Voting Power, Hierarchical Pivotal Sets, and Random Dictatorships

David M. Pennock
Yahoo! Research New York
111 West 40th Street, 17th floor
New York, NY 10018
pennockd@yahoo-inc.com

Lirong Xia
Department of Computer Science
Duke University
Durham, NC 27708, USA
lxia@cs.duke.edu

Abstract

In many traditional social choice problems, analyzing the voting power of the voters in a given profile is an important part. Usually the voting power of an agent is measured by whether the agent is *pivotal*. In this paper, we introduce two extensions of the set of pivotal agents to measure agents' voting power in a given profile. The first, which is called *hierarchical pivotal sets*, captures the voting power for an agent to make other agents pivotal. The second, which is called *coalitional pivotal sets*, is based on the fact that each agent is given a weight that is computed similarly to the *Shapley-Shubik power index*. We also introduce random dictatorships induced by the two types of pivotal sets to approximate full random dictatorships. We show that the random dictatorships induced by the hierarchical pivotal sets are *strategic-pivot-proof*, that is, no agent can make herself become one of the possible dictators by voting differently.

We then focus on the hierarchical pivotal sets when the hierarchical level goes to infinity. We prove that for any voting rule that satisfies *anonymity* and *unanimity*, and for any given profile, the union of the hierarchical pivotal sets are a sound and complete characterization of the non-redundant agents. We also show that if the voting rule does not satisfy anonymity, then this characterization might not be complete. Finally, we investigate algorithmic aspects of computing the hierarchical pivotal sets.

1 Introduction

Voting has been used in multiagent systems as a popular way to aggregate agents' preferences over a set of alternatives. Recently, a burgeoning field *computational social choice* was formed to study the computational aspects of voting. In computational social choice, one central problem is to investigate the possibility of using computational complexity as a barrier against manipulation. Researchers have been interested in the computational complexity of computing whether a single agent or a coalition of agents have enough power to replace the winner with their favorite alternative by casting votes strategically in collaboration. See [6] and [8] for nice recent surveys.

Looking back in the literature, the study of voting power has been favored in Political Science and Economics for a long time. It has been playing a central role in at least two other main research directions in addition to the study of manipulation. The first direction is the study of rational choice of voters, motivated by the "paradox of not voting", which dates back to Downs' seminal work [5]. The paradox states that when the number of voters is large, the voting power for a single voter to influence the outcome is negligible. Therefore, nobody should bother to vote, which sharply contradicts the much higher turnout in real-life elections. The paradox of not voting has influenced the study of voting in Political Science for more than half a century, and is still popular nowadays. Many research papers have been devoted to explaining the paradox from both theoretical and empirical sides, yet none of them has been successful so far. See [9; 10] for recent surveys.

The second research direction is the study of a class of coalitional games called *weighted voting games*. In a weighted voting game, each voter has a weight, and a coalition of voters is winning if the sum of their weights is higher than a quota (which is usually set to be half of the total weight). It is important to study the power of the voters for many purposes, e.g., for dividing the profit. One of the most important measurements is the *Shapley-Shubik power index* [13], where a voter's power is measured by (informally speaking) her marginal contribution in making coalitions of voters win.

In all the above research directions, a voter's voting power is determined by whether or not she is *pivotal*. That is, in a given profile, a voter is pivotal if and only if she can change the winner by casting a different vote, assuming that the other voters do not change their votes[1]. However, the mere "pivotal or not" measurement is often not discriminative enough. As the paradox of not voting says, the set of pivotal voters is always too small or even empty when the number of voters is large. This argument is supported by some recent work on the probability that a coalition of voters have power to change the outcome [12; 14].

Our conceptual contributions. In this paper, we introduce two new ways to measure a voter's power in a given profile for

[1] In the study of voting power, we do not consider the voter's incentive to cast a different vote.

a given voting rule. Both ways are extensions of the set of pivotal agents, and are much more discriminative. Therefore, we believe that these extensions provides new angles of the voters' strategic behavior in the three traditional research directions mentioned above. The first extension, which is called *hierarchical pivotal sets*, captures the power for a voter to make other voters pivotal. Given a profile, the level-1 hierarchical set is composed of all pivotal voters; for any $k \geq 2$, the level-k hierarchical set is composed of all voters who can change the level-$(k-1)$ hierarchical set by voting differently. The second extension is called *coalitional pivotal sets*. Such sets are subsets of voters who can change the winner by voting differently in collaboration. Based on the coalitional pivotal sets, we define power indices for the voters similarly to the Shapley-Shubik power index[2].

Our technical contributions. To illustrate the applications of these extensions, we define random dictatorships based on them to approximate the fully random dictatorship (which first chooses a voter uniformly at random, then select the winner to be the top-ranked alternative of the chosen voter). Fully random dictatorship is the only randomized voting rule that satisfies anonymity, Pareto-optimality, and strategy-proofness [11]. We prove that the random dictatorships based on hierarchical pivotal sets are *strategic-pivot-proof*, that is, no agent can make herself one of the possible dictators by voting differently.

Our main technical contribution is the following characterization of the hierarchical pivotal sets. We prove that for any voting rule that satisfies *anonymity* and *unanimity* and any profile, the voters in the hierarchical pivotal sets are not redundant (a voter is redundant if he/she is never pivotal in any profile). And conversely, any non-redundant voter must be in the level-k pivotal set for some $k \leq n+1$, where n is the number of voters. Therefore, in terms of hierarchical pivotal sets, for any anonymous voting rule, in any profile, any voter has some voting power to (directly or indirectly) change the winner. This provides a new perspective towards understanding the paradox of not voting. However, we also show that there exists a voting rule that does not satisfy anonymity, such that for any given profile, not all non-redundant voters are in the union of all hierarchical pivotal sets.

Finally, we investigate algorithmic aspects of computing the hierarchical pivotal sets.

2 Preliminaries

Let $\mathcal{C}$ be a finite set of *alternatives* (or *candidates*). A *vote* V is a linear order over $\mathcal{C}$, i.e., a transitive, antisymmetric, and total relation over $\mathcal{C}$. The set of all linear orders over $\mathcal{C}$ is denoted by $L(\mathcal{C})$. An n-voter profile P over $\mathcal{C}$ is a collection of n linear orders over $\mathcal{C}$, that is, $P = (V_1, \ldots, V_n)$, where for every $j \leq n$, $V_j \in L(\mathcal{C})$. In this paper, we let m denote the number of alternatives and let n denote the number of voters (agents) in a profile. Let $N = \{1, \ldots, n\}$. For any subset $S \subseteq N$, we let P_S denote the sub-profile of P that consists of the votes of the voters in S; let $P_{-S} = P_{N \setminus S}$. When $S = \{i\}$, we write P_{-i} instead of $P_{-\{i\}}$. The set of all n-profiles over $L(\mathcal{C})$ is denoted by $\mathrm{F}_n(\mathcal{C})$. In this paper, a *(voting) rule* r maps any n-profile to a single winning alternative, called the *winner*. Some commonly used voting rules are listed below.

- *Positional scoring rules.* Given a *scoring vector* $\vec{v} = (v_1, \ldots, v_m)$ of m integers, for any vote $V \in L(\mathcal{C})$ and any $c \in \mathcal{C}$, let $s_{\vec{v}}(V, c) = v_i$, where i is the rank of c in V. For any profile $P = (V_1, \ldots, V_n)$, let $s_{\vec{v}}(P, c) = \sum_{j=1}^{n} s_{\vec{v}}(V_j, c)$. The rule will select an alternative $c \in \mathcal{C}$ so that $s_{\vec{v}}(P, c)$ is maximized. Some examples of positional scoring rules are *plurality*, for which the scoring vector is $(1, 0, \ldots, 0)$, and *veto*, for which the scoring vector is $(1, \ldots, 1, 0)$. Plurality is also called *majority* when there are only two alternatives.
- *Single transferable vote (STV).* The election has m rounds. In each round, the alternative that gets the minimal plurality score drops out, and is removed from all of the votes. The last-remaining alternative is the winner.
- *Ranked pairs.* This rule first creates an entire ranking of all the alternatives. Let $D_P(c_i, c_j)$ denote the number of votes where $c_i \succ c_j$ minus the number of votes where $c_j \succ c_i$ in the profile P. In each step, we consider a pair of alternatives c_i, c_j that we have not previously considered, which has the highest $D_P(c_i, c_j)$ among the remaining pairs. We then fix the order $c_i \succ c_j$, unless it violates transitivity. We continue until all pairs of alternatives have been considered. The alternative at the top of the ranking wins.
- *Dictatorship.* For every $n \in \mathbb{N}$ there exists a voter $j \leq n$ such that the winner is always the alternative that is ranked in the top position in V_j. Voter j is called a *dictator*.

A voting rule r satisfies *anonymity*, if the winner under r does not depend on the name of the voters. That is, for any permutation M over N and any profile $P = (V_1, \ldots, V_n)$, we have $r(P) = r(M(P)) = r(V_{M(1)}, \ldots, V_{M(n)})$. r satisfies *unanimity*, if for any profile P in which all voters rank the same alternative c in their top positions, $r(P) = c$.

In this paper, we let a *random dictatorship* denote a mapping $D_r : \mathrm{F}_n(\mathcal{C}) \to 2^N$, where r is a "default" voting rule that is used to select the winner in case $D_r(P) = \emptyset$. That is, D_r selects a set of "possible dictators" to be randomized over. D_r naturally induces a mapping that assigns each profile to a probability distribution over $\mathcal{C}$ as follows. For any profile P, if $D_r(P) = \emptyset$, then it selects $r(P)$ with probability 1; if $D_r(P) \neq \emptyset$, then it first selects a voter j from $D_r(P)$ uniformly at random, then let the winner be the top-ranked alternative in V_j. A *fully random dictatorship* is a random dictatorship that always outputs N. A *weighted random dictatorship* D_r^w maps a profile to a probability distribution over N, or $\emptyset$. Similarly to random dictatorships, a weighted random dictatorship naturally induces a mapping that assigns each profile to a probability distribution over $\mathcal{C}$: if $D_r^w(P) = \pi \neq \emptyset$, then it selects a voter j from $D_r^w(P)$ according to the distribution π and let the winner to be the top-ranked alternative in V_j; and if $D_r^w(P) = \emptyset$, then it selects $r(P)$ with probability 1.

3 Pivotal sets and random dictatorships

In this section, we introduce two extensions of pivotal sets and their induced (weighted) random dictatorships, and discuss their relationships.

[2]The concept of coalitional pivotal sets is not new, for example, it is implicitly considered in the coalitional manipulation problems. However, as far as we know, this is the first time it is used to define voting power.

3.1 Hierarchical pivotal sets

Given a voting rule r and a profile P, we define the level-1 pivotal set $\text{PS}_r^1(P) \subseteq N$ to be the set of all pivotal voters. That is, $j \in \text{PS}_r^1(P)$ if and only if there exists a vote V_j' such that $r(P_{-j}, V_j') \neq r(P)$. Let the level-1 random dictatorship $D_r^1(P)$ be a mapping such that $D_r^1(P) = \text{PS}_r^1(P)$.

We argue that D_r^1 prevents voters' strategic behavior to some extent, by showing that any voter j who is not in $D_r^1(P)$ cannot make herself become a member in $D_r^1(P_{-j}, V_j')$ by casting a different vote V_j'. By definition, voter j is not pivotal. Therefore, for any pair of votes V_j' and V_j^*, $r(P_{-j}, V_j') = r(P_{-j}, V_j^*)$, which means that $j \notin D_r^1(P_{-j}, V_j')$. Formally, we have the following definition for random dictatorships.

Definition 1 *A random dictatorship D_r is* strategic-pivot-proof, *if for any profile P, any voter j, and any vote V_j', we have $j \in D_r(P_{-j}, V_j') \Longrightarrow j \in D_r(P)$.*

That is, D_r is strategic-pivot-proof if for any profile, any voter who is not selected by D_r cannot cast a different vote to make himself/herself one of the possible dictators. Of course for a strategic-pivot-proof random dictatorship, the voter might still have power and incentive to cast a different vote to change the set of possible dictators, even though she is not in it anyway. Therefore, it seems that strategic-pivot-proofness is weaker than the usual strategy-proofness. We note that they are actually not comparable. Exploring their relationship is an interesting direction for future research.

The level-1 pivotal set and its induced random dictatorship are not the end of the story. To capture the voting power for a voter to change the level-1 pivotal set, we can define level-2 pivotal sets to be composed of all voters who can change the level-1 pivotal set by voting differently. More generally, for any natural number k, we define the level-k pivotal set $\text{PS}_r^k(P) \subseteq N$ recursively as follows.

Definition 2 *For any voting rule r, any $k \in \mathbb{N}$, and any profile P, we define the level-k pivotal set $PS_r^k(P) \subseteq N$ recursively as follows.*

- *$j \in PS_r^1(P)$ if and only if there exists a vote V_j' such that $r(P) \neq r(P_{-j}, V_j')$.*
- *$j \in PS_r^k(P)$ if and only if there exists a vote V_j' such that $PS_r^{k-1}(P) \neq PS_r^{k-1}(P_{-j}, V_j')$. That is, voter j can change the level-$(k-1)$ pivotal set by voting differently.*

Here k is called the *hierarchical level.* Level-k pivotal sets capture voters' indirect power in the current profile P. The higher the hierarchical level is, the more indirectly the voters in it can influence the outcome for P. We note that the level-k pivotal sets for different profiles can be different.

Let D_r^k denote the random dictatorship such that $D_r^k(P) = \bigcup_{i=1}^k \text{PS}_r^i(P)$. In Section 4 we will show that for any voting rule r that satisfies anonymity and unanimity, D_r^k is an approximation to the fully random dictatorship after all redundant voters are removed. We note that the fully random dictatorship is strategy-proof.

Example 1 *There are two alternatives $\{a, b\}$, 5 voters, and we use the majority rule. Table 1 shows the level-k pivotal sets for all profiles, for $k = 1, 2, 3, 4$. Because the majority rule is anonymous, as we will show later in the paper (Lemma 1), the level-k pivotal set can be represented by a set of votes instead of a set of voters. A pivotal set is denoted by "b" if it is exactly the set of all voters whose votes are $b \succ a$; similarly for "a"; "all" denotes the set of all voters. For example, if two voters vote for $a \succ b$ and three voters vote for $b \succ a$, then the level-3 pivotal set consists of exactly the two voters whose votes are $a \succ b$.*

# of $a \succ b$	Pivotal sets				
	1	2	3	4	...
0	$\emptyset$	$\emptyset$	all	$\emptyset$	...
1	$\emptyset$	b	all	b	...
2	b	all	a	all	...
3	a	all	b	all	...
4	$\emptyset$	a	all	a	...
5	$\emptyset$	$\emptyset$	all	$\emptyset$	...

Table 1: The pivotal sets under majority.

Proposition 1 *For any $k \in \mathbb{N}$, D_r^k is strategic-pivot-proof.*

Proof: For any $j \notin \bigcup_{i=1}^k \text{PS}_r^i(P)$ and any vote V_j', we prove that for any $i \leq k$, $j \notin \text{PS}_r^i(P_{-j}, V_j')$. For the sake of contradiction, let $i \leq k$ and V_j' be such that $j \in \text{PS}_r^i(P_{-j}, V_j')$. By the definition of PS_r^i, there exists a vote V_j^* such that $\text{PS}_r^{i-1}(P_{-j}, V_j') \neq \text{PS}_r^{i-1}(P_{-j}, V_j^*)$. Therefore, either $\text{PS}_r^{i-1}(P_{-j}, V_j') \neq \text{PS}_r^{i-1}(P)$ or $\text{PS}_r^{i-1}(P_{-j}, V_j^*) \neq \text{PS}_r^{i-1}(P)$. In both cases $j \in \text{PS}_r^i(P)$, which contradicts the assumption. □

3.2 Coalitional pivotal sets and Shapley-Shubik power index

When defining hierarchical pivotal sets, we are concerned with the voting power for a single voter to (indirectly) change the winner. It is natural to consider the voting power for a coalition of voters to change the winner by voting collaboratively. We first define the set of pivotal coalitions.

Given a profile P, a subset $S \subset N$ is a *pivotal coalition*, if there exists a profile P_S' for the voters in S such that $r(P) \neq r(P_{-S}, P_S')$. We define the indicator function v_r^P as follows. For any coalition $S \subseteq N$, if S is a pivotal coalition, then $v_r^P(S) = 1$; otherwise $v_r^P(S) = 0$. For any voting rule r and any profile P, let $\text{CPS}_r(P)$ denote the set of all pivotal coalitions, that is, $\text{CPS}_r(P) = \{S \subseteq N : v_r^P(S) = 1\}$. Obviously, if a set of voters S can change the winner, then any superset of S can also change the winner. Therefore, for any r and any profile P, $\text{CPS}_r(P)$ is *upward-closed*, that is, for any $S \in \text{CPS}_r(P)$ and any S' such that $S \subseteq S'$, we have $S' \in \text{CPS}_r(P)$.

Example 2 *There are three alternatives $\{a, b, c\}$. Let P=($a \succ b \succ c$, $a \succ c \succ b$, $c \succ a \succ b$). We have $CPS_{\text{Plu}}(P) = \{\{1\}, \{2\}, \{1, 2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}$ and $CPS_{\text{Veto}}(P) = \{\{1\}, \{1, 2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}$.*

Now, a voter's voting power can be defined similarly to the Shapley-Shubik power index [13]. We now define a *power index w_r* that measures a voter's marginal contribution in making coalitions pivotal. Let $w_r : \text{F}_n(\mathcal{C}) \times N \rightarrow \mathbb{R}_{\geq 0}$ be a mapping such that for any profile $P \in \mathrm{F}_n(\mathcal{C})$ and any $j \leq n$, we have:

$$w_r(P, j) = \sum_{S \subseteq N \setminus \{j\}} \frac{|S|!(n - |S| - 1)!}{n!}(v_r^P(S \cup \{j\}) - v_r^P(S))$$

Proposition 2 *For any rule r that does not always select the same alternative and any profile P, $\sum_{j=1}^n w_r(P, j) = 1$.*

To the best of our knowledge, this is the first time that Shapley-Shubik power index is considered in the context of preference aggregation by voting rules.

Based on the power index w_r, we define a weighted random dictatorship D_r^w as follows. If $\mathrm{CPS}_r(P) = \emptyset$ (or equivalently, r always selects the same alternative), then $D_r^w(P) = \emptyset$. Otherwise, for any profile P, $D_r^w(P)$ is the distribution over N that chooses j with probability $w_r(P, j)$.

Example 3 *Let P be the same profile as defined in Example 2. $D_{\mathrm{Plu}}^w(P)$ chooses 1 and 2 with the same probability $1/2$; $D_{\mathrm{Veto}}^w(P)$ chooses 1 with probability $2/3$, and chooses 2 and 3 with the same probability $1/6$.*

3.3 Relationships between the two pivotal sets

The next theorem states that the smallest k such that the level-k pivotal set is non-empty equals to the size of the smallest coalitional pivotal set for P.

Theorem 1 *For any voting rule r and any profile P,*
$$\min\{k : PS_r^k(P) \neq \emptyset\} = \min_{S \in CPS_r(P)}\{|S|\}$$

Proof: Let $k^* = \arg\min_k\{\mathrm{PS}_r^k(P)\}$ and $k' = \min_{S \in \mathrm{CPS}_r(P)}\{|S|\}$. We first prove that $k^* \leq k'$. Suppose for the sake of contradiction that $k^* > k'$. Without loss of generality, $S = \{1, \ldots, k'\}$, and let $P'_S = (V'_1, \ldots, V'_{k'})$ be the votes such that $r(P) \neq r(P_{-S}, P'_S)$. For any $k \leq k'$, let $P_k = (V'_1, \ldots, V'_k, V_{k+1}, \ldots, V_n)$, that is, P_k is obtained from P by replacing the first k votes by $V'_1, \ldots, V'_k$, respectively. Because $k' < k^*$, for any $k \leq k'$, $\mathrm{PS}_r^k(P) = \emptyset$. Therefore, for any $k \leq k' - 1$, changing the vote of voter 1 from V_1 to V'_1 does not change the level-k pivotal set. That is, for any $k \leq k' - 1$, $\mathrm{PS}_r^k(P_1) = \emptyset$. Similarly, it is easy to see that for any $i \leq k' - 1$, for any $k \leq k' - i$, $\mathrm{PS}_r^k(P_i) = \emptyset$. Specifically, $\mathrm{PS}_r^1(P_{k'-1}) = \emptyset$. It follows from $\mathrm{PS}_r^1(P) = \emptyset$ and for any $i \leq k' - 1$, $\mathrm{PS}_r^1(P_i) = \emptyset$, that $r(P) = r(P_1) = r(P_2) = \ldots = r(P_{k'})$. This contradicts the assumption that $r(P) \neq r(P_{k'})$. Consequently, $k^* \leq k'$.

Next, we prove that $k' \leq k^*$. It suffices to prove that for any $k \leq k' - 1$, $\mathrm{PS}_r^k(P) = \emptyset$. We have following stronger claim, whose proof is omitted due to the space constraint.

Claim 1 *For any $2 \leq q \leq k'$, any P' that differs from P on no more than $k' - q$ votes, and any $k \leq q - 1$, $PS_r^k(P') = \emptyset$.*

Let $q = k'$ in Claim 1, we have that $\mathrm{PS}_r^{k'-1}(P) = \emptyset$, which means that $k^* \geq k'$. Therefore, $k^* = k'$. □

4 Hierarchical pivotal sets for anonymous voting rules

In the remainder of the paper, we focus on hierarchical pivotal sets. It is easy to see that if a voter is not pivotal in *any* profile, then for any k and any profile P, she is not in the level-k pivotal set. Such a voter is said to be *redundant*.

Definition 3 *Given a voting rule r, a voter j is* redundant*, if for any profile P and any vote V'_j, $r(P) = r(P_{-j}, V'_j)$.*

If a voter is redundant, then effectively her vote can be completely ignored. Therefore, for any profile, none of the voters in the union of its hierarchical pivotal sets (as $k \to \infty$) is redundant. That is, the union of the hierarchical pivotal sets for any profile is a *sound* characterization of the non-redundant voters. We ask the following two natural questions. The first question asks whether or not the union of the hierarchical pivotal sets for a given profile P is a *complete* characterization of the non-redundant voters.

Question 1 *Given a voting rule r, is it true that for any non-redundant voter j and any profile P, there exists $k \in \mathbb{N}$ such that j is in the level-k pivotal set for P?*

The second question concerns the asymptotic property of level-k pivotal sets when k goes to infinity. Given a profile P, we are asked whether the level-k pivotal sets for P will converge (to the empty set), when k goes to infinity.

Question 2 *Given a voting rule r, does there exist $K \in \mathbb{N}$ such that for any $k \geq K$, the level k-pivotal set is $\emptyset$?*

In this section, we give an affirmative answer to Question 1 for any voting rule that satisfies anonymity and unanimity, and a negative answer to Question 2 for the majority rule. We first prove a lemma, which states that for any anonymous voting rule r, if a voter j is in the level-k pivotal set for a profile P, then other voters who cast the same vote as j's vote are also in the level-k pivotal set for P. This lemma will be frequently used in this paper. Due to the space constraint, some proofs are omitted.

Lemma 1 *For any anonymous voting rule r, any profile P, any $k \in \mathbb{N}$, and any pair of voters i, j with $V_i = V_j$, $i \in PS_r^k(P)$ if and only if $j \in PS_r^k(P)$.*

Lemma 1 states that for any anonymous voting rule r and any profile P, a voter's membership in the level-k pivotal set can be characterized by her vote. Therefore, for any anonymous voting rule r and any profile, the level-k pivotal set can be represented by the set of all votes that are cast by some level-k pivotal voters. We will use this observation later in the paper, especially in Section 6. The next theorem gives an affirmative answer to Question 1 for any voting rule that satisfies anonymity and unanimity.

Theorem 2 *Let r be a voting rule that satisfies anonymity and unanimity. For any n-profile P and any voter j, there exists $k \leq \min_{S \in CPS_r(P)}\{|S|\} + 1 \leq n + 1$ such that $j \in PS_r^k(P)$.*

Proof: Let $K = \min_{S \in \mathrm{CPS}_r(P)}\{|S|\}$. For the sake of contradiction, without loss of generality for any $k \leq K + 1$, $1 \notin \mathrm{PS}_r^k(P)$. By Theorem 1, there exists $k^* \leq K$ such that $\mathrm{PS}_r^{k^*}(P) \neq \emptyset$. Let $j^* \in \mathrm{PS}_r^{k^*}(P)$ and W be the vote of voter j^*. Let $P' = (P_{-1}, W)$, that is, P' is the profile obtained from P by letting voter 1 vote for W. Because $1 \notin \mathrm{PS}_r^{k^*}(P)$ and $1 \notin \mathrm{PS}_r^{k^*+1}(P)$, we have that $1 \notin \mathrm{PS}_r^{k^*}(P')$. It follows from Lemma 1 that for any voter j whose vote is W in P', $j \notin \mathrm{PS}_r^{k^*}(P')$. Specifically, $j^* \notin \mathrm{PS}_r^{k^*}(P')$, which means that $\mathrm{PS}_r^{k^*}(P') \neq \mathrm{PS}_r^{k^*}(P)$. Therefore, $1 \in \mathrm{PS}_r^{k^*+1}(P)$. This contracts the assumption that $1 \notin \mathrm{PS}_r^{k^*+1}(P)$. □

Theorem 2 is quite positive. It implies that if we remove all redundant voters, D_r^k can be used to approximate the fully random dictatorship, which is strategy-proof. It is a very interesting topic to study how good this approximation is, which we left as an open problem.

For Question 2, suppose the level-k pivotal set converges as k goes to infinity, we first prove that it must converge to $\emptyset$.

Proposition 3 *For any anonymous voting rule r, if there exists k such that for any n-profile P, $PS_r^k(P) = PS_r^{k+1}(P)$, then for any n-profile P, $PS_r^k(P) = \emptyset$.*

However, Proposition 3 does not guarantee the existence of k such that $\text{PS}_r^k(P) = \text{PS}_r^{k+1}(P)$. In fact, the next proposition shows that such a k might not exist for the majority rule, which satisfies anonymity and unanimity. Therefore, the answer to Question 2 is negative.

Proposition 4 *Let there be two alternatives $\{a, b\}$, 5 voters, and we use the majority rule. There does not exist $k \in \mathbb{N}$ such that for any profile P, the level-k pivotal set for P is $\emptyset$.*

Proof: From Table 1 in Example 1, it is easy to see that for any profile, its level-2 and level-4 pivotal sets are identical and are different from level-3 pivotal sets. Therefore, for any profile, none of the level-k pivotal sets converges as k goes to infinity. □

5 Hierarchical pivotal sets for non-anonymous voting rules

In this section, we focus on non-anonymous voting rules. Surprisingly, for some voting rules that do not satisfy anonymity, the answer to Question 1 is negative.

Proposition 5 *Let $m = 4$ and $n = 3$. There exists a non-anonymous voting rule r that satisfies the following conditions.*

- *No voter is redundant.*
- *For any $k \in \mathbb{N}$ and any profile P such that $|P| = 3$, the level-k pivotal set for P is non-empty.*
- *For any voter j, there exists a profile P such that $|P| = 3$ and for any $k \in \mathbb{N}$, j is not in the level-k pivotal set for P.*

Proof: Let the four alternatives be $\{a, b, c, d\}$. Let $l = [a \succ b \succ c \succ d]$. We define a voting rule r as follows. $r(l, l, \neg) = r(\neg, l, \neg) = a$, $r(l, \neg, l) = r(l, \neg, \neg) = b$, $r(\neg, l, l) = r(\neg, \neg, l) = c$, $r(l, l, l) = r(\neg, \neg, \neg) = d$.

Here "$\neg$" means any linear order that is different from l. For example, $r(\neg, l, l) = c$ means that for any 3-profile where voter 1's voter is not l, and the votes of voter 2 voter and 3 are both l, the winner is c. The voting rule is illustrated in Figure 1(a), where each vertex represents a set of 3-profiles and the alternative associated with it is the winner for these profiles. An edge between two vertices A and B in the graph means that for any profile P in A, there exists a profile P' in B such that P' can be obtained from P by changing exactly one vote. An edge is bold if the winners for its two endpoints are the same. We have the following claim (whose proof is omitted due to the space constraint.)

Claim 2 *For any $k \in \mathbb{N}$ and any profile P, $PS_r^k(P) = PS_r^{k+1}(P)$, and is illustrated in Figure 1 (b).*

It follows from Claim 2 that r satisfies all the properties in the description of the proposition. □

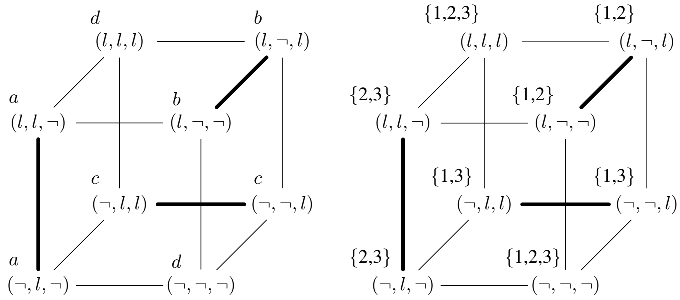

(a) The voting rule r. (b) The level-k pivotal set for any k.

Figure 1: The voting rule r and the hierarchical pivotal sets.

6 Computing hierarchical pivotal sets

In this section, we investigate the computational complexity of computing level-k pivotal sets. We first relate the problem of computing level-1 pivotal sets to the *unweighted coalitional manipulation (UCM)* problems with a single manipulator. An instance of UCM is a tuple (r, P^{NM}, c, M), where r is a voting rule, P^{NM} is the non-manipulators' profile, c is the manipulators' preferable alternative, and M is the set of manipulators. We are asked whether there exists a profile P^M for the manipulators such that $r(P^{NM} \cup P^M) = c$. Let UCM_1 denote the UCM problems with a single manipulator, that is, $|M| = 1$.

Proposition 6 *For any voting rule r, if UCM_1 is in P, then computing $PS_r^1(P)$ is also in P.*

Following the results of computing UCM_1 for common voting rules [3; 2; 4; 7; 16; 15], we immediately obtain the following corollary.

Corollary 1 *For any $r \in$ {Copeland, Veto, Plurality with runoff, Cup, Maximin, Bucklin, Borda} and any profile P, there exists a polynomial-time algorithm that computes $PS_r^1(P)$.*[3]

For STV and ranked pairs, UCM_1 is NP-complete [2; 15]. The next two theorems show that computing the level-1 pivotal sets for them are NP-complete. It is not clear whether there exists a general reduction that works for any voting rule.

Theorem 3 *It is NP-complete to compute $PS_r^1(P)$ for r=STV.*

Proof: It is easy to check that computing $\text{PS}_r^1(P)$ for STV is in NP. We prove the NP-hardness by a reduction from a special kind of UCM_1 problems for STV, where c is ranked in the top position in at least one vote in P^{NM}. This problem has been shown to be NP-complete [2]. For any UCM_1 instance $(\text{STV}, P^{NM}, c, \{n\})$ where c is ranked in the top position in at least one vote in P^{NM} ($|P^{NM}| = n - 1$), we construct the following instance of computing the level-1 pivotal set. Let $\mathcal{C}$ denote the set of alternatives in the UCM_1 instance.

Alternatives: $\mathcal{C} \cup \{d\}$, where d is an auxiliary alternative.

Profile: Let P denote a profile of $2n - 1$ votes as follows. The first $n-1$ votes are obtained from P^{NM} by putting d right below c. The next n votes ranks d in the first position (other alternatives are ranked arbitrarily). We are asked whether $n \in \text{PS}_{\text{STV}}^1(P)$.

[3]The definition of these voting rules can be found in e.g. [15].

It is easy to check that $\text{STV}(P) = d$. Suppose the UCM_1 instance has a solution, denoted by V. Then, let V' denote the linear order over $\mathcal{C} \cup \{d\}$ obtained from V by ranking d in the bottom position. Let P' denote the profile where voter n changes her vote to V'. We note that d is ranked in the top position for $n-1$ time in P'. Therefore, d is never eliminated in the first $|\mathcal{C}| - 1$ rounds. Moreover, for any $j \leq |\mathcal{C}| - 1$, the alternative that is eliminated in the jth round for P' is exactly the same as the alternative that is eliminated in the jth round for P. In the lasts round, c is ranked in the top position for n time, which means that $\text{STV}(P') = c \neq d$. Hence, $n \in \text{PS}^1_{\text{STV}}(P)$.

On the other hand, if $n \in \text{PS}^1_{\text{STV}}(P)$, then there exist a vote V' such that by changing her vote to V', voter n can change the winner under STV. Let $P' = (P_{-n}, V')$. Again, because d is ranked in the top position for at least $n-1$ time in P', it will only be eliminated in the last round. We recall that c is ranked in the first position in at least one vote in P^{NM}, and d is ranked right below c in the corresponding vote in P'. Therefore, d beats all alternatives in $\mathcal{C} \setminus \{c\}$ in their pairwise elections, which means that in the last round the only remaining alternatives must be c and d. Let V be a linear order obtained from V' by removing d. It follows that V is a solution to the UCM_1 instance.

Therefore, computing the level-1 pivotal set for STV is NP-complete. $\square$

Theorem 4 (proof omitted due to the space constraint) *It is NP-complete to compute $PS^1_r(P)$ for r=RP (ranked pairs).*

For any anonymous voting rule, when m is bounded above by a constant, we can find a dynamic-programming algorithm that computes the level-k pivotal set. The algorithm is based on the following two key observations. First, when the number of alternatives is bounded above by a constant, the number of essentially different profiles is polynomial. Second, by Lemma 1, a level-k pivotal set can be represented succinctly by a set of votes (instead of voters). The details of the algorithm is omitted due to the space constraint.

7 Future research

There are many interesting directions for future research. For example, in this paper we have three open problems. How can we compare the strategic-pivot-proofness and strategy-proofness? How good/bad it is to use D^k_r to approximate the fully random dictatorship? What is the computational complexity of computing level-k ($k \geq 2$) pivotal sets for common voting rules? Moreover, we believe that defining and computing voting power in the traditional voting setting (in contrast to the weighted voting games) is an important topic. It would be worthwhile studying applications of the two types of voting powers proposed in this paper (especially the Shapley-Shubik power index), for example, in defining other (weighted) random dictatorships or in the coalition formation of the manipulators. Besides these topics, we can definitely examine other ways of defining voting power, for example by using the Banzhaf power index [1].

Acknowledgements

Lirong Xia acknowledges a James B. Duke Fellowship and Vincent Conitzer's NSF CAREER 0953756 and IIS-0812113, and an Alfred P. Sloan fellowship for support. We thank all anonymous IJCAI-11 and WSCAI reviewers for their helpful suggestions and comments.

References

[1] John F. Banzhaf. Weighted voting doesn't work: A mathematical analysis. *Rutgers Law Review*, 19(2):317–343, 1965.

[2] John Bartholdi, III and James Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.

[3] John Bartholdi, III, Craig Tovey, and Michael Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[4] Vincent Conitzer, Tuomas Sandholm, and Jérôme Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3): 1–33, 2007.

[5] Anthony Downs. *An Economic Theory of Democracy*. New York: Harper & Row, 1957.

[6] Piotr Faliszewski, Edith Hemaspaandra, and Lane A. Hemaspaandra. Using complexity to protect elections. *Commun. ACM*, 53:74–82, 2010.

[7] Piotr Faliszewski, Edith Hemaspaandra, and Henning Schnoor. Copeland voting: ties matter. In *Proc. of AAMAS*, pages 983–990, 2008.

[8] Piotr Faliszewski and Ariel D. Procaccia. AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.

[9] Timothy J. Feddersen. Rational choice theory and the paradox of not voting. *JEP*, 18(1):99–112, 2004.

[10] Benny Geys. 'Rational' theories of voter turnout: A review. *Political Studies Review*, 4(1):16–35, 2006.

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

[12] Ariel D. Procaccia and Jeffrey S. Rosenschein. Average-case tractability of manipulation in voting via the fraction of manipulators. In *Proc. AAMAS*, pages 718–720, 2007.

[13] Lloyd Shapley and Martin Shubik. A method for evaluating the distribution of power in a committee system. *The American Political Science Review*, 48(3):787–792, 1954.

[14] Lirong Xia and Vincent Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proc. EC*, pages 109–118, 2008.

[15] Lirong Xia, Michael Zuckerman, Ariel D. Procaccia, Vincent Conitzer, and Jeffrey Rosenschein. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proc. IJCAI*, pages 348–353, 2009.

[16] Michael Zuckerman, Ariel D. Procaccia, and Jeffrey S. Rosenschein. Algorithms for the coalitional manipulation problem. *Artificial Intelligence*, 173(2):392–412, 2009.