
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.

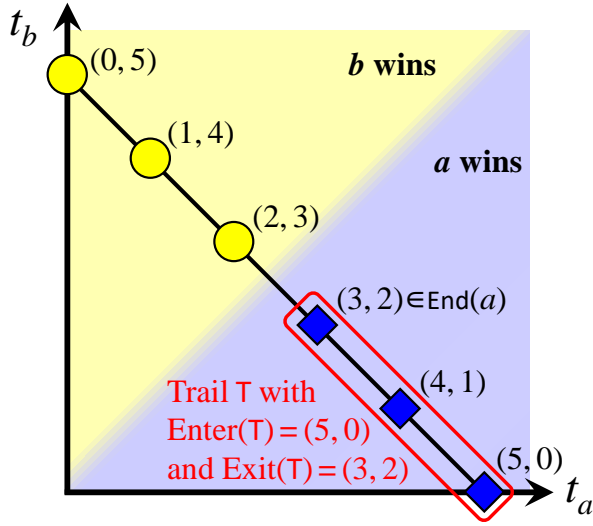


Figure 1: A trail for two candidates. A graph of number of votes for candidate a ($= t_a$) versus votes for candidate b ($= t_b$). Each point in the line is a histogram where the total number of votes is $n = 5$. The set $\{(5, 0), (4, 1), (3, 2)\}$ forms a trail. We denote by $\text{End}(a)$ (used in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

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

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

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

□

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 $\delta(n) = \frac{1}{\sqrt{an+b}}$.

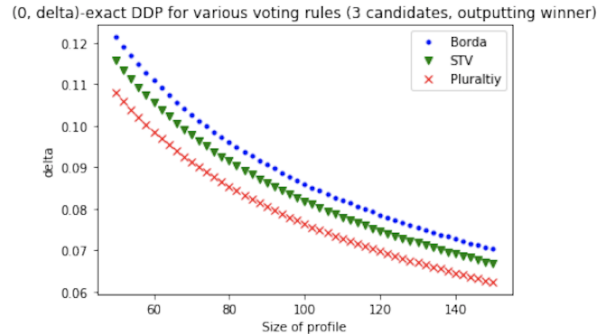


Figure 2: The δ values in $(0, \delta, \Delta)$ -eDDP for Borda, STV, and plurality in our concrete estimates.

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.

A Additional Related Literature

The first works on DP described how one can create mechanisms for answering standard statistical queries on a database (e.g., number of records with some property or histograms) in a way that satisfies the DP definition. This ignited a vast and rapidly evolving line of research on extending the set of mechanisms and achieving different DP guarantees—we refer the reader to [Dwork and Roth, 2014] for an (already outdated) survey—to a rich literature of relaxations to the definition, e.g., [Bhaskar *et al.*, 2011; Leung and Lui, 2012a; Duan, 2009; Bassily *et al.*, 2013], that capture among others, noiseless versions of privacy, as well as works studying the trade-offs between privacy and utility of various mechanisms [McSherry and Talwar, 2007; Blum *et al.*, 2008; Hardt and Talwar, 2010; Bassily and Smith, 2015; Ghosh *et al.*, 2009].

Generalized Scoring Rules (GSRs) is a class of voting rules that include many commonly studied voting rules, such as Plurality, Borda, Copeland, Maximin, and STV [Xia and Conitzer, 2008]. It has been shown that for any GSR the probability for a group of manipulators to be able to change the winner has a phase transition [Xia and Conitzer, 2009; Mossel *et al.*, 2013]. An axiomatic characterization of GSRs is given in [Xia and Conitzer, 2009]. The most robust GSR with respect to a large class of statistical models has been characterized [Caragiannis *et al.*, 2014]. Recently GSRs have been extended to an arbitrary decision space, for example to choose a set of winners or rankings over candidates [Xia, 2015].

Relaxations to Differential Privacy and Noiseless Functions Relaxations to differential privacy have been proposed to allow functions with less to no noise to achieve a DP-style notion of privacy. Kasiviswanathan and Smith [Kasiviswanathan and Smith, 2008] formally proved that differential privacy holds in presence of arbitrary adversarial information, and formulated a Bayesian definition of differential privacy which makes adversarial information explicit. Hall *et al.* [Hall *et al.*, 2012] suggested adding noise to only certain values (such as low-count components in histograms) to achieve a relaxed notion of Random Differential Privacy with higher accuracy. Taking advantage of (assumed) inherent randomness in the database, several works have also put forward DP-style definitions which allow for noiseless functions. Duan [Duan, 2009] showed that sum queries of databases with i.i.d. rows can be outputted without noise. Bhaskar *et al.* [Bhaskar *et al.*, 2011] introduced Noiseless Privacy for database distributions with i.i.d. rows, whose parameters depend on how far the query is from a function which only depends on a subset of the database. Motivated by Bayesian mechanism design, Leung and Lui [Leung and Lui, 2012b], suggested noiseless sum queries and introduced Bayesian differential privacy for database distributions with independent rows, where the auxiliary information is some number of revealed rows.

These ideas were generalized and extended by Bassily *et al.* who introduced *distributional differential privacy* (DDP) [Bassily *et al.*, 2013; Groce, 2014]. Informally, given a distribution (X, Z) , where X is the adversary’s uncertainty in the database distribution and Z is a parameter used for proving composition theorems (i.e. computing DDP when outputting results from two functions that are both DDP with some parameters), we say a function \mathbf{M} is $(\epsilon, \delta, \Delta = \{(X, Z)\})$ -DDP if its output distribution $\mathbf{M}(X)|Z$ can be simulated by a simulator that is given the database missing one row. In these works, noiseless functions which have been shown to satisfy DDP are exact sums, truncated histograms, and *stable* functions where with large probability, the output is the same given neighboring databases.

B Distributional Differential Privacy for Voting

B.1 Equivalence of our DDP definition and that of [Bassily *et al.*, 2013]

For completeness we present the DDP definition of [Bassily *et al.*, 2013]. However, this definition is harder to work with, and harder to explain conceptually. Thus, we choose to present Definition 2. Below, we will show that in our setting, these two definitions are equivalent.

Definition 8 (Distributional Differential Privacy (DDP) [Bassily *et al.*, 2013]). *A function $\mathbf{M} : \mathcal{U}^* \rightarrow \mathcal{R}$ is $(\epsilon, \delta, \Delta)$ -distributional differentially private (DDP) if there is a simulator \mathbf{Sim} such that for all $D = (\pi, Z) \in \Delta$, $\vec{X} \sim \pi$, for all i , $(x, z) \in \text{Supp}(X_i, Z)$ (where X_i denotes the random variable that is the i th component of \vec{X} , X_{-i} denotes the r.v. that is \vec{X} without the i th component, and $\text{Supp}(\cdot)$ denotes the support of a distribution), and all sets $\mathcal{S} \subseteq \mathcal{R}$,*

$$\Pr_{\vec{X} \sim \pi} (\mathbf{M}(\vec{X}) \in \mathcal{S} | X_i = x, Z = z) \leq e^\epsilon \Pr_{\vec{X} \sim \pi} (\mathbf{Sim}(X_{-i}) \in \mathcal{S} | X_i = x, Z = z) + \delta$$

and

$$\Pr_{\vec{X} \sim \pi} (\mathbf{Sim}(X_{-i}) \in \mathcal{S} | X_i = x, Z = z) \leq e^\epsilon \Pr_{\vec{X} \sim \pi} (\mathbf{M}(\vec{X}) \in \mathcal{S} | X_i = x, Z = z) + \delta$$

Lemma 2 (Equivalence of definitions). *For any \mathcal{U} , let $\Delta \subseteq \Pi(\mathcal{U})$ and $\Delta' = (\Delta, Z = \emptyset)$ (where Z is a parameter in the [Bassily et al., 2013] definition). Suppose \mathbf{M} is $(\epsilon, \delta, \Delta')$ -(simulation-based) DDP [Bassily et al., 2013], then \mathbf{M} is $(2\epsilon, (1 + e^\epsilon)\delta, \Delta)$ -DDP for our Definition 2. Conversely, if \mathbf{M} is $(\epsilon, \delta, \Delta)$ -DDP for Definition 2 then \mathbf{M} satisfies $(\epsilon, \delta, \Delta')$ -(simulation-based) DDP.*

Lemma 2. We prove the first statement, that is, if \mathbf{M} is $(\epsilon, \delta, \Delta')$ -(simulation-based) DDP [Bassily et al., 2013], then \mathbf{M} is $(2\epsilon, (1 + e^\epsilon)\delta, \Delta)$ -DDP of Definition 2.

By the definition of \mathbf{M} being $(\epsilon, \delta, \Delta')$ -(simulation-based) DDP, the simulator \mathbf{Sim} has to satisfy the below inequalities for any $(\pi, Z) \in \Delta'$, any i , and $x \in \mathbf{Supp}(X_i)$ (for $\vec{X} \sim \pi$). With $Z = \emptyset$, we can write the inequalities in the DDP definition without Z as

$$\Pr_{\vec{X} \sim \pi} (\mathbf{M}(\vec{X}) \in \mathcal{S} | X_i = x) \leq e^\epsilon \Pr_{\vec{X} \sim \pi} (\mathbf{Sim}(X_{-i}) \in \mathcal{S} | X_i = x) + \delta$$

$$\Pr_{\vec{X} \sim \pi} (\mathbf{Sim}(X_{-i}) \in \mathcal{S} | X_i = x) \leq e^\epsilon \Pr_{\vec{X} \sim \pi} (\mathbf{M}(\vec{X}) \in \mathcal{S} | X_i = x) + \delta \quad (5)$$

(We make $\vec{X} \sim \pi$ implicit to ease presentation.) Now consider any $x' \in \mathbf{Supp}(X_i)$, possibly different from the x above. By the definition of DDP, the inequalities should also hold for x' , i.e.

$$\Pr(\mathbf{M}(\vec{X}) \in \mathcal{S} | X_i = x') \leq e^\epsilon \Pr(\mathbf{Sim}(X_{-i}) | X_i = x') + \delta$$

Since the simulator is not given i th entry of the database, its output does not depend on the value of the i th row. Moreover, if database rows are independent, the distributions $X_{-i} | X_i = x' = X_{-i} | X_i = x$. Thus $\Pr(\mathbf{Sim}(X_{-i}) | X_i = x') = \Pr(\mathbf{Sim}(X_{-i}) \in \mathcal{S} | X_i = x)$. So,

$$\Pr(\mathbf{M}(\vec{X}) \in \mathcal{S} | X_i = x') \leq e^\epsilon \Pr(\mathbf{Sim}(X_{-i}) \in \mathcal{S} | X_i = x) + \delta$$

$$\Pr(\mathbf{M}(\vec{X}) \in \mathcal{S} | X_i = x') \leq e^\epsilon (e^\epsilon \Pr(\mathbf{M}(\vec{X}) \in \mathcal{S} | X_i = x) + \delta) + \delta \text{ (By Equation 5 above.)}$$

$$\Pr(\mathbf{M}(\vec{X}) \in \mathcal{S} | X_i = x') \leq e^{2\epsilon} \Pr(\mathbf{M}(X) \in \mathcal{S} | X_i = x) + e^\epsilon \delta + \delta$$

Thus, we have shown that for all $x, x' \in \mathbf{Supp}(X_i)$ (and all i),

$$\Pr(\mathbf{M}(X) \in \mathcal{S} | X_i = x') \leq e^{2\epsilon} \Pr(\mathbf{M}(X) \in \mathcal{S} | X_i = x) + (e^\epsilon + 1)\delta$$

So, \mathbf{M} is $(2\epsilon, (1 + e^\epsilon)\delta, \Delta)$ -DDP, proving the first statement.

We now prove the second statement. That is, if \mathbf{M} is $(\epsilon, \delta, \Delta)$ -DDP of Definition 2 then \mathbf{M} is $(\epsilon, \delta, \Delta')$ -(simulation-based) DDP. To do so, we define the simulator \mathbf{Sim} to be the algorithm which inserts any $x' \in \mathbf{Supp}(X_i)$ to the missing i th row of the database, and apply \mathbf{M} to the result. By independence of rows, $\Pr(\mathbf{Sim}(X_{-i}) | X_i = x) = \Pr(\mathbf{Sim}(X_{-i}) | X_i = x')$ by our definition of \mathbf{Sim} , equal to $\Pr(\mathbf{M}(X) | X_i = x')$. Then, for any $X \in \Delta$, i , and $x, x' \in \mathbf{Supp}(X_i)$,

$$\Pr(\mathbf{Sim}(X_{-i}) \in \mathcal{S} | X_i = x) = \Pr(\mathbf{M}(X) \in \mathcal{S} | X_i = x') \leq e^\epsilon \Pr(\mathbf{M}(X) \in \mathcal{S} | X_i = x) + \delta$$

by inequality of Definition 2. This proves the second statement. \square

B.2 Proof of Theorem 1: DDP of Histogram

Theorem 1 (DDP of 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 $\epsilon \geq 2 \ln(1 + \frac{1}{p_{\min} n})$, \mathbf{Hist} for n voters is $(\epsilon, \delta, \Delta)$ -DDP where $\delta = \exp(-\Omega(np_{\min}[\min(2 \ln(2), \epsilon)]^2))$.*

Proof. At a high level, the proof is similar to Theorem 8 of [Leung and Lui, 2012b].

Fix $\pi \in \Delta$. Since votes are i.i.d. and all $i \in [n]$ are equivalent, we simplify $\Pr_{\vec{X} \sim \pi}(\mathbf{Hist}(\vec{X}) \in \mathcal{S} | X_i = x)$ as $\Pr(\mathbf{Hist}(x, X_{-1}) \in \mathcal{S})$, where X_{-1} refers to \vec{X} without the first vote.

We need to show that for all $x_i, x_j \in \{x_1, \dots, x_l\}$, and all $\mathcal{S} \subseteq \mathbb{N}^l$:

$$\Pr(\mathbf{Hist}(x_i, X_{-1}) \in \mathcal{S}) \leq e^\epsilon \Pr(\mathbf{Hist}(x_j, X_{-1}) \in \mathcal{S}) + \delta$$

We observe that for any set \mathcal{B} and x :

$$\Pr(\mathbf{Hist}(x, X_{-1}) \in \mathcal{S}) = \Pr(\mathbf{Hist}(x, X_{-1}) \in \mathcal{S} \cap \bar{\mathcal{B}}) + \Pr(\mathbf{Hist}(x, X_{-1}) \in \mathcal{S} \cap \mathcal{B}) \quad (6)$$

$$\leq \Pr(\mathbf{Hist}(x, X_{-1}) \in \mathcal{S} \cap \bar{\mathcal{B}}) + \Pr(\mathbf{Hist}(x, X_{-1}) \in \mathcal{B}) \quad (7)$$

Let \mathcal{B} be the set of all histogram $t \in \mathbb{N}^l$ where $t_i > p_i(n-1)e^{\epsilon/2}$ and $t_j < p_j(n-1)e^{-\epsilon/2}$. Fix a choice of $\epsilon > 2 \ln(1 + \frac{1}{p_{\min} n})$. We claim that for $\delta = \exp(\Omega(np_{\min}(\min(2 \ln(2), \epsilon))^2))$, the following hold:

$$\text{Claim 1: } \Pr(\mathbf{Hist}(x_i, X_{-1}) \in \mathcal{S} \cap \bar{\mathcal{B}}) \leq e^\epsilon \Pr(\mathbf{Hist}(x_j, X_{-1}) \in \mathcal{S} \cap \bar{\mathcal{B}})$$

$$\text{Claim 2: } \Pr(\mathbf{Hist}(x_i, X_{-1}) \in \mathcal{B}) \leq \delta$$

If both claims are true, then by Inequality (7),

$$\begin{aligned} \Pr(\mathbf{Hist}(x_i, X_{-1}) \in \mathcal{S}) &\leq \Pr(\mathbf{Hist}(x_i, X_{-1}) \in \mathcal{S} \cap \bar{\mathcal{B}}) + \Pr(\mathbf{Hist}(x_i, X_{-1}) \in \mathcal{B}) \\ &\leq e^\epsilon \Pr(\mathbf{Hist}(x_j, X_{-1}) \in \mathcal{S} \cap \bar{\mathcal{B}}) + \delta \\ &\leq e^\epsilon \Pr(\mathbf{Hist}(x_j, X_{-1}) \in \mathcal{S}) + \delta \end{aligned}$$

which proves the theorem. Below we will prove both claims.

Claim 1 proof:

Since all entries in random variable X_{-1} are i.i.d., the random variable

$\mathbf{Hist}(X_{-1})$ which outputs the histogram of the database has distribution equal to the multinomial distribution on $n-1$ trials and l events:

$$\Pr(\mathbf{Hist}(X_{-1}) = (t_1, \dots, t_l)) = \frac{(n-1)!}{t_1! \dots t_l!} p_1^{t_1} \dots p_l^{t_l}$$

where t_i is the count of entries with the value x_i and p_i is the probability for an entry to have the value x_i .

Thus,

$$\Pr(\mathbf{Hist}(x_i, X_{-1}) = (t_1, \dots, t_l)) = \frac{(n-1)!}{t_1! \dots (t_i-1)! \dots t_l!} p_1^{t_1} \dots p_i^{t_i-1} \dots p_l^{t_l}$$

and

$$\Pr(\mathbf{Hist}(x_j, X_{-1}) = (t_1, \dots, t_l)) = \frac{(n-1)!}{t_1! \dots (t_j-1)! \dots t_l!} p_1^{t_1} \dots p_i^{t_j-1} \dots p_l^{t_l}$$

So, for every $t = (t_1, \dots, t_l) \in \mathcal{S} \cap \bar{\mathcal{B}}$:

$$\begin{aligned} \frac{\Pr(\mathbf{Hist}(x_i, X_{-1}) = t)}{\Pr(\mathbf{Hist}(x_j, X_{-1}) = t)} &= \frac{\frac{(n-1)!}{t_1! \dots (t_i-1)! \dots t_l!} p_1^{t_1} \dots p_i^{t_i-1} \dots p_l^{t_l}}{\frac{(n-1)!}{t_1! \dots (t_j-1)! \dots t_l!} p_1^{t_1} \dots p_i^{t_j-1} \dots p_l^{t_l}} \\ &= \frac{t_i p_j}{p_i t_j} \end{aligned}$$

$$\begin{aligned}
&= \frac{t_i}{p_i(n-1)} \frac{p_j(n-1)}{t_j} \\
&\quad \text{By definition of } \mathcal{B}, t_i > p_i(n-1)e^{\epsilon/2} \text{ or } t_j < p_j(n-1)e^{-\epsilon/2}, \\
&\quad \text{so } t \in \overline{\mathcal{B}} \text{ has } t_i \leq p_i(n-1)e^{\epsilon/2} \text{ and } t_j \geq p_j(n-1)e^{-\epsilon/2} \\
&\leq \frac{p_i(n-1)e^{\epsilon/2}}{p_i(n-1)} \frac{p_j(n-1)}{p_j(n-1)e^{-\epsilon/2}} = e^{\epsilon/2} \times e^{\epsilon/2} = e^\epsilon
\end{aligned}$$

This proves Claim 1.

Claim 2 proof: Recall \mathcal{B} is the set of all histogram $t \in \mathbb{N}^l$ where $t_i > p_i(n-1)e^{\epsilon/2}$ and $t_j < p_j(n-1)e^{-\epsilon/2}$. For any $i \in \{1, \dots, l\}$ let $\mathbf{Hist}(x, X_{-1})_i$ denote i th component of the random variable $\mathbf{Hist}(x, X_{-1})$.

$$\begin{aligned}
&\Pr(\mathbf{Hist}(x_i, X_{-1}) \in \mathcal{B}) \\
&= \Pr(\mathbf{Hist}(x_i, X_{-1})_i > p_i(n-1)e^{\epsilon/2} \text{ or } \mathbf{Hist}(x_i, X_{-1})_j < p_j(n-1)e^{-\epsilon/2}) \\
&\leq \Pr(\mathbf{Hist}(x_i, X_{-1})_i > p_i(n-1)e^{\epsilon/2}) + \Pr(\mathbf{Hist}(x_i, X_{-1})_j < p_j(n-1)e^{-\epsilon/2}) \quad (\text{By union bound}) \\
&= \Pr(1 + \mathbf{Bin}(n-1, p_i) > p_i(n-1)e^{\epsilon/2}) + \Pr(\mathbf{Bin}(n-1, p_j) < p_j(n-1)e^{-\epsilon/2}) \\
&\quad (\text{Where } \mathbf{Bin}(n, p) \text{ denotes binomial r.v. with } n \text{ trials and success probability } p) \\
&= \Pr(\mathbf{Bin}(n-1, p_i) > p_i(n-1)(e^{\epsilon/2} - (p_i(n-1))^{-1})) \\
&\quad + \Pr(\mathbf{Bin}(n-1, p_j) < p_j(n-1)e^{-\epsilon/2})
\end{aligned}$$

The random variable $\mathbf{Bin}(n-1, p_i)$ has mean $\mu = p_i(n-1)$. When

$$2 \ln\left(1 + \frac{1}{p_i(n-1)}\right) < 2 \ln\left(1 + \frac{1}{p_{\min}(n-1)}\right) < \epsilon \leq 2 \ln(2) < 2 \ln\left(2 + \frac{1}{p_i(n-1)}\right)$$

we have $0 < \beta = e^{\epsilon/2} - (p_i(n-1))^{-1} - 1 < 1$. By Chernoff bound,

$$\begin{aligned}
\Pr(\mathbf{Bin}(n-1, p_i) > (1 + \beta)\mu) &\leq e^{-\mu\beta^2/3} \\
&= \exp(-\Omega(p_i(n-1)(e^{\epsilon/2} - (p_i(n-1))^{-1} - 1)^2)) \\
&= \exp(-\Omega(p_i n \epsilon^2))
\end{aligned}$$

The random variable $\mathbf{Bin}(n-1, p_j)$ has mean $\mu = p_j(n-1)$. By Chernoff bound, for any $0 < \beta = 1 - e^{-\epsilon/2} < 1$ (ie. $\epsilon > 0$),

$$\begin{aligned}
\Pr(\mathbf{Bin}(n-1, p_j) < (1 - \beta)\mu) &\leq e^{-\mu\beta^2/2} \\
&= \exp(-\Omega(p_j(n-1)(1 - e^{-\epsilon/2})^2)) \\
&= \exp(-\Omega(p_j n \epsilon^2))
\end{aligned}$$

So that:

$$\begin{aligned}
\Pr(\mathbf{Hist}(x_i, X_{-1}) \in \mathcal{B}) &\leq \Pr(\mathbf{Bin}(n-1, p_i) > p_i(n-1)(e^{\epsilon/2} - (p_i(n-1))^{-1})) \\
&\quad + \Pr(\mathbf{Bin}(n-1, p_j) < p_j(n-1)e^{-\epsilon/2}) \\
&\leq \exp(-\Omega(p_i n \epsilon^2)) + \exp(-\Omega(p_j n \epsilon^2)) \\
&\leq \exp(-\Omega(p_{\min} n \epsilon^2)) = \delta
\end{aligned}$$

for $2 \ln(1 + \frac{1}{p_{\min}(n-1)}) < \epsilon \leq 2 \ln(2)$. To get rid of the upper bound on ϵ , notice when $\epsilon = 2 \ln(2)$, $\delta = \exp(-\Omega(p_{\min} n (2 \ln(2))^2))$ suffices to satisfy the inequality

$$\Pr(\mathbf{Hist}(x_i, X_{-i}) \in \mathcal{S}) \leq e^\epsilon \Pr(\mathbf{Hist}(x_j, X_{-i}) \in \mathcal{S}) + \delta$$

Thus, when $\epsilon > 2 \ln(2)$, the same $\delta = \exp(-\Omega(np_{\min}[\min(2 \ln(2), \epsilon)]^2)) = \exp(-\Omega(p_{\min} n (2 \ln(2))^2))$ also suffices, as a larger ϵ only makes the right hand side of the inequality larger.

This proves Claim 2. □

The definition of distributional differential privacy, like differential privacy, is immune to post-processing. This means that if \mathbf{M} is $(\epsilon, \delta, \Delta)$ -DDP, and \mathbf{f} is a function on the output of \mathbf{M} , then $\mathbf{f} \circ \mathbf{M}$ (their composition) is also $(\epsilon, \delta, \Delta)$ -DDP. Note that post-processing immunity is not a property of exact privacy, since exact privacy describes tight bounds on ϵ, δ .

Lemma 3 (Immunity to Post-processing). *Suppose $\mathbf{M} : \mathcal{U}^* \rightarrow \mathcal{R}$ is $(\epsilon, \delta, \Delta)$ -DDP. Let $\mathbf{f} : \mathcal{R} \rightarrow \mathcal{R}'$ be a deterministic function. Then $\mathbf{f} \circ \mathbf{M} : \mathcal{U}^* \rightarrow \mathcal{R}'$ is also $(\epsilon, \delta, \Delta)$ -DDP.*

Proof. For any $\pi \in \Delta$, $x, x' \in \mathbf{Supp}(X_i)$ and $\mathcal{S} \subseteq \mathcal{R}'$, let $\mathcal{W} = \{w : \mathbf{f}(w) \in \mathcal{S}\}$. Then

$$\begin{aligned} & \Pr(\mathbf{f}(\mathbf{M}(\vec{X})) \in \mathcal{S} \mid X_i = x) \\ &= \Pr(\mathbf{M}(\vec{X}) \in \mathcal{W} \mid X_i = x) && \text{(By definition of } \mathcal{W} \text{)} \\ &\leq e^\epsilon \Pr(\mathbf{M}(\vec{X}) \in \mathcal{W} \mid X_i = x') + \delta && \text{(By } \mathbf{M} \text{ being } (\epsilon, \delta, \Delta)\text{-DDP)} \\ &= e^\epsilon \Pr(\mathbf{f}(\mathbf{M}(\vec{X})) \in \mathcal{S} \mid X_i = x) + \delta && \text{(By definition of } \mathcal{W} \text{)} \end{aligned}$$

□

By post-processing immunity, the parameters proven in Theorem 1 also apply to functions whose outputs are based on the histogram of the database, such as most voting rules.

C Exact Privacy of Voting Rules: Two Candidate-Case (Cont'd)

C.1 Proof of Lemma 1: Trails Technique

Lemma 1 (Trails) *Let T be a trail with direction (j, k) , and let π be a distribution where votes are independently distributed. For any $i, x_j, x_k \in \mathbf{Supp}(X_i)$,*

$$\begin{aligned} & \Pr(\mathbf{Hist}(\vec{X}) \in \mathsf{T} \mid X_i = x_j) - \Pr(\mathbf{Hist}(\vec{X}) \in \mathsf{T} \mid X_i = x_k) \\ &= \Pr(\mathbf{Hist}(\vec{X}) = \mathit{Exit}(\mathsf{T}) \mid X_i = x_j) - \Pr(\mathbf{Hist}(\vec{X}) = \mathit{Enter}(\mathsf{T}) \mid X_i = x_k) \end{aligned}$$

Proof for Lemma 1. 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 $j \in [l]$

$$\Pr(\mathbf{Hist}(\vec{X}) = t \mid X_i = x_j) = \Pr(\mathbf{Hist}(X_{-i}) = t - x_j)$$

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

$$\Pr(\mathbf{Hist}(\vec{X}) = t_z \mid X_i = x_j)$$

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

In other words,

$$\begin{aligned}
&\Pr(\mathbf{Hist}(\vec{X}) \in \mathbb{T} | X_i = x_j) - \Pr(\mathbf{Hist}(\vec{X}) \in \mathbb{T} | X_i = x_k) \\
&= \Pr(\mathbf{Hist}(\vec{X}) = t_q | X_i = x_j) - \Pr(\mathbf{Hist}(\vec{X}) = t_0) \\
&\quad + \sum_{0 \leq z < q} \Pr(\mathbf{Hist}(\vec{X}) = t_z | X_i = x_j) - \Pr(\mathbf{Hist}(\vec{X}) = t_{z+1} | X_i = x_k) \\
&= \Pr(\mathbf{Hist}(\vec{X}) = t_q | X_i = x_j) - \Pr(\mathbf{Hist}(\vec{X}) = t_0 | X_i = x_k) \\
&\quad \text{(Every term in the summation of differences cancels out.)} \\
&= \Pr(\mathbf{Hist}(\vec{X}) = \text{Exit}(\mathbb{T}) | X_i = x_j) - \Pr(\mathbf{Hist}(\vec{X}) = \text{Enter}(\mathbb{T}) | X_i = x_k)
\end{aligned}$$

□

C.2 Full proof for Theorem 2: Biased Majority

Theorem 2 (Exact DDP for Majority Rules) 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 $\exists \pi \in \Delta$ s.t. $\pi(a) = \alpha$; otherwise $\delta = \exp(-\Omega(n))$.

Proof. (Full proof for Theorem 2).

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 lower bound on δ). Thus, we only consider $\mathcal{S} = \{a\}$ (when the 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.):

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

We first discuss the case that $\mathcal{S} = \{a\}$ where its 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 . Thus,

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

and

$$\begin{aligned}
&\Pr(\mathbf{Hist}(\vec{X}) = \text{Exit}(\mathbb{T}_a) | X_i = a) \\
&= \Pr(\mathbf{Hist}(\vec{X}) = (\lceil \alpha n \rceil, \lfloor (1-\alpha)n \rfloor) | X_i = a) \\
&= \Pr(\mathbf{Hist}(\vec{X}) = (\lceil \alpha n \rceil - 1, \lfloor (1-\alpha)n \rfloor)) \\
&= p^{\lceil \alpha n \rceil - 1} (1-p)^{\lfloor (1-\alpha)n \rfloor} \frac{(n-1)!}{\lceil \alpha n - 1 \rceil! \cdot \lfloor (1-\alpha)n \rfloor!} \\
&= \Theta \left[\frac{1}{\sqrt{n}} \cdot \left(\frac{pn}{\lceil \alpha n - 1 \rceil} \right)^{\lceil \alpha n - 1 \rceil} \cdot \left(\frac{(1-p)n}{\lfloor (1-\alpha)n \rfloor} \right)^{\lfloor (1-\alpha)n \rfloor} \right] \quad \text{(Stirling's formula)} \\
&= \Theta \left(\sqrt{\frac{1}{n}} \left[\left(\frac{p}{\alpha} \right)^\alpha \left(\frac{1-p}{1-\alpha} \right)^{1-\alpha} \right]^n \right)
\end{aligned}$$

The case for $\mathcal{S} = \{b\}$ is similar. We note that $\left(\frac{p}{\alpha}\right)^\alpha \left(\frac{1-p}{1-\alpha}\right)^{1-\alpha} \leq 1$, and equality holds if and only if $p = \alpha$. Finally, we take the maximum of all δ 's over $\pi \in \Delta$. \square

D Exact Privacy of Voting Rules: General Case (Cont'd)

In all proofs of this section, we will use r instead of M to denote GSR voting rules.

D.1 Full proof for Theorem 3

Theorem 3 (Dichotomy of Exact DDP for GSR) Fix $m \geq 2$ and $\Delta \subseteq \Pi(\mathcal{L}(C))$ with $|\Delta| < \infty$. For any n , any GSR 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.

Theorem 3, (Exact) DDP for GSR. To present the result, we first introduce an equivalent definition of GSR that is similar to the ones used in [Xia and Conitzer, 2009; Mossel et al., 2013].

Definition 9 (The (H, g_H) definition of GSR). A GSR over m candidates is defined by a set of hyperplanes $H = \{\vec{h}_1, \dots, \vec{h}_R\} \subseteq \mathbb{R}^m$ and a function $g_H : \{+, 0, -\}^{|H|} \rightarrow C$. For any anonymous profile $\vec{p} \in \mathbb{R}^m$, we let $H(\vec{p}) = (\text{Sign}(\vec{h}_1 \cdot \vec{p}), \dots, \text{Sign}(\vec{h}_R \cdot \vec{p}))$, where $\text{Sign}(x)$ is the sign (+, - or 0) of a number x . We let the winner be $g_H(H(\vec{p}))$.

That is, to determine the winner, we first use each hyperplane in H to classify the profile \vec{p} , to decide whether \vec{p} is on the positive side (+), negative side (-), or is contained in the hyperplane (0). Then g_H is used to choose the winner from $H(\vec{p})$. We refer to this definition the (H, g_H) definition. Also see Example 4 for how (H, g_H) works. In the next claim, we show the equivalence of two definitions of GSR.

Claim 1. The (H, g_H) definition of GSR is equivalent to the (f, g) definition of GSR in Definition 5.

Proof for Claim 1. We first show that any (H, g_H) GSR can be represented by a (f, g) GSR in the following way: for each ranking V , we let $f(V) = (\vec{h}_1 \cdot \vec{e}_V, \vec{h}_2 \cdot \vec{e}_V, \dots, \vec{h}_R \cdot \vec{e}_V, 0)$. Then, the g function mimics g_H by only focusing on orderings between the k th component of $f(P)$ and the last component, which is always 0, for all $k \leq R$. More precisely, ordering between the k th component of $f(P)$ and 0 uniquely determines $\text{Sign}(\vec{h}_k \cdot \vec{p})$.

We now prove that any (f, g) GSR can be represented by an (H, g_H) GSR. For any pair of distinct component $k_1, k_2 \leq K$, we introduce a hyperplane $\vec{h}_{k_1, k_2} = ([f(V)]_{k_1} - [f(V)]_{k_2})_{V \in \mathcal{L}(C)}$. Therefore, for any profile \vec{p} , $\vec{h}_{k_1, k_2} \cdot \vec{p} = [f(\vec{p})]_{k_1} - [f(\vec{p})]_{k_2}$. The sign of $\vec{h}_{k_1, k_2} \cdot \vec{p}$ corresponds to the order between $[f(\vec{p})]_{k_1}$ and $[f(\vec{p})]_{k_2}$. Then, g_H mimics g . \square

We are now ready to present our theorem on GSRs. We will characterize eDDP under uniform distribution and give an exponential upper bound on DDP under some other distributions. For any pair of $\vec{\pi}$ and \vec{h} , we let $\text{Dist}(\vec{\pi}, \vec{h}) = \frac{\vec{\pi} \cdot \vec{h}}{\|\vec{h}\|_2}$ to denote the distance between hyperplane $\vec{h} \cdot \vec{p} = 0$ and vector $\vec{\pi}$.

We first show that w.l.o.g. we can assume that all hyperplanes in H passes $\vec{1}$.

Lemma 4. A GSR satisfies canceling-out, if and only if there exists another equivalent GSR $r = (H, g_H)$, where all hyperplanes passes $\vec{1}$.

Proof. The “if” direction is straightforward. To prove the “only if” part, it suffices to prove that g_H does not depend on outcomes of hyperplanes in H that does not pass $\vec{1}$. W.l.o.g. let $\vec{h}_1 \in H$ denote the hyperplane that does not pass $\vec{1}$, that is, $\vec{h}_1 \cdot \vec{1} \neq 0$. We will prove that for any $\vec{u}_{-1} \in \{-1, 0, 1\}^{L-1}$ and any $u_1, u'_1 \in \{-1, 0, 1\}$, such that there exist profiles P, Q with $H(P) = (u_1, \vec{u}_{-1})$ and $H(Q) = (u'_1, \vec{u}_{-1})$, we have $g_H(u_1, \vec{u}_{-1}) = g_H(u'_1, \vec{u}_{-1})$.

For the sake of contradiction, suppose this does not hold and let P, Q be the profiles such that $H(P)$ and $H(Q)$ differ on the first coordinate, and $r(P) \neq r(Q)$. Then, for sufficiently large n we have that $H(P + n\mathcal{L}(C)) = H(Q + n\mathcal{L}(C))$. This is because for any $\vec{h} \in H$ that passes $\vec{1}$, we have $\vec{h} \cdot (P + n\mathcal{L}(C)) = \vec{h} \cdot P = \vec{h} \cdot (Q + n\mathcal{L}(C))$. For any $\vec{h} \in H$ that does not pass $\vec{1}$, we have $\vec{h} \cdot (P + n\mathcal{L}(C)) = \vec{h} \cdot P + n\vec{h} \cdot \vec{1}$, and when n is sufficiently large, the

sign of $\vec{h} \cdot (P + n\mathcal{L}(C))$ is the same as the sign of $n\vec{h} \cdot 1$, which is the sign of $\vec{h} \cdot (Q + n\mathcal{L}(C))$. This means that $\text{Sign}(\vec{h} \cdot P) = \text{Sign}(\vec{h} \cdot (P + n\mathcal{L}(C))) = \text{Sign}(\vec{h} \cdot (Q + n\mathcal{L}(C))) = \text{Sign}(\vec{h} \cdot Q)$, which is a contradiction. \square

Let r be a GSR, P^* be the locally stable profile and a be the candidate, V, W be the rankings as in the statement of Definition 6. W.l.o.g. suppose V is the first type ranking and W is the second type ranking. In other words, V (respectively, W) is the first (respectively, second) coordinate in the m -profiles space. We will show that the exact DDP bound is achieved when S is the set of all profiles where the winner is a .

We recall that for any profile P , a pair of different votes V, W . and a length $q \in \mathbb{N}$, $\mathsf{T}_{P,V,W,q}$ is the trail starting at P , going along the $V - W$ direction, and contains q profiles. We let $\mathsf{T}_{P,V,W,\infty} = \max_q \mathsf{T}_{P,V,W,q}$ denote the longest $V - W$ trail starting at P . For a GSR r , we define $\text{End}(a) = \{\text{Exit}(\mathsf{T}_{P,V,W,\infty}) : \forall V, W \in \mathcal{U}, r(P) = a\}$. In other words, there are no W votes in $\text{End}(a)$.

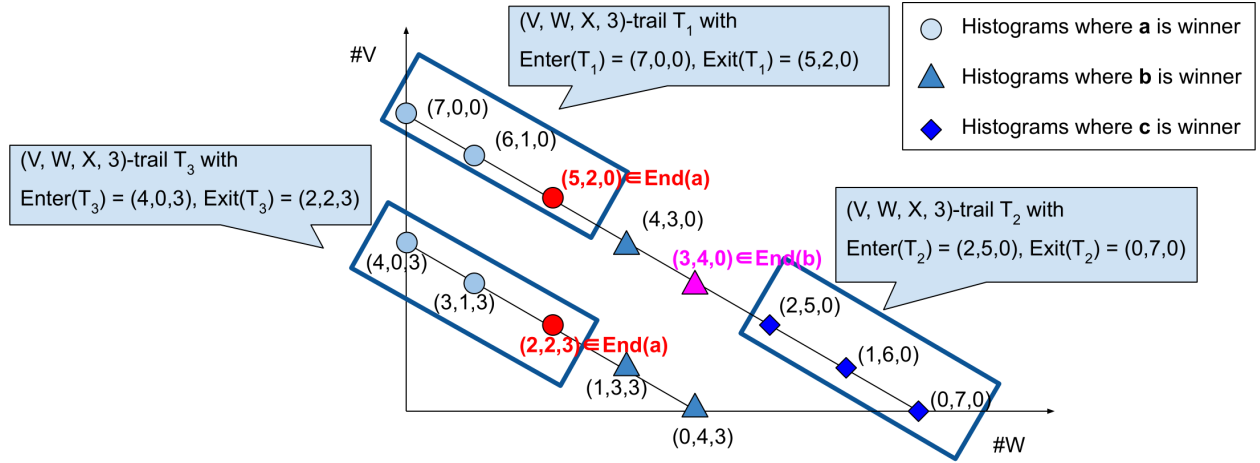


Figure 3: Example of $\text{End}(a)$ and $\text{End}(b)$, for 3-candidate case. The 3 kinds of votes other than V, W and X are not shown to simplify notations. Number of unshown votes are considered as constant.

Because r satisfies monotonicity, for any profile P such that $r(P) = a$, we must have that a is the winner under all profiles in the V - W trail starting at P . Therefore, S can be partitioned into multiple non-overlapping trails, each of which starts at a different profile, where a is the winner, and a is no longer the winner if we go one step into the W - V direction. Formally, we let $\text{End}(a)$ (shown in Figure 3) denote all n -profiles P such that (1) $r(P) = a$ and (2) $r(P + W - V) \neq a$. Then, we define a partition \mathcal{S}_a as follows.

$$\mathcal{S}_a = \{P : r(P) = a\} = \bigcup_{P \in \text{End}(a)} \mathsf{T}_{P,V,W,\infty}$$

It follows from Lemma 1 that

$$\Pr(P \in \mathcal{S}_a | X_1 = V) - \Pr(P \in \mathcal{S}_a | X_1 = W) = \sum_{P \in \text{End}(a) : P(V) > 0} \Pr(P - V).$$

We will define a subset of n -profile, \mathcal{R}_n and prove the lower bound on it. For a locally stable profile P^* (with constant γ in the statement of Definition 6), let $\vec{p}_0 = P^* - \vec{1} \cdot \frac{|P^*|}{m!}$. That is, \vec{p}_0 be obtained from P^* by subtracting a constant in each component, such that $\vec{p}_0 \cdot \vec{1} = 0$. For any n , we define \mathcal{R}_n to be the set of n -profiles that are in the $\gamma\sqrt{n}$ neighborhood of $\frac{n}{m!} \cdot \vec{1} + \vec{p}_0 \cdot \sqrt{n}$ w.r.t. L_∞ norm for last $m! - 2$ dimensions. That is,

$$\mathcal{R}_n = \left\{ P : P[V] = 0 \text{ and } \forall j \geq 3, \left| P[j] - \left(\frac{n}{m!} + \vec{p}_0[j] \cdot \sqrt{n} \right) \right| \leq \gamma\sqrt{n} \right\}$$

Throughout the proof in Theorem 3, we will use $\vec{\pi}$ to denote the database distribution D , and $\pi[j]$ denote the probability of j -th kind of ranking. Here $P[V]$ is the number of V votes in P and $P[j]$ is the number of j -th type of vote in P .

For any $P \in \mathcal{R}_n$, we let $\text{Piv}(P) = \text{End}(a) \cap \mathbb{T}_{P,V,W,\infty}$ denote the intersection of $\text{End}(a)$ and the V - W trail starting at P . That is, $\text{Piv}(P) = P + l(V - W)$ for some $l \in \mathbb{Z}$, $r(\text{Piv}(P)) = a$, and $r(\text{Piv}(P) - V + W) \neq a$.

We next prove that the number of V votes in $\text{Piv}(P)$ and the number of W votes in $\text{Piv}(P)$ are close—the difference is $O(\sqrt{n})$.

Claim 2. For any $P \in \mathcal{R}_n$, we have $|\text{Piv}(P)[V] - \text{Piv}(P)[W]| = O(\sqrt{n})$.

Proof. Let $Q^+ = \text{Piv}(P)$ and $Q^- = \text{Piv}(P) - V + W$. We note that $\text{Piv}(P)$ is at the boundary of S , which means that $r(Q^+) \neq r(Q^-)$. Therefore, because r is a GSR, the line segment between Q^+ and Q^- must contain the intersection of $\mathbb{T}_{P,V,W,\infty}$ and a hyperplane $\vec{h} \in H$. Therefore, it suffices to show that the difference in number of V votes and number of W votes at the intersection of $\mathbb{T}_{P,V,W,\infty}$ and any hyperplane \vec{h} is $O(\sqrt{n})$.

We recall that by Lemma 4, all hyperplanes for r pass $\vec{1}$. For any $\vec{h} \in H$, we recall that we assumed that V and W corresponds to the first and second coordinate, respectively. Because $\vec{h} \cdot (P + l(V - W)) = 0$, we have $(h_2 - h_1)l = \vec{h} \cdot P = \vec{h} \cdot (P - \vec{1} \cdot \frac{n}{m!}) = O(\sqrt{n})$. This means that $|l| = |\text{Piv}(P)[V] - \text{Piv}(P)[W]| = O(\sqrt{n})$. \square

Claim 3. For any $P \in \mathcal{R}_n$, there is a V - W trail passing P .

Proof. According to the cancelling out property of r , we can construct profile $P' = P - \frac{n - |P^*| \sqrt{n}}{m!}$, which is equivalent to P . For any profile $P \in \mathcal{R}_n$, we have $|P[j] - (\frac{n}{m!} + \vec{p}_0[j] \cdot \sqrt{n})| \leq \gamma \sqrt{n}$, which is equivalent with $|P'[j] - P^*[j] \cdot \sqrt{n}| \leq \gamma \sqrt{n}$, which means $\frac{P'}{\sqrt{n}}$ is in the γ neighborhood of profile P^* in terms of the 3-rd to $m!$ -th dimensions. According to the (H, g_H) definition of GSR, we know $r(P^*) = r(P')$ and the claim follows by local stability of P^* . \square

We will show that the probability of a subset of $\text{End}(a)$ —the pivotal profiles on trails starting at profiles in \mathcal{R}_n —is $\Theta(1/\sqrt{n})$ for the condition that π is uniform over \mathcal{U} . Let $\mathcal{R}_n^- \subseteq \mathbb{R}^{m!-2}$ and for any $\vec{p}_- \in \mathcal{R}_n^-$, we define $\text{Piv}(\vec{p}_-) = \text{Piv}(P)$, where $P \in \mathcal{R}_n$ and $P[3, \dots, m!] = \vec{p}_-$.

$$\begin{aligned} & \sum_{P \in \text{End}(a)} \Pr(P - V) \geq \sum_{P \in \mathcal{R}_n} \Pr(\text{Piv}(P) - V) \\ &= \sum_{\vec{p}_- \in \mathcal{R}_n^-, |P|=n-1} \left(\Pr(P[3, \dots, m!] = \vec{p}_-) \right. \\ & \quad \left. \Pr(P[1] = \text{Piv}(\vec{p}_-)[1] - 1, \Pr(P[2] = \text{Piv}(\vec{p}_-)[2] | P[3, \dots, m!] = \vec{p}_-) \right) \\ &= \sum_{\vec{p}_- \in \mathcal{R}_n^-, |P|=n-1} A(\vec{p}_-) B(\vec{p}_-) \end{aligned}$$

where $A(\vec{p}_-) = \Pr(P[3, \dots, m!] = \vec{p}_-)$ and

$$B(\vec{p}_-) = \Pr(P[1] = \text{Piv}(\vec{p}_-)[1] - 1, \Pr(P[2] = \text{Piv}(\vec{p}_-)[2] | P[3, \dots, m!] = \vec{p}_-)$$

It follows that $B(\vec{p}_-)$ is equivalent to probability of flipping a coin ($\frac{\pi[W]}{\pi[V] + \pi[W]}$ probability for head) for $\text{Piv}(\vec{p}_-)[1] + \text{Piv}(\vec{p}_-)[2] - 1$ times, with $\text{Piv}(\vec{p}_-)[1] - 1$ heads and $\text{Piv}(\vec{p}_-)[2]$ tails. The next lemma gives a lower bound to $\sum_{\vec{p}_- \in \mathcal{R}_n^-, |P|=n-1} A(\vec{p}_-) B(\vec{p}_-)$ when π is a uniform distribution.

Lemma 5. $\sum_{\vec{p}_- \in \mathcal{R}_n^-, |P|=n-1} A(\vec{p}_-) B(\vec{p}_-) = \Omega\left(\frac{1}{\sqrt{n}}\right)$ if π is uniform over \mathcal{U} .

Proof. We first bound the total number of V and W votes in $P \in \mathcal{R}_n$ in the next claim.

Claim 4. $\text{Piv}(\vec{p}_-)[1] + \text{Piv}(\vec{p}_-)[2] - 1 = \Theta(n)$ for all $\vec{p}_- \in \mathcal{R}_n^-$.

Proof.

$$\left| \text{Piv}(\vec{p}_-)[1] + \text{Piv}(\vec{p}_-)[2] - \frac{2n}{m!} \right| = \sum_{j=3}^{m!} \left| P[j] - \frac{n}{m!} \right| \leq \sum_{j=3}^{m!} (\gamma \sqrt{n} + |\vec{p}_0[j]| \sqrt{n}) \leq (\gamma + 1) m! \sqrt{n}$$

□

According to Claim 2 & 4, we know that $B(\vec{p}_-)$ is equivalent to probability of flipping a fair coin for $\frac{2n}{m!} + c_1\sqrt{n}$ times and get $\frac{n}{m!} + c_2\sqrt{n}$, where c_1 and c_2 are bounded constants. In the next claim, we give a tight bound to $B(\vec{p}_-)$ for uniform distributed entries.

Claim 5. $B(\vec{p}_-) = \Theta\left(\sqrt{\frac{1}{n}}\right)$ for any $\vec{p}_- \in \mathcal{R}_n^-$

Proof. Letting $n' = \frac{2n}{m!} + c_1\sqrt{n}$, $c' = c_2 - \frac{c_1}{2}$ and assuming n' is a even number, for the lower bound, we have,

$$\begin{aligned}
B(\vec{p}_-) &= \left(\frac{1}{2}\right)^{\frac{2n}{m!} + c_1\sqrt{n}} \binom{\frac{2n}{m!} + c_1\sqrt{n}}{\frac{n}{m!} + c_2\sqrt{n}} = \left(\frac{1}{2}\right)^{n'} \binom{n'}{n'/2 + c'\sqrt{n}} \\
&= \left(\frac{1}{2}\right)^{n'} \cdot \binom{n'}{n'/2} \cdot \frac{\frac{n'}{2} \times \cdots \times (\frac{n'}{2} - c'\sqrt{n'} + 1)}{(\frac{n'}{2} + c'\sqrt{n'} - 1) \times \cdots \times \frac{n'}{2}} \\
&> \frac{1}{2^{n'}} \binom{n'}{n'/2} \cdot \left(\frac{n'/2 - c'\sqrt{n'}}{n'/2}\right)^{c'\sqrt{n'}} \\
&= \Omega\left(\frac{1}{\sqrt{n}}\right) \quad (\text{applying Stirling's Formula})
\end{aligned} \tag{8}$$

Upper bound can be obtained using similar technique as lower bound. □

The next claim gives a lower bound on $\sum_{\vec{p}_- \in \mathcal{R}_n^-} A(\vec{p}_-)$. The proof uses the main technique of Lindeberg-Levy Central Limit Theorem [Greene, 2003].

Claim 6. $\sum_{\vec{p}_- \in \mathcal{R}_n^-} A(\vec{p}_-) = \Omega(1)$.

Proof of Claim 6. We first define a set of $m! - 2$ dimensions random variables that $Y_i = (Y_i[1], \dots, Y_i[m! - 2])$, where $Y_i[j] = 1$ if ranking j happens to i -th row and $Y_i[j] = 0$ otherwise. According to the definition of profile, we have $P[j + 2] = \sum_{i=1}^n Y_i[j]$ and $\mathbb{E}(P[j]) = \frac{n}{m!}$ for uniform case. We further define a $m! - 2$ dimensional random vector \vec{u} such that $\vec{u}[j] = (P[j + 2] - \frac{n}{m!}) / \sqrt{n}$, which is the scaled average of Y_1, \dots, Y_n . According to Lindeberg-Levy Central Limit Theorem [Greene, 2003], we know that the distribution of \vec{u} converges in probability to multivariate normal distribution $\mathcal{N}(0, \Sigma)$, where

$$\Sigma = \begin{bmatrix} \frac{m!-1}{(m!)^2} & -\frac{1}{(m!)^2} & \cdots & -\frac{1}{(m!)^2} \\ -\frac{1}{(m!)^2} & \frac{m!-1}{(m!)^2} & \cdots & -\frac{1}{(m!)^2} \\ \vdots & \vdots & \ddots & \vdots \\ -\frac{1}{(m!)^2} & -\frac{1}{(m!)^2} & \cdots & \frac{m!-1}{(m!)^2} \end{bmatrix}.$$

Since each diagonal element in Σ is strictly larger than the sum of the absolute value of all other elements in the same row, we know that Σ is non-singular according to Levy-Desplanques Theorem [Horn and Johnson, 1990]. According to Varah *et al.* [Varah, 1975], we obtain a bound on Σ^{-1} 's L_∞ norm as,

$$\|\Sigma^{-1}\|_\infty \leq \frac{1}{\min_i \left(|\Sigma_{ii}| - \sum_{j \neq i} |\Sigma_{ij}| \right)} \leq \frac{(m!)^2}{2}.$$

For any $m! - 2$ dimensional random vector \vec{u} constructed from a profile P using the procedure that $\vec{u}[j] = (P[j + 2] - \frac{n}{m!}) / \sqrt{n}$, we have,

$$P \in \mathcal{R}_n^- \quad \text{if and only if} \quad \vec{u} \in \mathbb{U} = \{\vec{u} : |\vec{u}[j] - \vec{p}_0[j]| \leq \gamma, \forall j \in [m! - 2]\}.$$

Thus, for all $\vec{u} \in \mathbb{U}$ we know about its Probability Density Function (PDF) that,

$$\begin{aligned}
\text{PDF}(\vec{u}) &= \frac{1}{\sqrt{(2\pi)^{m^1-2}|\Sigma|}} \exp\left(-\frac{1}{2}\vec{u}^T \Sigma^{-1} \vec{u}\right) \\
&= \frac{1}{\sqrt{(2\pi)^{m^1-2}|\Sigma|}} \exp\left(-\frac{1}{2}|\vec{u}^T \Sigma^{-1} \vec{u}|\right) \\
&\geq \frac{1}{\sqrt{(2\pi)^{m^1-2}|\Sigma|}} \exp\left(-\frac{1}{2}\|\vec{u}^T \Sigma^{-1}\|_\infty \cdot \|\vec{u}\|_1\right) \quad (\text{Holder's Inequality}) \\
&\geq \frac{1}{\sqrt{(2\pi)^{m^1-2}|\Sigma|}} \exp\left(-\frac{1}{2}\|\vec{u}^T\|_\infty \cdot \|\Sigma^{-1}\|_\infty \cdot \|\vec{u}\|_1\right) \\
&\geq \frac{1}{\sqrt{(2\pi)^{m^1-2}|\Sigma|}} \left[\exp\left(\frac{(m!)^2}{4}\right)\right]^{-\|\vec{u}\|_\infty^2} \\
&= \Omega(1).
\end{aligned}$$

Thus, letting $\text{Vol}(\cdot)$ be the volume function,

$$\sum_{\vec{p}_- \in \mathcal{R}_n^-} A(\vec{p}_-) \geq \text{Vol}(\mathbb{U}) \cdot \min_{\vec{u} \in \mathbb{U}} \text{PDF}(\vec{u}) \geq \gamma^{m^1-2} \cdot \Omega(1) = \Omega(1).$$

□

Lemma 5 follows by combining Claim 6 and Claim 5. □

Recalling Lemma 1, for the case that π is uniform over all ranking, we have,

$$\begin{aligned}
\delta &= \max_{x, x', \mathcal{S}} \Pr(\mathbf{M}(X) \in \mathcal{S} | X_1 = x) - \Pr(\mathbf{M}(X) \in \mathcal{S} | X_1 = x') \\
&\leq \Pr(\mathbf{M}(X) \in \mathcal{S}_a | X_1 = W) - \Pr(\mathbf{M}(X) \in \mathcal{S}_a | X_1 = V) \\
&= \sum_{P \in \text{End}(a)} \Pr(P - V) = \Omega\left(\frac{1}{\sqrt{n}}\right).
\end{aligned}$$

Then, we derive an upper bound of δ using the similar technique of lower bound (π can be non-uniform for this bound). We first define \mathcal{R}'_n , a subset of n -profile space, where event $P \in \mathcal{R}'_n$ will be proved to happen with high probability.

$$\mathcal{R}'_n = \left\{P : P[V] = 0 \text{ and } \forall j \geq 3, |P[j] - (n \cdot \pi[j])| \leq n^{3/4}\right\}.$$

Then, we recall Lemma 1, for the case that π such that $\min_i \pi[i] > 0$, we have,

$$\begin{aligned}
\delta &= \max_{V, W, \mathcal{S}} \Pr(P \in \mathcal{S} | X_1 = V) - \Pr(P \in \mathcal{S} | X_1 = W) \\
&\leq \max_{V, W} \sum_{i=1}^m \Pr(P \in \mathcal{S}_i | X_1 = V) - \Pr(P \in \mathcal{S}_i | X_1 = W) = \sum_{i=1}^m \sum_{P \in \text{End}(x_i)} \Pr(P - V).
\end{aligned}$$

where $\mathcal{S}_i = \{X : r(X) = x_i\} = \bigcup_{P \in \text{End}(x_i)} \mathbb{T}_{P, V, W, \infty}$. The next claim gives an upper bound on the number of pivotal profiles sharing one End.

Claim 7. For any profile P in \mathcal{R}'_n , there are at most $|H|$ pivotal profiles following $V - W$ direction.

Proof. We know from the (H, g_H) definition of GSR that r 's output only changes while passing at least one hyperplane. Considering a trail \mathbb{T}_{P_0} enter at $(P_0[1] + P_0[2], 0, P_0[3], \dots, P_0[m!])$ and exit at $(0, P_0[1] + P_0[2], P_0[3], \dots, P_0[m!])$ (P_0 is an arbitrary n -profile). Thus, there are at most $|H|$ pivotal profiles sharing the same end point because \mathbb{T}_{P_0} passes hyperplanes at most $|H|$ times. □

Using the partition of \mathcal{R}'_n and arbitrarily selected candidate a , we have,

$$\begin{aligned}
\sum_{P \in \text{End}(x_i)} \Pr(P - V) &\leq |H| \left(\sum_{P \in \mathcal{R}'_n} \Pr(\text{Piv}(P) - V) + \sum_{P \in \text{End}(x_i) \setminus \mathcal{R}'_n} \Pr(\text{Piv}(P) - V) \right) \\
&\leq |H| \left(\sum_{\vec{p}_- \in \mathcal{R}'_n, |P|=n-1} A(\vec{p}_-)B(\vec{p}_-) + \sum_{\vec{p}_- \notin \mathcal{R}'_n, |P|=n-1} A(\vec{p}_-)B(\vec{p}_-) \right) \\
&\leq |H| \left(\max_{\vec{p}_- \in \mathcal{R}'_n} B(\vec{p}_-) \cdot \sum_{\vec{p}_- \in \mathcal{R}'_n} A(\vec{p}_-) + \max_{\vec{p}_- \notin \mathcal{R}'_n} B(\vec{p}_-) \cdot \sum_{\vec{p}_- \notin \mathcal{R}'_n} A(\vec{p}_-) \right) \\
&= O\left(\frac{1}{\sqrt{n}}\right) \cdot O(1) + O(1) \cdot O\left(\frac{1}{\sqrt{n}}\right) \quad (\text{by applying Claim 9}) \\
&= O\left(\frac{1}{\sqrt{n}}\right)
\end{aligned}$$

The next claim gives an upper bound to $\sum_{\vec{p}_- \notin \mathcal{R}'_n} A(\vec{p}_-)$.

Claim 8. $\sum_{\vec{p}_- \notin \mathcal{R}'_n} A(\vec{p}_-) = O\left(\frac{1}{\sqrt{n}}\right)$.

Proof. Let $Y_j^{(i)}$ = "the i -th agent gives vote of type j ". One can see that $P[j] = \sum_{i=1}^n Y_j^{(i)}$, $\mathbb{E}(P[j]) = n\pi[j]$ and $\text{Var}(P[j]) = n\pi[j](1 - \pi[j])$. Thus,

$$\begin{aligned}
\sum_{\vec{p}_- \notin \mathcal{R}'_n} A(\vec{p}_-) &= \Pr \left[\bigcup_{j=3}^{m!} \left\{ |P[j] - n \cdot \pi[j]| \leq n^{3/4} \right\} \right] \\
&\leq \sum_{j=3}^{m!} \Pr \left[\left\{ |P[j] - \mathbb{E}(P[j])| \leq n^{3/4} \right\} \right] \\
&\leq \sum_{j=3}^{m!} \frac{n\pi[j](1 - \pi[j])}{n^{3/2}} \quad (\text{by Chebyshev's Inequality}) \\
&= O\left(\frac{1}{\sqrt{n}}\right)
\end{aligned}$$

□

Then, all we need is an upper bound on $B(\vec{p}_-)$, and we first prove that the length of $V - W$ sequence is $\Theta(n)$ for all $P \in \mathcal{R}'_n$.

Claim 9. $\text{Piv}(\vec{p}_-)[1] + \text{Piv}(\vec{p}_-)[2] - 1 = \Theta(n)$ for all $P \in \mathcal{R}'_n$.

Proof.

$$|\text{Piv}(\vec{p}_-)[1] + \text{Piv}(\vec{p}_-)[2] - n(\pi[W] + \pi[V])| = \sum_{j=3}^{m!} |P[j] - n \cdot \pi[j]| \leq \sum_{j=3}^{m!} n^{3/4} \leq m! \cdot n^{3/4}$$

□

Then, using the same technique of Claim 5, we know that,

$$B(\vec{p}_-) = \Theta\left(\sqrt{\frac{1}{n}}\right) \quad \text{for all } p_- \in \mathcal{R}'_n$$

Thus, combining all results above, we have,

$$\delta \leq \sum_{i=1}^m \sum_{P \in \text{End}(x_i)} \Pr(P - V) = \sum_{i=1}^m \sum_{P \in \text{End}(x_i)} \Pr(P - V) = O\left(\frac{1}{\sqrt{n}}\right)$$

Next, we will give a exponential (tighter) upper bound on δ when π does not belong to any hyperplanes. We first give a generalized definition of pivotal profile.

Definition 10 (Generalized Pivotal Profile). *Profile P is a (generalized) pivotal profile if there exist pair of votes V and W such that $r(P) \neq r(P - V + W)$.*

Then, we define a distance function $\text{Dist}^*(P, h)$ to be a generalized distance between profile P and hyperplane h . We define

$$\text{Dist}^*(P, \vec{h}) = \inf_{P' \in \vec{h}} \|P - P'\|_2,$$

where $\vec{h} = \{P \in h : \exists \text{ unit vector } \vec{e} \text{ s.t. } r(P' - \vec{e}) \neq r(P' + \vec{e})\}$. In the next lemma we will show generalized pivotal profiles only lays close to hyperplanes. We fist gives definition of distance function $\text{Dist}(\cdot, \cdot)$:

1. for hyperplane h and a point (n -profile) P , $\text{Dist}(h, P) = \frac{\vec{h} \cdot P}{\|\vec{h}\|_2}$, which is the Euclidian distance between P and hyperplane $\vec{h} \cdot \vec{p} = 0$.
2. for 2 points (n -profile) P_1 and P_2 , $\text{Dist}(h, P)$, returns the Euclidian distance between P_1 and P_2 .

Claim 10. *For any GSR $r = (H, g_H)$ and one of its generalized pivotal profile P , there must exist one hyperplane $\vec{h} \in H$ such that $\text{Dist}(h, P) \leq \sqrt{2}$.*

Proof. Recalling the definition of generalized pivotal profiles, we know the GSR winner will change at the 1 neighborhood of P . Thus, there must exist a hyperplane $\vec{h} \in H$ and pair of votes V, W such that $\text{Sign}[\vec{h} \cdot P] \neq \text{Sign}[\vec{h} \cdot (P + V - W)]$ and $\text{Dist}(h, P) \leq \text{Dist}(P, P + V - W) = \sqrt{2}$. \square

Lemma 6. *Let D be the distribution on profiles (databases of votes), where each entry is iid according to distribution π over linear orders on m candidates. GSR $r(H, h_H)$ is $(0, \delta, \Delta = \{(D, \emptyset)\})$ -DDP when only the winner is announced, where*

$$\delta = O\left[\exp\left(-\frac{[\min_{h \in H} \text{Dist}^*(\vec{\pi}, h)]^2}{3(m!)^2 (\max_{i \in [m]} \pi[i])} \cdot n\right)\right] = O\left[e^{-\Omega(n)}\right].$$

Proof. We first define the set of all generalized pivotal profiles \mathbb{P}_{piv} . For any $P \in \mathbb{P}_{\text{piv}}$, we know that there exist hyperplane $h \in H$ such that $\text{Dist}^*(h, P) \leq \sqrt{2}$. According to triangular inequality, we have $\text{Dist}^*(n\vec{\pi}, P) \geq \text{Dist}^*(n\vec{\pi}, h) - \text{Dist}(h, P) \geq n\text{Dist}^*(\vec{\pi}, h) - \sqrt{2}$. The second \geq sign comes from the fact that all hyperplanes passes $\vec{0}$. Thus, there must exist one dimension j that $|P[j] - n\pi[j]| \geq \frac{n\text{Dist}^*(\vec{\pi}, h) - \sqrt{2}}{m!}$. Then, we bound δ as,

$$\begin{aligned} \delta &= \max_{V, W, \mathcal{S}} [\Pr(P \in \mathcal{S}_i | X_1 = V) - \Pr(P \in \mathcal{S}_i | X_1 = W)] \\ &\leq \sum_{P \in \mathbb{P}_{\text{piv}}} \left[\max_V \Pr(P \in \mathbb{P}_{\text{piv}} | X_1 = V) \right] \\ &\leq \max_{V, h, j} \Pr\left(|P[j] - n\pi[j]| \geq \frac{n\text{Dist}^*(\vec{\pi}, h) - \sqrt{2}}{m!} \middle| X_1 = V\right) \\ &\leq \max_{h, j} \Pr\left(|P[j] - n\pi[j]| \geq \frac{n\text{Dist}^*(\vec{\pi}, h) - \sqrt{2}}{m!} - 1\right) \\ &= O\left[\exp\left(-\frac{[\min_{h \in H} \text{Dist}^*(\vec{\pi}, h)]^2}{3(m!)^2 (\max_{i \in [m]} \pi[i])} \cdot n\right)\right] \text{ by applying Chernoff bound.} \end{aligned}$$

\square

Theorem 3 follows by combining all three bounds derived above. □

D.2 Proof for Corollary 1

Proposition 2. *All positional scoring rules and all Condorcet consistent and monotonic rules satisfy all properties required by Theorem 3.*

Proof of Proposition 2. 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 $c_1 \rightarrow c_2 \rightarrow \dots \rightarrow c_{m-2} \rightarrow c_1$. 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.

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

Corollary 1 follows by the definition of voting rules and the definition of positional scoring rules.

D.3 Exact DDP for Histogram

As a complementary result to the DDP result for histograms, we present the histogram's eDDP with $\epsilon = 0$.

Theorem 4 (Exact DDP of Histogram). *Fix $l \geq 2$, $\mathcal{U} = \{x_1, \dots, x_l\}$, and $\Delta \subseteq \Pi(\mathcal{U})$. Let $p_{\min} = \min_{\pi \in \Delta} (\pi(x_i) + \pi(x_j))$. For all $n \in \mathbb{N}$, **Hist** of n voters is $(0, \delta(n) = \Theta\left(\sqrt{\frac{1}{np_{\min}}}\right), \Delta)$ -eDDP.*

Sketch. First we present the case for $l = 2$.

Lemma 7 (Exact DDP for Histogram, when $l = 2$). *Fix $\mathcal{U} = \{x_1, x_2\}$ and $\Delta \subseteq \Pi(\mathcal{U})$. The histogram for n voters is $(0, \Theta(1/\sqrt{n}), \Delta)$ -eDDP.*

Lemma 7. Consider some $\pi \in \Delta$, and let $p = \pi(a)$. Without loss of generality let $x = x_1$ and $x' = x_2$ (otherwise, rename them). Then, the maximizing set \mathcal{S} in Equation (3) is exactly the set of histograms such that

$$\Pr_{\vec{X} \in \pi} (\mathbf{Hist}(\vec{X}) \in \mathcal{S} | X_i = x_1) > \Pr(\mathbf{Hist}(\vec{X}) \in \mathcal{S} | X_i = x_2)$$

Since votes are i.i.d., these follow the binomial distribution (with n trials). Below we find that \mathcal{S} is the set of histograms $(k, n - k)$ where $k > pn$.

$$\begin{aligned} \Pr(\mathbf{Hist}(\vec{X}) = (k, n - k) | X_i = x_1) &> \Pr(\mathbf{Hist}(\vec{X}) = (k, n - k) | X_i = x_2) \\ \implies p^{k-1}(1-p)^{n-k} \frac{(n-1)!}{(n-k)!(k-1)!} &> p^k(1-p)^{n-k-1} \frac{(n-1)!}{(n-k-1)!k!} \\ \implies k &> pn \end{aligned}$$

Thus, $\mathcal{S} = \{t = (k, n - k) : k > pn\}$. This set forms a trail T which starts from $\text{Enter}(T) = (n, 0)$ and exits at $\text{Exit}(T) = (pn + 1, n - (pn + 1))$. Thus,

$$\begin{aligned} \delta &= \Pr(\mathbf{Hist}(\vec{X}) \in \mathcal{S} | X_i = x_1) - \Pr(\mathbf{Hist}(\vec{X}) \in \mathcal{S} | X_i = x_2) && \text{(Equation (3))} \\ &= \Pr(\mathbf{Hist}(\vec{X}) = \text{Exit}(T) | X_i = x_1) - \Pr(\mathbf{Hist}(\vec{X}) = \text{Enter}(T) | X_i = x_2) && \text{(Lemma 1)} \\ &= \Pr(\mathbf{Hist}(\vec{X}) = (pn + 1, n - (pn + 1)) | X_i = x_1) - \Pr(\mathbf{Hist}(\vec{X}) = (n, 0) | X_i = x_2) \\ &= p^{pn}(1-p)^{n-pn-1} \frac{(n-1)!}{(pn)!(n-pn-1)!} \\ &\quad \text{(When one row is fixed to } x_2, \text{ the probability of histogram being } (n, 0) \text{ is zero.)} \\ &= \Theta(1/\sqrt{n}) && \text{(By applying Stirling's formula)} \end{aligned}$$

□

We can generalize the result to $l > 2$, by using the trail technique. Again we assume WLOG that $x = x_1$ and $x' = x_2$. Let $t = (t_1, \dots, t_l)$ be the histogram, where t_i counts the number of occurrences of x_i . We observe that, when votes are i.i.d, t_3, \dots, t_l are independent of t_1, t_2 when conditioned on the sum $s = t_1 + t_2$. This means that we can compute δ for general l , as a sum

$$\delta = \sum_{0 < s \leq n} \delta_s \Pr(\mathbf{Bin}(n, \pi(x_1) + \pi(x_2)) = s)$$

Where δ_s is the δ -value for $l = 2$, when there are s votes. Using Chernoff bound we see that $\mathbf{Bin}(n, \pi(x_1) + \pi(x_2))$ is concentrated at its mean $n(\pi(x_1) + \pi(x_2))$. Plugging in the result for $l = 2$, we get $\delta = \Theta\left(\frac{1}{\sqrt{n(\pi(x_1) + \pi(x_2))}}\right)$. \square

Full proof

Below we present the full proof of Theorem 4, using Lemma 7 which showed the case for $l = 2$.

Proof of Theorem 4, Exact DDP of Histogram. Consider any $\pi \in \Delta$, and let $p_i = \pi(x_i)$. Like in the $l = 2$ case, without loss of generality, we can let $x = x_1$ and $x' = x_2$ (otherwise, rename them). Then, the maximizing set \mathcal{S} (similar to when $l = 2$) is exactly the set of histograms such that

$$\Pr_{\vec{X} \sim \pi}(\mathbf{Hist}(\vec{X}) \in \mathcal{S} | X_i = x_1) > \Pr_{\vec{X} \sim \pi}(\mathbf{Hist}(\vec{X}) \in \mathcal{S} | X_i = x_2)$$

(We will implicitly assume $\vec{X} \sim \pi$ from now on) Since we have i.i.d. votes, the histogram follows the multinomial distribution (with n trials). For any $0 < s \leq n$, (t_3, \dots, t_l) where $t_3 + \dots + t_l = n - s$, and $k \leq s$:

$$\begin{aligned} \Pr(\mathbf{Hist}(\vec{X}) = (k, s - k, t_3, \dots, t_l) | X_i = x_1) &> \Pr(\mathbf{Hist}(\vec{X}) = (k, s - k, t_3, \dots, t_l) | X_i = x_2) \\ p_1^{k-1} p_2^{n-k} p_3^{t_3} \dots p_l^{t_l} \frac{(n-1)!}{(k-1)!(s-k)!t_3! \dots t_l!} &> p_1^k p_2^{n-k-1} p_3^{t_3} \dots p_l^{t_l} \frac{(n-1)!}{(s-k-1)!k!t_3! \dots t_l!} \\ \frac{p_2}{s-k} &> \frac{p_1}{k} \\ k &> \left(\frac{p_1}{p_1 + p_2}\right) s \end{aligned}$$

Thus, the set $\mathcal{S} = \left\{ t = (k, s - k, t_3, \dots, t_l) : k > \left(\frac{p_1}{p_1 + p_2}\right) s \right\}$.

Let $p = \frac{p_1}{p_1 + p_2}$. For each $0 < s \leq n$ and (t_3, \dots, t_l) which sum to $n - s$ (i.e. $t_3 + \dots + t_l = n - s$), let $\mathbb{T}_{s, (t_3, \dots, t_l)}$ be the trail starting from $\text{Enter}(\mathbb{T}_{s, (t_3, \dots, t_l)}) = (s, 0, t_3, \dots, t_l)$ and exiting at $\text{Exit}(\mathbb{T}_{s, (t_3, \dots, t_l)}) = (ps + 1, s - (ps + 1), t_3, \dots, t_l)$. The set \mathcal{S} then can be partitioned into such trails. Thus,

$$\begin{aligned} \delta &= \Pr(\mathbf{Hist}(\vec{X}) \in \mathcal{S} | X_i = x_1) - \Pr(\mathbf{Hist}(\vec{X}) \in \mathcal{S} | X_i = x_2) \\ &= \sum_{\mathbb{T}_{s, (t_3, \dots, t_l)}} \Pr(\mathbf{Hist}(\vec{X}) \in \mathbb{T}_{s, (t_3, \dots, t_l)} | X_i = x_1) - \Pr(\mathbf{Hist}(\vec{X}) \in \mathbb{T}_{s, (t_3, \dots, t_l)} | X_i = x_2) \\ &= \sum_{\mathbb{T}_{s, (t_3, \dots, t_l)}} \Pr(\mathbf{Hist}(\vec{X}) = \text{Exit}(\mathbb{T}_{s, (t_3, \dots, t_l)}) | X_i = x_1) \\ &\quad - \Pr(\mathbf{Hist}(\vec{X}) = \text{Enter}(\mathbb{T}_{s, (t_3, \dots, t_l)}) | X_i = x_2) \tag{By Lemma 1} \\ &= \sum_{0 < s \leq n} \sum_{\substack{(t_3, \dots, t_l) \\ t_3 + \dots + t_l = n - s}} \Pr(\mathbf{Hist}(\vec{X}) = (ps + 1, s - (ps + 1), t_3, \dots, t_l) | X_i = x_1) \\ &\quad - \Pr(\mathbf{Hist}(\vec{X}) = (s, 0, t_3, \dots, t_l) | X_i = x_2) \end{aligned}$$

Now let us consider these two probabilities. Consider the distribution X_{-i} , which is \vec{X} but without the i th row. Let the random variables of the individual components of $\mathbf{Hist}(X_{-i})$ be (a_1, \dots, a_l) . Since votes are i.i.d., for any

(t_1, \dots, t_l) ,

$$\begin{aligned}
\Pr(\mathbf{Hist}(\vec{X}) = (t_1, \dots, t_l) | X_i = x_1) &= \Pr(\mathbf{Hist}(X_{-i}) = (t_1 - 1, t_2, t_3, \dots, t_l)) \\
&= \Pr((a_1, \dots, a_l) = (t_1 - 1, t_2, t_3, \dots, t_l)) \quad (\text{Recall these } a\text{'s are components of } \mathbf{Hist}(X_{-i})) \\
&= \Pr((a_1, \dots, a_l) = (t_1 - 1, t_2, t_3, \dots, t_l) | a_1 + a_2 = s) \times \Pr(a_1 + a_2 = s) \\
&= \Pr((a_1, a_2) = (t_1 - 1, t_2) | a_1 + a_2 = s) \\
&\quad \times \Pr((a_3, \dots, a_l) = (t_3, \dots, t_l) | a_1 + a_2 = s) \times \Pr(a_1 + a_2 = s) \\
&\quad (\text{By Lemma 8, } (a_1, a_2) \text{ and } (a_3, \dots, a_l) \text{ are independent conditioned on } a_1 + a_2 = s)
\end{aligned}$$

Similar to the $l = 2$ case, $\Pr(\mathbf{Hist}(\vec{X}) = (s, 0, t_3, \dots, t_l) | X_i = x_2) = 0$. This is because when one vote is fixed to x_2 , it is impossible to have zero in the second component in the histogram (which is the number of occurrences of x_2). Thus,

$$\begin{aligned}
\delta &= \sum_{0 < s \leq n} \sum_{\substack{(t_3, \dots, t_l) \\ t_3 + \dots + t_l = n - s}} \Pr((a_1, a_2) = (ps, s - (ps + 1)) | a_1 + a_2 = s) \\
&\quad \times \Pr((a_3, \dots, a_l) = (t_3, \dots, t_l) | a_1 + a_2 = s) \times \Pr(a_1 + a_2 = s) \\
&= \sum_{0 < s \leq n} \Pr((a_1, a_2) = (ps, s - (ps + 1)) | a_1 + a_2 = s) \times \Pr(a_1 + a_2 = s) \\
&\quad \times \sum_{\substack{(t_3, \dots, t_l) \\ t_3 + \dots + t_l = n - s}} \Pr((a_3, \dots, a_l) = (t_3, \dots, t_l) | a_1 + a_2 = s) \\
&\quad (\text{Factor out the common terms } \Pr((a_1, a_2) = (ps, s - (ps + 1)) | a_1 + a_2 = s) \text{ and } \Pr(a_1 + a_2 = s)) \\
&= \sum_{0 < s \leq n} \Pr((a_1, a_2) = (ps, s - (ps + 1)) | a_1 + a_2 = s) \times \Pr(a_1 + a_2 = s) \\
&\quad (\text{For any } s, \text{ the second sum equals one.})
\end{aligned}$$

Where $\Pr((a_1, a_2) = (ps, s - (ps + 1)) | a_1 + a_2 = s)$ is the δ value for histogram when $l = 2$, the vote distribution is $\pi' \in \Pi(\{x_1, x_2\})$, where $\pi'(x_1) = \frac{p_1}{p_1 + p_2}$, and number of voters is s (we refer to Lemma 7 of the $l = 2$ case for this claim). We denote this δ by δ_s . Moreover,

$$\Pr(a_1 + a_2 = s) = \Pr(\mathbf{Bin}(n, p_1 + p_2) = s)$$

We denote $p' = p_1 + p_2$. Then, $\mathbf{Bin}(n, p')$ is the binomial distribution with n trials and probability $p' = p_1 + p_2$ (recall that $p_i = \pi(x_i)$). Then

$$\begin{aligned}
\delta &= \sum_{0 < s \leq n} \delta_s \Pr(\mathbf{Bin}(n, p') = s) \\
&= \sum_{\substack{s \geq (1 - \sqrt{\frac{3}{4}})np' \\ s \leq (1 + \sqrt{\frac{3}{4}})np'}} \Pr(\mathbf{Bin}(n, p') = s) \times \delta_s + \sum_{\substack{s < (1 - \sqrt{\frac{3}{4}})np' \\ s > (1 + \sqrt{\frac{3}{4}})np'}} \Pr(\mathbf{Bin}(n, p') = s) \times \delta_s
\end{aligned}$$

Lower bound of δ :

$$\delta \geq \sum_{\substack{s \geq (1 - \sqrt{\frac{3}{4}})np' \\ s \leq (1 + \sqrt{\frac{3}{4}})np'}} \Pr(\mathbf{Bin}(n, p') = s) \times \delta_s$$

Since δ_s decreases with larger s (more votes implies more privacy), $\delta_{(1+\sqrt{\frac{3}{4}})np'}$ is the minimum.

$$\begin{aligned} &\geq \delta_{(1+\sqrt{\frac{3}{4}})np'} \times \sum_{\substack{s \geq (1-\sqrt{\frac{3}{4}})np' \\ s \leq (1+\sqrt{\frac{3}{4}})np'}} \Pr(\mathbf{Bin}(n, p') = s) \\ &= \delta_{(1+\sqrt{\frac{3}{4}})np'} \times \left[1 - \Pr\left(\mathbf{Bin}(n, p') > \left(1 + \sqrt{\frac{3}{4}}np'\right)\right) - \Pr\left(\mathbf{Bin}(n, p') < \left(1 - \sqrt{\frac{3}{4}}np'\right)\right) \right] \end{aligned}$$

By Chernoff bound for binomial distribution, for any $0 < \beta < 1$, we have:

$$\Pr(\mathbf{Bin}(n, p') > (1 + \beta)\mu) \leq e^{-\frac{\beta^2\mu}{3}}$$

$$\Pr(\mathbf{Bin}(n, p') < (1 - \beta)\mu) \leq e^{-\frac{\beta^2\mu}{2}}$$

Where $\mu = np'$ is the mean of $\mathbf{Bin}(n, np')$. Now let $\beta = \sqrt{\frac{3}{4}}$, which is between 0 and 1. Then,

$$\begin{aligned} 1 &\geq \left[1 - \Pr\left(\mathbf{Bin}(n, p') > \left(1 + \sqrt{\frac{3}{4}}np'\right)\right) - \Pr\left(\mathbf{Bin}(n, p') < \left(1 - \sqrt{\frac{3}{4}}np'\right)\right) \right] \\ &\geq 1 - e^{-\frac{3}{4}\frac{\mu}{3}} - e^{-\frac{3}{4}\frac{\mu}{2}} \\ &= 1 - e^{-\frac{np'}{2}} - e^{-\frac{3np'}{2}} \end{aligned}$$

(For large enough n , $np' \geq 1$, so $e^{-\frac{np'}{2}} \leq e^{-1/2}$ and $e^{-\frac{3np'}{2}} \leq e^{-3/2}$)

$$\geq 1 - e^{-1/2} - e^{-3/2} \geq \frac{1}{10}$$

Which means $\left[1 - \Pr\left(\mathbf{Bin}(n, p') > \left(1 + \sqrt{\frac{3}{4}}np'\right)\right) - \Pr\left(\mathbf{Bin}(n, p') < \left(1 - \sqrt{\frac{3}{4}}np'\right)\right) \right] = \Theta(1)$.

By Stirling formula, we have

$$\begin{aligned} \delta_{(1+\sqrt{\frac{3}{4}})np'} &= \Theta\left(\frac{1}{\sqrt{\left(1 + \sqrt{\frac{3}{4}}\right)np'}}\right) \\ &= \Theta\left(\sqrt{\frac{1}{np'}}\right) \end{aligned}$$

(Recall we assumed the maximizing x, x' are x_1, x_2 , up to renaming the x_i 's, and that $p' = p_1 + p_2$)

$$= \Theta\left(\sqrt{\frac{1}{np_{\min}}}\right) \quad (\text{In general, } p_{\min} = \min_{i \neq j \in [l]}(p_i + p_j).)$$

Which gives us the lower bound $\delta \geq \Theta\left(\sqrt{\frac{1}{np_{\min}}}\right)$.

Upper bound of δ :

$$\begin{aligned} \delta &= \sum_{\substack{s \geq (1-\sqrt{\frac{3}{4}})np' \\ s \leq (1+\sqrt{\frac{3}{4}})np'}} \Pr(\mathbf{Bin}(n, p') = s) \times \delta_s \\ &\quad + \sum_{\substack{s < (1-\sqrt{\frac{3}{4}})np' \\ s > (1+\sqrt{\frac{3}{4}})np'}} \Pr(\mathbf{Bin}(n, p') = s) \times \delta_s \end{aligned}$$

$$\text{Since } \delta_s \leq 1 \text{ for all } s \text{ and } \sum_{\substack{s \geq (1-\sqrt{\frac{3}{4}})np' \\ s \leq (1+\sqrt{\frac{3}{4}})np'}} \Pr(\mathbf{Bin}(n, p') = s) \leq 1$$

$$\begin{aligned}
&\leq \max_{(1-\sqrt{\frac{3}{4}})np' \leq s \leq (1+\sqrt{\frac{3}{4}})np'} (\delta_s) + \sum_{\substack{s < (1-\sqrt{\frac{3}{4}})np' \\ s > (1+\sqrt{\frac{3}{4}})np'}} \Pr(\mathbf{Bin}(n, p') = s) \\
&= \delta_{(1-\sqrt{\frac{3}{4}})np'} + \Pr\left(\mathbf{Bin}(n, p') < \left(1 - \sqrt{\frac{3}{4}}\right)np'\right) + \Pr\left(\mathbf{Bin}(n, p') > \left(1 + \sqrt{\frac{3}{4}}\right)np'\right) \\
&\leq \delta_{(1-\sqrt{\frac{3}{4}})np'} + e^{-\frac{np'}{2}} + e^{\frac{3np'}{2}} \quad (\text{By Chernoff bound for binomial}) \\
&\leq \delta_{(1-\sqrt{\frac{3}{4}})np'} + 2\sqrt{\frac{1}{np'}} \quad (\text{Since } np' \geq 0, \text{ both } e^{-\frac{np'}{2}}, e^{\frac{3np'}{2}} \leq \sqrt{\frac{1}{np'}}) \\
&\text{By Stirling's formula, } \delta_{(1-\sqrt{\frac{3}{4}})np'} = \Theta\left(\frac{1}{\sqrt{(1-\sqrt{\frac{3}{4}})np'}}\right) \\
&= \Theta\left(\sqrt{\frac{1}{np'}}\right)
\end{aligned}$$

As is with the lower bound, in general (without assuming $(x, x') = (x_1, x_2)$), we have $p' = p_{\min} = \min_{i \neq j \in [l]} (p_i + p_j)$. Since both lower and upper bounds of δ are $\Theta\left(\sqrt{\frac{1}{np_{\min}}}\right)$, $\delta = \Theta\left(\sqrt{\frac{1}{np_{\min}}}\right)$. \square

Lemma 8 (Conditional independence). *Let $\mathcal{U} = \{x_1, \dots, x_l\}$ and $\pi \in \Delta(\mathcal{U})$. Let $\#x_i$ denote the r.v. of the number of occurrences of the vote x_i in π . Then, for all $0 \leq s \leq n$, the random variables $(\#x_1, \#x_2)$ and $(\#x_3, \dots, \#x_l)$ are independent conditioned on $\#x_1 + \#x_2 = s$. In other words, for any (t_1, \dots, t_l) such that $\sum_i t_i = n$, we have*

$$\begin{aligned}
&\Pr((\#x_1, \dots, \#x_l) = (t_1, \dots, t_l) \mid \#x_1 + \#x_2 = s) \\
&= \Pr((\#x_1, \#x_2) = (t_1, t_2) \mid \#x_1 + \#x_2 = s) \times \Pr((\#x_3, \dots, \#x_l) = (t_3, \dots, t_l) \mid \#x_1 + \#x_2 = s)
\end{aligned}$$

Proof for Lemma 8. We equivalently show that

$$\begin{aligned}
&\Pr((\#x_3, \dots, \#x_l) = (t_3, \dots, t_l) \mid \#x_1 + \#x_2 = s) \\
&= \Pr((\#x_3, \dots, \#x_l) = (t_3, \dots, t_l) \mid \#x_1 + \#x_2 = s \wedge (\#x_1, \#x_2) = (t_1, t_2))
\end{aligned} \tag{9}$$

Now, conditioned on there being exactly s people who voted x_1 or x_2 , let $D_1 > D_2 > \dots > D_s$ denote the random variables of the indices of the votes in the profile which voted for x_1 or x_2 , in ascending order. By total probability, the left hand side of Equation 9 is:

$$\begin{aligned}
&\Pr((\#x_3, \dots, \#x_l) = (t_3, \dots, t_l) \mid \#x_1 + \#x_2 = s) \\
&= \sum_{d_1 > d_2 > \dots > d_s} \Pr((\#x_3, \dots, \#x_l) = (t_3, \dots, t_l) \mid \#x_1 + \#x_2 = s \wedge (D_1, \dots, D_s) = (d_1, \dots, d_s)) \\
&\quad \times \Pr((D_1, \dots, D_s) = (d_s, \dots, d_1) \mid \#x_1 + \#x_2 = s)
\end{aligned}$$

We already assume there are exactly s votes for x_1 or x_2 , so

$$\begin{aligned}
&\Pr((\#x_3, \dots, \#x_l) = (t_3, \dots, t_l) \mid \#x_1 + \#x_2 = s) \\
&= \sum_{d_1 > d_2 > \dots > d_s} \Pr((\#x_3, \dots, \#x_l) = (t_3, \dots, t_l) \mid (D_1, \dots, D_s) = (d_1, \dots, d_s)) \\
&\quad \times \Pr((D_1, \dots, D_s) = (d_1, \dots, d_s))
\end{aligned}$$

The right hand side of Equation 9 is:

$$\begin{aligned}
&\Pr((\#x_3, \dots, \#x_l) = (t_3, \dots, t_l) \mid \#x_1 + \#x_2 = s \wedge (\#x_1, \#x_2) = (t_1, t_2)) \\
&= \Pr((\#x_3, \dots, \#x_l) = (t_3, \dots, t_l) \mid \#x_1 + \#x_2 = s \wedge (\#x_1, \#x_2) = (t_1, t_2)) \\
&= \Pr((\#x_3, \dots, \#x_l) = (t_3, \dots, t_l) \mid (\#x_1, \#x_2) = (t_1, t_2)) \quad (\text{Since we assume } t_1 + t_2 = s)
\end{aligned}$$

Rule	Winner	Mean Square Error ($n \in [50]$)
Borda	$\delta(n) = \frac{1}{\sqrt{1.347n + 0.5263}}$	0.0566844201243
STV	$\delta(n) = \frac{1}{\sqrt{1.495n + 0.02669}}$	0.0542992943035
Maximin	$\delta(n) = \frac{1}{\sqrt{1.553n + 4.433}}$	0.0377631805983
Plurality	$\delta(n) = \frac{1}{\sqrt{1.717n - 0.09225}}$	0.0477175838906
2-approval	$\delta(n) = \frac{1}{\sqrt{1.786n + 0.3536}}$	0.0454223047191

Table 2: δ values in $(0, \delta, \Delta)$ -eDDP for some commonly-used voting rules under the i.i.d. uniform distribution. $m = 3$ and $n = 10$ to 50.

$$\begin{aligned}
&= \sum_{d_1 > d_2 > \dots > d_s} \Pr((\#x_3, \dots, \#x_l) = (t_3, \dots, t_l) \mid (\#x_1, \#x_2) = (t_1, t_2) \wedge (D_1, \dots, D_s) = (d_1, \dots, d_s)) \\
&\quad \times \Pr((D_1, \dots, D_s) = (d_1, \dots, d_s) \mid (\#x_1, \#x_2) = (t_1, t_2)) \quad (\text{By total probability,})
\end{aligned}$$

Since each vote is independent, $(\#x_3, \dots, \#x_l)$ is independent of $(\#x_1, \#x_2)$. Moreover, the vote indices (D_1, \dots, D_s) are independent of $(\#x_1, \#x_2)$. As votes are i.i.d., $(\#x_1, \#x_2)$ does not depend on the value of (d_1, \dots, d_s) . Thus,

$$\begin{aligned}
&\Pr((\#x_3, \dots, \#x_l) = (t_3, \dots, t_l) \mid \#x_1 + \#x_2 = s \wedge (\#x_1, \#x_2) = (t_1, t_2)) \\
&= \sum_{d_1 > d_2 > \dots > d_s} \Pr((\#x_3, \dots, \#x_l) = (t_3, \dots, t_l) \mid (D_1, \dots, D_s) = (d_1, \dots, d_s)) \\
&\quad \times \Pr((D_1, \dots, D_s) = (d_1, \dots, d_s))
\end{aligned}$$

This concludes that the left hand side and right hand side probabilities of Equation 9 are equal. The random variables $(\#x_1, \#x_2)$ are independent conditioned on $(\#x_1, \#x_2)$. \square

E Concrete Estimate of the Privacy Parameters

In this section 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 by doing an exhaustive search of all 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, by the value a for outputting the winner. The larger the a , the smaller the δ value and more private. The resulting ranking from most to least private is:

$$2\text{-approval} \triangleright \text{Plurality} \triangleright \text{Maximin} \triangleright \text{STV} \triangleright \text{Borda}$$

We show in Table 2 the fitted δ curves with the mean square error in the fit.

Additional References

[Mossel *et al.*, 2013] Elchanan Mossel, Ariel D. Procaccia, and Miklos Z. Racz. A Smooth Transition From Powerlessness to Absolute Power. *Journal of Artificial Intelligence Research*, 48(1):923–951, 2013