therefore, there exists ϵ to satisfy the condition in local stability.

The same profile can be used to prove the local stability of all Condorcet consistent and monotonic rules. \square

Then, Corollary 1 follows by combining the results for all three properties. \square

Another commonly-used GSR called STV does not satisfy monotonicity, which means that Theorem 3 does not apply. However, empirical results (Section E) suggest that STV is likely also $(0, \Theta(1/\sqrt{n}), \Delta)$ -eDDP for this distribution.

6 CONCRETE ESTIMATION OF THE PRIVACY PARAMETERS

We present an example of computing concrete estimates of $(0, \delta, \Delta)$ -exact DDP values for several GSRs. For this example, we let $\Delta = \{\pi\}$ such that $\pi \in \Pi(\{x_1, x_2, x_3\})$ and $\pi(x_i) = \pi(x_j) = 1/3$ (i.e., votes are i.i.d. and uniform). We generated these concrete estimates via exhaustive search over possible profiles for 3 candidates and $n \leq 50$ votes, and computing the δ values exactly for each n . Since we know that $\delta = \Theta(1/\sqrt{n})$, we fit these values to $\delta(n) = \frac{1}{\sqrt{an+b}}$ via linear regression. We rank voting rules from most to least private. The larger the a , the smaller the δ value and thus more private:

2-approval \triangleright Plurality \triangleright Maximin \triangleright STV \triangleright Borda

We showed in Table 1 (Section 1, also see Table 2 in Appendix E for more information) the fitted δ curves. Figure 2 shows the comparison between Plurality, Borda, and STV voting rules w.r.t. their δ values in $(0, \delta, \Delta)$ -eDDP, when fitted to $\delta(n) = \frac{1}{\sqrt{an+b}}$.




Figure 2: The δ values in $(0, \delta, \Delta)$ -eDDP for Borda, STV, and plurality in our concrete estimates.

7 SUMMARY AND FUTURE WORK

We address the limitation of DP in deterministic voting rules by introducing and characterizing (exact) DDP for voting rules, leading to an encouraging message about the good privacy of commonly-studied voting rules and a framework to compare them w.r.t. eDDP. There are many directions for future work. An immediate open question for theoretical study is to extend our studies to general (ϵ, δ) , and non-i.i.d. distributions, as well as to other high-stakes social choice procedures such as matching and resource allocation. On the practical side, it could be informative to study the eDDP of other data that is often published during an election, such as demographic information, and interpret their consequences.

References

- [Bassily and Smith, 2015] Raef Bassily and Adam Smith. Local, Private, Efficient Protocols for Succinct Histograms. *STOC*, 2015.
- [Bassily *et al.*, 2013] Raef Bassily, Adam Groce, Jonathan Katz, and Adam Smith. Coupled-worlds privacy: Exploiting adversarial uncertainty in statistical data privacy. *FOCS*, 2013.
- [Bhaskar *et al.*, 2011] Raghav Bhaskar, Abhishek Bhowmick, Vipul Goyal, Srivatsan Laxman, and Abhradeep Thakurta. Noiseless Database Privacy. *Asiacrypt*, 7073, 2011.
- [Birrell and Pass, 2011] Eleanor Birrell and Rafael Pass. Approximately strategy-proof voting. *IJCAI*, 2011.
- [Blum *et al.*, 2008] Avrim Blum, Katrina Ligett, and Aaron Roth. A learning theory approach to noninteractive database privacy. *STOC*, 2008.
- [Bradshaw and Howard, 2018] Samantha Bradshaw and Philip N Howard. Challenging truth and trust: A global inventory of organized social media manipulation. *The Computational Propaganda Project*, 2018.
- [Caragiannis *et al.*, 2014] Ioannis Caragiannis, Ariel D. Procaccia, and Nisarg Shah. Modal Ranking: A Uniquely Robust Voting Rule. In *AAAI*, 2014.
- [Clayworth and Noble, 2016] Jason Clayworth and Jason Noble. Iowa caucus coin flip count unknown. *OnPolitics*, 1:2016, 2016.
- [Duan, 2009] Yitao Duan. Privacy without Noise. *CIKM*, 2009.
- [Dwork and Roth, 2014] Cynthia Dwork and Aaron Roth. The Algorithmic Foundations of Differential Privacy, 2014.
- [Dwork *et al.*, 2006] Cynthia Dwork, F. McSherry, Kobbi Nisim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. *TCC*, 2006.
- [Dwork, 2006] Cynthia Dwork. Differential Privacy. *ICALP*, 2006.
- [Garfinkel *et al.*, 2018] Simson Garfinkel, John M Abowd, and Christian Martindale. Understanding database reconstruction attacks on public data. *Queue*, 2018.
- [Georges-Théodule, 1952] Guilbaud Georges-Théodule. Les théories de l'intérêt général et le problème logique de l'agrégation [theories of the general interest and the logical problem of aggregation]. *Economie appliquée*, 1952.
- [Ghosh *et al.*, 2009] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. Universally utility-maximizing privacy mechanisms. *STOC*, 2009.
- [Greene, 2003] William H Greene. *Econometric analysis*. Pearson Education India, 2003.
- [Groce, 2014] Adam Groce. New Notions and Mechanisms for Statistical Privacy, PhD Thesis, 2014.
- [Hall *et al.*, 2012] Rob Hall, Alessandro Rinaldo, and Larry Wasserman. Random Differential Privacy. *Journal of Privacy and Confidentiality*, 2012.
- [Hardt and Talwar, 2010] Moritz Hardt and Kunal Talwar. On the Geometry of Differential Privacy. *STOC*, 2010.
- [Hay *et al.*, 2017] M. Hay, L. Elagina, and G. Miklau. Differentially private rank aggregation. *SDM*, 2017.
- [Horn and Johnson, 1990] Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 1990.
- [Kasiviswanathan and Smith, 2008] Shiva Prasad Kasiviswanathan and Adam Smith. A note on differential privacy: Defining resistance to arbitrary side information. *CoRR abs/0803.3946*, 2008.
- [Lee, 2015] David T. Lee. Efficient, private, and e-strategy proof elicitation of tournament voting rules. *IJCAI*, 2015.
- [Leung and Lui, 2012a] Samantha Leung and Edward Lui. Bayesian mechanism design with efficiency, privacy, and approximate truthfulness. *International Workshop on Internet and Network Economics*,
- [Leung and Lui, 2012b] Samantha Leung and Edward Lui. Bayesian mechanism design with efficiency, privacy, and approximate truthfulness. *WINE*, 2012.
- [Mattei and Walsh, 2013] Nicholas Mattei and Toby Walsh. PrefLib: A Library of Preference Data. In *Algorithmic Decision Theory*, Lecture Notes in Artificial Intelligence, 2013.
- [McSherry and Talwar, 2007] Frank McSherry and Kunal Talwar. Mechanism Design via Differential Privacy. *FOCS*, 2007.
- [Shang *et al.*, 2014] Shang Shang, Tiance Wang, Paul Cuff, and Sanjeev Kulkarni. The Application of Differential Privacy for Rank Aggregation: Privacy and Accuracy. *Information Fusion*, 2014.
- [US Census Bureau, 2019] US Census Bureau. Voting and registration tables, Jun 2019.
- [Varah, 1975] James M Varah. A lower bound for the smallest singular value of a matrix. *Linear Algebra and its Applications*, 1975.
- [Verini, 2007] James Verini. Big brother inc. *Vanity Fair*, Dec 2007.
- [Wang *et al.*, 2019] Jun Wang, Sujoy Sikdar, Tyler Shepherd, Zhibing Zhao, Chunheng Jiang, and Lirong Xia. Practical Algorithms for STV and Ranked Pairs with Parallel Universes Tiebreaking. In *AAAI*, 2019.
- [Xia and Conitzer, 2008] Lirong Xia and Vincent Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Electronic Commerce*, 2008.
- [Xia and Conitzer, 2009] Lirong Xia and Vincent Conitzer. Finite Local Consistency Characterizes Generalized Scoring Rules. *IJCAI*, 2009.
- [Xia, 2015] Lirong Xia. Generalized Decision Scoring Rules: Statistical, Computational, and Axiomatic Properties. *EC*, 2015.