Efficient, Private, and ϵ-Strategyproof Elicitation of Tournament Voting Rules

David T. Lee
Stanford University
davidtlee@stanford.edu

Abstract

Voting is commonly used as a method for aggregating information in crowdsourcing and human computation. In many settings, one would like to use voting rules which can be efficiently elicited, preserve voter privacy, and are robust to strategic manipulation. In this paper, we give algorithms which elicit approximate winners in a way which provably satisfies all three of these requirements simultaneously. Our results hold for *tournament voting rules*, which we define to be the voting rules which can be expressed solely as a function of the table of pairwise comparisons containing the number of voters preferring one candidate to another[1]. Tournament voting rules include many common voting rules such as the Borda, Copeland, Maximin, Nanson, Baldwin, Kemeny-Young, Ranked Pairs, Cup, and Schulze voting rules. Our results significantly expand the set of voting rules for which efficient elicitation was known to be possible and improve the known approximation factors for ϵ-strategyproof voting in the regime where the number of candidates is large.

1 Introduction

Crowdsourcing and human computation systems often contain components which aggregate collected information to arrive at rankings over a set of alternatives. One method of aggregation is to use social choice functions, commonly known as voting rules [Brams and Fishburn, 2002; Brandt *et al.*, 2012]. In this setting, alternatives are viewed as candidates in an election and participants are viewed as voters, each of which submits a ranking over the alternatives according to their beliefs or preferences. A voting rule is then a (possibly stochastic) function which maps the set of rankings to an output winner.

One scenario in which voting rules have been applied to crowdsourcing is when a ground truth exists [Mao *et al.*, 2013], for which voting rules can be interpreted as maximum likelihood estimators for various noise models over user beliefs [Conitzer and Sandholm, 2013]. Another natural scenario for the application of voting rules, due to its democratic nature, is that of crowdsourcing for political decision-making [Lee *et al.*, 2014]. For both of these scenarios, three important problems exist:

- *Preference Elicitation:* Can the correct output ranking be found without requiring all voters to submit their complete preferences, e.g. by only eliciting a small sample of voters or a small number of pairwise comparisons?
- *Strategic Manipulation:* Can the correct output ranking be found when voters are strategic agents who may submit false rankings in order to improve their own utility?
- *Participant Privacy:* Can votes be elicited in a manner which preserves participant privacy?

These are particularly important when crowdsourcing is applied to political decision-making. Inefficient elicitation can render the crowdsourcing effort infeasible since the number of alternatives produced by the crowd may be too large for voters to even pick their top choices, not to mention ranking them. Poor manipulation properties can result in voters who strategically optimize for their own preferences at the cost of the general electorate. Loss of privacy can result in sensitive voter information being leaked to the public.

We specifically highlight the importance of privacy in voting since it is the least studied of the three in existing literature. It is self-evident that privacy is important when a vote correlates clearly with sensitive information such as sexual orientation. However, in an information-rich world, it has become increasingly possible to obtain accurate predictions of a sensitive information from seemingly benign information [Narayanan and Shmatikov, 2008]. As a result, this age is also one in which privacy properties are increasingly important to participants of social systems.

Both preference elicitation and strategic manipulation are well-studied problems in the social choice literature and have been found to be subject to impossibility results. For many common voting rules, finding out how to elicit optimally is NP-complete [Conitzer and Sandholm, 2002]. Moreover, for many common voting rules, it is impossible to reduce the number of bits necessary to find the output winner or approximate winner by more than a constant factor [Conitzer and Sandholm, 2005; Service and Adams, 2012]. For strate-

[1] Tournament voting rules can also be described as C1 or C2 functions according to Fishburn's classification [Fishburn, 1977].

gic agents, it is known that any strategyproof voting rule is a dictatorship [Gibbard, 1973; Satterthwaite, 1975].

Many approaches have been taken to circumvent these impossibility results. Two of these are particularly relevant to our work. Lee et al. gave definitions for ϵ-Borda and ϵ-Condorcet winners, and showed that they could be elicited with $\tilde{O}(\frac{m}{\epsilon^2})$ and $\tilde{O}(\frac{m}{\epsilon^4})$ pairwise comparisons[2], a significant improvement over the $\Omega(mn)$ pairwise comparisons previously thought to be necessary. They also proposed a broader definition for approximate voting: an alternative is an ϵ-approximate winner if it could have been the winner given a change in at most an ϵ fraction of the pairwise comparisons involving any alternative [Lee *et al.*, 2014]. This definition, formally given in the Model section, is applicable to all voting rules which can be defined solely in terms of the *preference tournament*, which informally consists of the number of voters preferring any given candidate to another. Lee et al. show that any alternative which is an ϵ-Borda or ϵ-Condorcet winner under their initial definition is also one under this broader definition; however, they do not describe how to elicit approximate winners for other voting rules.

Birrell and Pass defined a voting rule to be ϵ-strategyproof if a non-truthful voter can only improve his utility by ϵ, regardless of the rankings provided by other voters [Birrell and Pass, 2011]. They showed that approximate winners[3] could be found in an ϵ-strategyproof manner; however, when the number of candidates is large ($m \geq \sqrt{\epsilon n}$, where m and n are the number of candidates and voters respectively), their approximation obtained is poor in that it admits any alternative as an "approximate winner".

1.1 Our Contribution

In this paper, we show that it is possible to elicit ϵ-approximate winners (as defined in Lee et al. 2014) for all tournament voting rules in a way which is efficient, preserves voter privacy (as formalized by the notion of differential privacy, detailed later), and is ϵ-strategyproof.

- We first give a variant of the algorithm in Lee et al. which finds an ϵ-Borda winner in $\tilde{O}(\frac{m}{\epsilon^2})$ total comparisons (an expected $\tilde{O}(\frac{1}{\epsilon^2})$ per voter). Our variant uses the same number of comparisons, but is also ϵ-strategyproof and ϵ-differentially private.
- We then show that sampling $\tilde{O}(\frac{1}{\epsilon^2})$ voter rankings, a process which takes $\tilde{O}(\frac{m}{\epsilon^2})$ total pairwise comparisons, is sufficient for finding an ϵ-approximate winner for *any tournament voting rule*. This includes the Borda, Copeland, Maximin, Nanson, Baldwin, Kemeny-Young, Ranked Pairs, Cup, and Schulze voting rules, most of which were previously thought to require $\Omega(mn)$ pairwise comparisons [Conitzer and Sandholm, 2005]. We note, however, that the comparisons elicited in this case are all concentrated in a few voters who must give their full rankings.
- Our third algorithm finds an ϵ-approximate winner for *any tournament voting rule* using $\tilde{O}(\frac{m\sqrt{n}}{\epsilon})$ total pairwise comparisons. Like the first, this algorithm also spreads the burden equally among all of the voters (an expected $\tilde{O}(\frac{\sqrt{n}}{\epsilon})$ per voter).
- Finally, we show that our algorithms also are resistant to collusion in that the approximation factors degrade gracefully when multiple voters collude together.

Beyond the contributions we make in extending efficient elicitation to all tournament voting rules, our algorithms are also the first to obtain ϵ-strategyproofness when the number of candidates are large. Specifically, our results give non-trivial approximations so long as $m \leq \tilde{O}(n)$, as compared to prior results which gave non-trivial approximations for $m \leq \sqrt{\epsilon n}$. Ours are also the first which obtain efficient elicitation, ϵ-differential privacy, and ϵ-strategyproofness simultaneously (the algorithms given in Lee et al. are not strategyproof, and the algorithms given in Birrell and Pass require complete preference profiles).

We note that, due to space constraints, detailed proofs can be found in the Appendix section of the full version[4].

2 Related Work

There are many other approaches that have been taken in the literature on circumventing impossibility results in preference elicitation. One approach considers elicitation when the set of rankings are restricted to single-peaked preferences [Conitzer, 2009; Goel and Lee, 2012]. Another approach studies elicitation experimentally using heuristics inspired from machine learning or assumes that preferences are drawn from a given distribution [Lu and Boutilier, 2011a; 2011b; Kalech *et al.*, 2011; Oren *et al.*, 2013; Caragiannis *et al.*, 2013; Chevaleyre *et al.*, 2011]. We specifically point out [Caragiannis *et al.*, 2013] in which they show that sampling $O(\frac{1}{\epsilon^2}\log\frac{m}{\delta})$ voter rankings is sufficient for determining the exact output ranking when voters rankings are assumed to be drawn from a Mallows model. This is similar to our second algorithm, except that we aim only at finding an ϵ-approximate winner, but in the more general case when voter rankings are arbitrary.

There have also been many approaches to circumvent impossibility results in strategic manipulation. Gibbard was the first to consider randomized voting rules [Gibbard, 1977], but showed that only trivial randomized voting rules were strategyproof. Procaccia studied the approximation factors possible with these voting rules [Procaccia, 2010], but showed that the approximations are not good. Other approaches include restricting the domain of voter preferences, e.g. to single-peaked preferences [Moulin, 1980], studying voting rules which are computationally difficult to manipulate [Bartholdi *et al.*, 1989; Conitzer and Sandholm, 2006], and considering voters who only have partial probabilistic information on the preference profile [Lu *et al.*, 2012].

[2]The notation $\tilde{O}(\cdot)$ notation is the same as the $O(\cdot)$ notation, but hides logarithmic factors.

[3]In Birrell and Pass, approximate winners are defined differently than in Lee et al: an alternative is an ϵ-approximate winner if it could have been the winner given any change in at most ϵ rankings.

[4]http://stanford.edu/~dtlee88/papers/ijcai-2015-ef-full.pdf

3 Model

Let C denote the set of alternatives and V the set of participants. Let m and n denote the number of alternatives and participants respectively. Each participant i has a ranking, i.e. a linear ordering, over the set of alternatives representing his or her preference. Using the notation of [Conitzer and Sandholm, 2005], we use $\succ_i$ to denote the ranking of participant i. If participant i prefers alternative x to y, we denote this by $x \succ_i y$. The set of all rankings is called a preference profile, and is denoted by $\vec{\succ}$. A voting rule is a (possibly stochastic) function which maps a given preference profile to a single winning alternative[5]. We will use the term *voting algorithm* to refer to the elicitation scheme inducing a voting rule.

A preference tournament T is a complete, directed graph on the set of alternatives. The weight of an edge (x, y) is denoted by $T_{xy} = |\{i \mid x \succ_i y\}|$, which represents the number of participants who prefer alternative x to y. We use T_x to denote the vector of length $m-1$ containing the weights from x to all other candidates. Preference tournaments are also commonly depicted using a table as in Example 1. We define a tournament voting rule to be any voting rule which can be defined solely on the preference tournament.

3.1 Approximate winners

We use the definition of an *approximate winner* used in Lee et al. It informally defines an alternative $\hat{w}$ to be an ϵ-approximate winner if it could have resulted from a preference tournament T' which differs from the true preference tournament T in at most ϵ fraction of the comparisons involving any candidate. Thus, an approximate winner can roughly be interpreted as an alternative that could have been the winner given a small perturbation in voter input.

Definition 1. *For a tournament voting rule f and preference tournament T, $\hat{w}$ is an ϵ-f winner if*

- *$\exists\, T'$ such that $\hat{w} = f(T')$ and,*
- *$\|T_x - T'_x\|_1 \leq \epsilon n(m-1)$ for any x.*

When the voting rule f is clear from the context, we simply refer to $\hat{w}$ as an ϵ-approximate winner.

The following example illustrates this concept.

Example 1. *Suppose that* 100 *participants supply rankings of* 3 *alternatives, and that these rankings result in the following preference tournament (the values in the table represents the number of participants preferring the left alternative to the top alternative):*

	A	B	C
A	-	52	61
B	48	-	71
C	39	29	-

The Borda winner is the alternative which wins the largest number of pairwise comparisons. A, B, and C win 113, 119, *and* 68 *comparisons respectively, so that the Borda winner is B. The Copeland winner is the alternative which wins the largest number of pairwise elections. A, B, and C win* 2, 1, *and* 0 *pairwise elections respectively, so A is the Copeland winner.*

Simple calculations show us that A is a 0.02*-Borda winner since swapping* 4 *of the comparisons preferring B to A would make A the winner, and there are* 200 *comparisons involving any given alternative. Similarly, C is a* 0.13*-Borda winner since swapping* 26 *of the comparisons preferring B to C would make C the winner. We also see that B is a* 0.015*-Copeland winner since swapping* 3 *of the comparisons preferring A to B would make B the winner. Similarly, C is a* 0.17*-Copeland winner since swapping* 12 *of the comparisons preferring A to C and* 22 *of the comparisons preferring B to C would make C the winner.*

We can see that the definitions capture the intuition that A is a good approximate Borda winner, B is a good approximate Copeland winner, but C is a bad approximate Borda and Copeland winner. □

3.2 Approximate strategyproofness

We consider participants who have utilities over the set of alternatives. The utility of voter i is denoted by $u_i(\cdot)$ and is assumed to satisfy $0 \leq u_i(x) \leq 1$ for all x. It is *consistent* with a ranking $\succ_i$ if $u_i(x) \geq u_i(y)$ if $x \succ_i y$. An ϵ-strategyproof voting rule is defined according to the definition given by Birrell and Pass, i.e. regardless of how a participant changes his votes (including if he knows all the votes cast by others), he cannot increase his expected utility by more than ϵ.

Definition 2. *A voting rule f is ϵ-strategyproof if for all voters i, all preference profiles $\vec{\succ}$, all alternative preference profiles $\vec{\succ}'$ which differ only in voter i's ranking, and all utility functions u_i consistent with $\succ_i$,*

$$\mathbb{E}[u_i(f(\vec{\succ}'))] \leq \mathbb{E}[u_i(f(\vec{\succ}))] + \epsilon$$

where the expectation $\mathbb{E}[\cdot]$ is over any randomness in the voting rule f. A voting algorithm is ϵ-strategyproof if the voting rule it induces is ϵ-strategyproof.

3.3 Differential privacy

Differential privacy is a standard way to formalize the loss of privacy of an individual after a computation using the individual's personal information is published [Dwork and Roth, 2014]. In the context of voting, loss of privacy occurs when others are able to determine the private preference rankings of an individual, which could be sensitive depending on the issue being voted on.

Informally, a differentially private voting rule is one in which an intruder is unable to determine an individual's ranking based on the published output ranking, even if he knows all the other voter rankings.

Definition 3. *A voting rule f is ϵ-differentially private if for all preference profiles $\vec{\succ}$ and $\vec{\succ}'$ that differ on a single voter's ranking, and all $S \subseteq C$,*

$$\mathbf{Pr}[f(\vec{\succ}) \in S] \leq (1+\epsilon)\mathbf{Pr}[f(\vec{\succ}') \in S].$$

A voting algorithm is ϵ-differentially private if the voting rule it induces is ϵ-differentially private.

[5] For some classical voting rules, such as Copeland, one would have to specify a tie-breaking rule to ensure a single winner.

Though we are not aware of work on differentially private voting, it was pointed out by Birrell and Pass that the approach they take to achieving ϵ-strategyproofness, that of adding noise to the original function computed, is a common approach taken to achieving differential privacy.

4 Eliciting an approximate Borda winner

In Lee et al., an algorithm was given which found an ϵ-Borda winner with probability $1-\delta$ using $N = O(\frac{m}{\epsilon^2}\log\frac{m}{\delta})$ comparisons. In this section, we show that a variant of their algorithm also obtains an ϵ-Borda winner using $N = O(\frac{m}{\epsilon^2}\log\frac{m}{\delta})$ comparisons, but in a way which guarantees ϵ-differential privacy and ϵ-strategyproofness when $n \geq \frac{1}{\epsilon}N$, i.e. $m \leq \tilde{O}(\epsilon^3 n)$. We emphasize that it is able to achieve these while also spreading out the work evenly among all the voters.

The algorithm (Alg. 1) works in the following way. With probability $\gamma = \min(\frac{1}{2}\delta, \frac{1}{m})$, return a random alternative as the winner. Otherwise, define N_x for $x \in C$ to be i.i.d. Binomial($N, \frac{2}{m}$) random variables[6]. For each x, sample N_x random comparisons in which a random voter compares x and another random candidate. Let S_x denote the number of these comparisons for which x wins. Then return the alternative which maximizes S_x as the winner (ties broken uniformly at random).

ALGORITHM 1: BORDA+RANDOM

Input: alternatives C, voters V, parameters N, γ
Output: A winning alternative
if $U \sim$ *Uniform*$[0,1) \leq \gamma$ **then**
    return an alternative uniformly at random;
else
    return an alternative according to BORDA(C, V, N);

ALGORITHM 2: BORDA

Input: alternatives C, voters V, parameter N
Output: A winning alternative
for $x \in C$ **do**
    Let $N_x \sim$ Binomial($N, \frac{2}{m}$);
    Initialize $S_x = 0$;
    for $i \leftarrow 1$ **to** N_x **do**
        Sample $y \in C \setminus \{x\}$ and $v \in V$ uniformly at random;
        if $x \succ_v y$ **then**
            $S_x = S_x + 1$;
return the alternative maximizing S_x (ties broken uniformly);

Theorem 1. *For any $\epsilon, \delta \in (0,1)$, Algorithm 1 with $N = O(\frac{m}{\epsilon^2}\ln\frac{m}{\delta})$ and $\gamma = \min(\frac{1}{2}\delta, \frac{1}{m})$ returns an ϵ-Borda winner with probability at least $1-\delta$. Moreover, it is also ϵ-differentially private and ϵ-strategyproof when $n \geq \frac{1}{\epsilon}N$.*[7]

[6]Choosing N_x i.i.d. was necessary for our proof, but it may be that a proof exists which can relax this. The binomial parameter $\frac{2}{m}$ was chosen so that $\mathbb{E}[\sum_x S_x] = N$.

[7]For simplicity, we use a single parameter ϵ for approximate winners, differential privacy, and approximate strategyproofness. It would not be hard to generalize our results to different parameters.

Proof. We first show that Algorithm 1 returns an ϵ-Borda winner with probability $1-\delta$. By Lemma 1, the alternative returned by BORDA (Algorithm 2) is an ϵ-Borda winner with probability at least $1-\frac{\delta}{2}$. Since the probability of returning an alternative randomly is only $\gamma \leq \frac{\delta}{2}$, our result follows by a union bound.

Lemmas 2 and 5 show that Algorithm 1 is ϵ-differentially private and ϵ-strategyproof. Lemma 2 first demonstrates that BORDA (Algorithm 2), which we denote by $\mathcal{B}$, satisfies

$$\mathbf{Pr}[\mathcal{B}(\vec{\succ}) = w] \leq (1+\epsilon)\mathbf{Pr}[\mathcal{B}(\vec{\succ}') = w] + \epsilon\gamma\frac{1}{m}$$

when $n \geq \frac{1}{\epsilon}N$ (see Definition 2). Lemma 5 then obtains our final result from this condition. These Lemmas can be found in the Appendix of the full version (see note in Section 1.1). □

4.1 Intuition

The intuition for this result is that, for alternatives which have a non-trivial probability of winning, their probability of winning is not greatly affected (by more than a $(1+\epsilon)$ factor) by a single voter's ranking. This is due to the random nature of the elicitation scheme. For alternatives which have a very small probability of winning under the elicitation scheme, their probability of winning is dominated by the probability of choosing a random alternative as the winner. This intuition is broadly applicable, and is the method used in our later proofs.

In this case, it is possible to show that mixing Algorithm 2 with choosing a random alternative is necessary to obtain ϵ-differential privacy. However, Algorithm 2 alone (without choosing a random alternative) is sufficient for obtaining ϵ-strategyproofness. This will not always be true in the general case. These statements are formalized in Lemmas 6 and 7 in the Appendix of the full version (see note in Section 1.1).

5 Eliciting approximations by sampling voters

In [Caragiannis *et al.*, 2013], it was shown that $O(\frac{1}{\epsilon^2}\log\frac{m}{\delta})$ sampled voter rankings allows one to find the true ranking when voters are assumed to be drawn from a Mallows distribution. In this section, we show that it is possible to find the ϵ-approximate winner of any tournament voting rule by sampling $N = O(\frac{1}{\epsilon^2}\log\frac{m}{\delta})$ voter rankings (Alg. 3). By using standard sorting algorithms, such as quicksort, one can elicit each ranking using $O(m\log m)$ pairwise comparisons for a total of $\tilde{O}(\frac{m}{\epsilon^2})$ comparisons. Our results hold for arbitrary voter rankings. We also show that this simple procedure is both ϵ-differentially private and ϵ-strategyproof when $n \geq \frac{m}{\epsilon^2}N$, i.e. $m \leq \tilde{O}(\epsilon^4 n)$.

Theorem 2. *For any $\epsilon, \delta \in (0,1)$, Algorithm 4 with $N = O(\frac{1}{\epsilon^2}\ln\frac{m}{\delta})$ and any tournament voting rule f returns an ϵ-f winner with probability $1-\delta$.*

Proof. Let T denote the true preference tournament of the voters, i.e. T_{xy} is the number of voters preferring alternative x to y. Recall that $\hat{w}$ is an ϵ-f winner if another preference

ALGORITHM 3: SAMPLE-VOTERS+RANDOM

Input: alternatives C, voters V, a tournament voting rule f, and parameters N, γ
Output: A winning alternative
if $U \sim \textit{Uniform}[0,1) \leq \gamma$ **then**
  return an alternative uniformly at random;
else
  return an alternative according to SAMPLE-VOTERS(C, V, f, N);

ALGORITHM 4: SAMPLE-VOTERS

Input: alternatives C, voters V, a tournament voting rule f, and parameter N
Output: An output winner
$Z_{xy} = 0$ for all $x, y \in C$;
Let S be N voters sampled with replacement from V;
for $v \in S$ **do**
  Use quicksort to elicit participant v's ranking of C;
  for $x, y \in C$ **do**
    if $x \succ_v y$ **then**
      $Z_{xy} = Z_{xy} + 1$;
    else
      $Z_{yx} = Z_{yx} + 1$;
Normalize Z by letting $T'_{xy} = \frac{n}{N} \cdot Z_{xy}$ for all $x, y \in C$;
return $f(T')$;

tournament T' exists such that $\hat{w} = f(T')$ and $\|T'_x - T_x\|_1 \leq \epsilon n(m-1)$ for any x. In Algorithm 4, the winner returned was defined as the output of $f(T')$, where T' is constructed from the sampled voter rankings. We show in Lemma 8 that T' does indeed satisfy $\|T'_x - T_x\|_1 \leq \epsilon n(m-1)$ with probability at least $1-\delta$, which gives us our result (see Appendix in full version linked to in Section 1.1). □

The following theorem shows that mixing SAMPLE-VOTERS with some small probability of returning a random alternative achieves ϵ-differential privacy and ϵ-strategyproofness. We note that a slight sacrifice is made in that an ϵ-approximate winner is only achieved in expectation (as opposed to with high probability).

Theorem 3. *For any $\epsilon \in (0,1)$, Algorithm 3 with $N = O(\frac{1}{\epsilon^2} \ln \frac{m}{\epsilon})$, $\gamma = \epsilon$, and any tournament voting rule f returns an ϵ-f winner in expectation. Moreover, it is also ϵ-differentially private and ϵ-strategyproof when $n \geq \frac{2m}{\epsilon^2} N$.*

Proof. Let $\epsilon' = \frac{\epsilon}{3}$. A winner is returned from SAMPLE-VOTERS with probability $1 - \epsilon'$. This winner is an ϵ'-approximate winner with probability at least $1-\epsilon'$ (from Theorem 2). In any other scenario, the alternative returned must be a 1-approximate winner at worst (since all comparisons can be changed with $\epsilon = 1$). Then the expected approximation factor is

$$(1-\epsilon')[(1-\epsilon')\epsilon' + \epsilon' \cdot 1] + \epsilon' \cdot 1 \leq 3\epsilon' = \epsilon$$

The proof that Algorithm 3 is ϵ-differentially private and ϵ-strategyproof is similar to that of Theorem 1. We first show that SAMPLE-VOTERS (Algorithm 4), which we denote by $\mathcal{B}$, satisfies

$$\mathbf{Pr}[\mathcal{B}(\vec{\succ}) = w] \leq (1+\epsilon)\mathbf{Pr}[\mathcal{B}(\vec{\succ}') = w] + \epsilon^2 \frac{1}{m}$$

when $n \geq \frac{m}{\epsilon^2} N$ (Lemma 9). Then we can apply Lemma 5 to obtain ϵ-differential privacy and ϵ-strategyproofness. These Lemmas can be found in the Appendix of the full version as noted in Section 1.1. □

It is useful to note that one could still obtain an ϵ-approximate winner with high probability ($\geq 1-\delta$) if one chooses $\gamma = O(\delta)$. However, this would require n to satisfy $n \geq \frac{m}{\epsilon\delta} N$ which restricts the applicability of such a result since δ is typically very small. We believe that it is possible to obtain a high probability result without further restricting n for many common voting rules by following the same intuition described previously. However, such a proof would likely need to be more specifically tailored to individual voting rules.

6 Eliciting approximations by sampled subsets

Even though Algorithm 3 is efficient, ϵ-strategyproof, and ϵ-differentially private for all tournament voting rules, it relies on a small number of voters who must contribute a large amount of information. While this may be useful in some situations, it is sometimes desirable to not require a large amount of information from any voter. This is particularly important when the number of alternatives is so large that it is not feasible for voters to even look through all the alternatives.

In this section, we show that one can efficiently elicit approximate winners for any tournament voting rule, while also spreading out the workload evenly among voters. However, it is still an open problem whether the given algorithm is able to achieve ϵ-differential privacy and ϵ-strategyproofness.

The algorithm we use asks each participant to provide a full ranking of k randomly chosen alternatives, a task which can be achieved in $O(k \log k)$ comparisons with standard sorting algorithms. We show that when $k = \tilde{O}(\frac{m}{\epsilon\sqrt{n}})$, it is possible to approximate all tournament voting rules. This can be elicited in a total of $\tilde{O}(\frac{m\sqrt{n}}{\epsilon})$ pairwise comparisons.

Theorem 4. *For any $\epsilon, \delta \in (0,1)$, Algorithm 5 with $k = O(\frac{m}{\epsilon\sqrt{n}} \sqrt{\log \frac{m}{\delta}})$ and any tournament voting rule f returns an ϵ-f winner with probability at least $1-\delta$.*

Proof. The proof is similar to that of Theorem 2. All we need to show is that T' as constructed in Algorithm 5 satisfies $\|T'_x - T_x\|_1 \leq \epsilon n(m-1)$ with probability at least $1-\delta$. We show this in Lemma 10 (see Appendix in full version linked to in Section 1.1). □

7 Group strategyproofness

Birrell and Pass also defined approximate group strategyproofness to show that the approximation guarantees degraded gracefully if the algorithms were used to protect against collusion from multiple players. In this section, we

ALGORITHM 5: SAMPLED-SUBSETS

Input: alternatives C, voters V, a tournament voting rule f, and parameter k
Output: An output winner
$Z_{xy} = 0$ for all $x, y \in C$;
for $v \in V$ **do**
  Let C_v denote k alternatives sampled without replacement;
  Use quicksort to elicit participant v's ranking of C_v;
  for $x, y \in C_v$ **do**
    if $x \succ_v y$ **then**
      $Z_{xy} = Z_{xy} + 1$;
    else
      $Z_{yx} = Z_{yx} + 1$;
Normalize Z by letting $T'_{xy} = n \cdot \frac{Z_{xy}}{Z_{xy}+Z_{yx}}$ for all $x, y \in C$;
return $f(T')$;

show that our algorithms also degrade quite nicely as the number of colluding voters increases.

Definition 4. *A voting rule f is (t, ϵ)-strategy-proof if for any set of t voters D, all preference profiles $\vec{\succ}$, all alternative preference profiles $\vec{\succ}'$ which differ only in the rankings of voters in D, and all utility functions u_i,*

$$\sum_{i \in D} \mathbb{E}[u_i(f(\vec{\succ}'))] \leq \sum_{i \in D} \mathbb{E}[u_i(f(\vec{\succ}))] + \epsilon$$

Intuitively, (t, ϵ)-strategyproofness states that no group of t voters can improve their *collective* utility by more than ϵ by deviating, regardless of the ranking other voters give.

The following theorems show that decreasing the number of pairwise comparisons used by a factor of t allows one to achieve (t, ϵ)-strategyproofness for Algorithms 1 and 3. The cost is that the approximation factor (for the winner obtained) goes up by a factor of $\sqrt{t}$. These results are shown with the same assumptions as in the prior results, i.e. $m \leq \tilde{O}(\epsilon^3 n)$ and $m \leq \tilde{O}(\epsilon^4 n)$ respectively. The intuition for this is that decreasing the number of comparisons adds more noise to the winner, which leads to stronger strategyproofness.

Theorem 5. *For any $\epsilon, \delta \in (0, 1)$, Algorithm 1 with $N = O(\frac{m}{t\epsilon^2} \ln \frac{m}{\delta})$ and $\gamma = \min(\frac{1}{2}\delta, \frac{1}{m})$ returns an $\epsilon\sqrt{t}$-Borda winner with probability at least $1 - \delta$. Moreover, it is also ϵ-differentially private and ϵ-strategyproof when $n \geq \frac{t}{\epsilon} N$.*

Proof. The proof is essentially identical to that of Theorem 1. The key difference is in Lemma 4 in which one wants to show that $f_k \leq (1 + \epsilon) f'_k$. With t colluding voters, $|p'_i - p_i| \leq \frac{2t}{mn}$, so that $f_k \leq f'_k \exp\{O(tN/n)\}$. □

Theorem 6. *For any $\epsilon \in (0, 1)$, Algorithm 3 with $N = O(\frac{1}{t\epsilon^2} \ln \frac{m}{\epsilon})$, $\gamma = \epsilon$, and any tournament voting rule f returns an $\epsilon\sqrt{t}$-f winner in expectation. Moreover, it is also ϵ-differentially private and ϵ-strategyproof when $n \geq \frac{2mt}{\epsilon^2} N$.*

Proof. The proof is essentially identical to that of Theorem 3. The key difference is in Lemma 9 in which one wants to find the probability that none of the t colluding voters get sampled. Let D be any set of t voters. Then,

$$\mathbf{Pr}[D \not\subset S] = \left(1 - \frac{t}{n}\right)^N \geq e^{-\frac{2tN}{n}} \geq e^{-\epsilon^2/m} \geq 1 - \frac{\epsilon^2}{m}$$

The rest of the arguments are identical. □

8 Discussion and Future Work

There are a few important observations to make with respect to limitations of the presented algorithms. First, our algorithms do not circumvent intractability. That is, for voting rules such as Kemeny, calculating $f(T')$ is intractable, which means that Algorithms 4 and 5 are also intractable.

Second, we note that in the definition of an approximate winner $\hat{w}$, it was only required that $\hat{w}$ be the output of a preference tournament T' which was close to the actual preference tournament T. It was pointed out in Lee et al. that this does not necessarily imply that a preference profile (set of rankings) exists which induces T'. An interesting direction is to generalize these results to a definition of an approximate winner which is defined directly on preference profiles using the Kendall-tau distance (see the discussion in Lee et al.).

Finally, we note that while our results scale well with the number of candidates, we believe that there is still room for improvement in the dependence on ϵ. For example, in Theorem 3, ϵ-strategyproofness and ϵ-differential privacy are only attainable when $n \geq O(\frac{1}{\epsilon^4})$, which may only be feasible for $\epsilon \sim 0.1$.

Acknowledgments

This work was supported in part by ARO Grant No. 116388, an NSF graduate research fellowship, and a Brown Institute for Media Innovation Magic Grant.

References

[Bartholdi *et al.*, 1989] John Bartholdi, Craig Tovey, and Michael Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6:227–241, 1989.

[Birrell and Pass, 2011] Eleanor Birrell and Rafael Pass. Approximately strategy-proof voting. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, 2011.

[Brams and Fishburn, 2002] Steven J. Brams and Peter C. Fishburn. Voting procedures. In Kenneth J. Arrow, Amartya Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, pages 173–236. Elsevier, 2002.

[Brandt *et al.*, 2012] Felix Brandt, Vincent Conitzer, and Ulrich Endriss. Computational social choice. In *Multiagent Systems*. MIT Press, 2012.

[Caragiannis *et al.*, 2013] Ioannis Caragiannis, Ariel D. Procaccia, and Nisarg Shah. When do noisy votes reveal the truth? In *Proceedings of the 14th ACM Conference on Electronic Commerce (ACM-EC)*, 2013.

[Chevaleyre *et al.*, 2011] Yann Chevaleyre, Jerome Lang, Nicolas Maudet, and Jerome Monnot. Compilation and communication protocols for voting rules with a dynamic set of candidates. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 2011.

[Conitzer and Sandholm, 2002] Vincent Conitzer and Tuomas Sandholm. Vote elicitation: Complexity and strategy-proofness. In *Proceedings of the 17th AAAI Conference*, 2002.

[Conitzer and Sandholm, 2005] Vincent Conitzer and Tuomas Sandholm. Communication complexity of common voting rules. In *Proceedings of the 6th ACM Conference on Electronic Commerce (ACM-EC)*, pages 78–87, 2005.

[Conitzer and Sandholm, 2006] Vincent Conitzer and Tuomas Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the 21st AAAI Conference*, pages 627–634, 2006.

[Conitzer and Sandholm, 2013] Vincent Conitzer and Tuomas Sandholm. Common voting rules as maximum likelihood estimators. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, 2013.

[Conitzer, 2009] Vincent Conitzer. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research*, 35:161–191, 2009.

[Dwork and Roth, 2014] Cynthia Dwork and Aaron Roth. *The Algorithmic Foundations of Differential Privacy*. NOW Publishers, 2014.

[Fishburn, 1977] Peter C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.

[Gibbard, 1973] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.

[Gibbard, 1977] Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, 1977.

[Goel and Lee, 2012] Ashish Goel and David T. Lee. Triadic consensus: A randomized algorithm for voting in a crowd. In *Proceedings of the 8th Workshop on Internet and Network Economics (WINE)*, 2012.

[Kalech *et al.*, 2011] Meir Kalech, Sarit Kraus, Gal A. Kaminka, and Claudia V. Goldman. Practical voting rules with partial information. *Journal of Autonomous Agents and Multi-Agent Systems*, 22(1):151–182, 2011.

[Lee *et al.*, 2014] David T. Lee, Ashish Goel, Tanja Aitamurto, and Helene Landemore. Crowdsourcing for participatory democracies: Efficient elicitation of social choice functions. In *The 4th Workshop on Social Computing and User Generated Content (SCUGC)*, 2014.

[Lu and Boutilier, 2011a] Tyler Lu and Craig Boutilier. Robust approximation and incremental elicitation in voting protocols. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2011.

[Lu and Boutilier, 2011b] Tyler Lu and Craig Boutilier. Vote elicitation with probabilistic preference models: empirical estimation and cost tradeoffs. In *Proceedings of the 2nd International Conference on Algorithmic Decision Theory (ADT)*, 2011.

[Lu *et al.*, 2012] Tyler Lu, Pingzhong Tang, Ariel D. Procaccia, and Craig Boutilier. Bayesian vote manipulation: Optimal strategies and impact on welfare. In *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2012.

[Mao *et al.*, 2013] Andrew Mao, Ariel D. Procaccia, and Yiling Chen. Better human computation through principled voting. In *Proceedings of the 27th Conference on Artificial Intelligence (AAAI)*, 2013.

[Moulin, 1980] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.

[Narayanan and Shmatikov, 2008] Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *Proceedings of the 29th IEEE Symposium on Security and Privacy (SP)*, pages 111–125, 2008.

[Oren *et al.*, 2013] Joel Oren, Yuval Filmus, and Craig Boutilier. Efficient vote elicitation under candidate uncertainty. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2013.

[Procaccia, 2010] Ariel Procaccia. Can approximation circumvent gibbard-satterthwaite? In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, pages 836–841, 2010.

[Satterthwaite, 1975] Mark Satterthwaite. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.

[Service and Adams, 2012] Travis C. Service and Julie A. Adams. Communication complexity of approximating voting rules. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 593–602, 2012.