Manipulating the Quota in Weighted Voting Games

Michael Zuckerman
School of Computer Science and Engineering, The Hebrew University of Jerusalem, Jerusalem 91904, Israel
*

Piotr Faliszewski
Department of Computer Science University of Rochester Rochester, NY 14627 USA
†

Yoram Bachrach
School of Computer Science and Engineering, The Hebrew University of Jerusalem, Jerusalem 91904, Israel
‡

Edith Elkind
School of Electronics and Computer Science University of Southampton, Southampton SO17 1BJ, United Kingdom
§

Abstract

Weighted voting games provide a popular model of decision making in multiagent systems. Such games are described by a set of players, a list of players' weights, and a quota; a coalition of the players is said to be winning if the total weight of its members meets or exceeds the quota. The power of a player in such games is traditionally identified with her Shapley–Shubik index or her Banzhaf index, two classical power measures that reflect the player's marginal contributions under different coalition formation scenarios. In this paper, we investigate by how much the central authority can change a player's power, as measured by these indices, by modifying the quota. We provide tight upper and lower bounds on the changes in the individual player's power that can result from a change in quota. We also study how the choice of quota can affect the relative power of the players. From the algorithmic perspective, we provide an efficient algorithm for determining whether there is a value of the quota that makes a given player a *dummy*, i.e., reduces his power (as measured by both indices) to 0. On the other hand, we show that checking which of the two values of the quota makes this player more powerful is computationally hard, namely, complete for the complexity class PP, which is believed to be significantly more powerful than NP.

Introduction

Cooperation and joint decision-making are key aspects of many interactions among self-interested agents. The collaborating agents may have different preferences, so they need a method to agree on a common course of action. One possible solution to this problem is to use a (weighted) voting procedure. Under such a procedure, each agent is assigned a numerical weight, and a coalition is deemed to be winning if its total weight exceeds a given quota.

An important issue in weighted voting is how to measure the *power* of each voter, i.e., her ability to affect the final outcome. In particular, this question is critical when the agents have to decide how to distribute the payoffs resulting from their joint action: A natural approach is to pay each agent according to his contribution, i.e., his voting power.

An agent's ability to affect the result of the game is not always directly proportional to her weight. For example, in a game where the quota is so high that the only winning coalition is the one that involves all agents, each agent can veto the decision, and hence all agents have equal power. Thus, to measure the power, instead of using agents' weights, one typically employs one of the so-called power index functions. Perhaps the most prominent ones are the Shapley–Shubik index (Shapley and Shubik 1954) and the Banzhaf index (Dubey and Shapley 1979; Banzhaf 1965). Intuitively, they both measure the probability that a given agent is critical to a forming coalition, i.e., that the coalition would become winning if the agent joined in. The difference between these two power indices comes from different coalition formation models.

The value of an agent's power index reflects his ability to affect the outcome and may determine his payoffs. Therefore, selfish agents may try to increase their power, as measured by these power indices, by employing some form of manipulative behavior, such as, e.g., splitting their weight between several identities; this form of manipulation was recently studied in (Bachrach and Elkind 2008). Similarly, the central authority, may want to minimize or maximize the influence of a particular agent by modifying the rules of the game, e.g., by changing the quota. The goal of this paper is to study the effects on the agents' power caused by a malicious central authority.[1] Plausible goals for the center include maximizing or minimizing a given player's power-index value (in particular, making a given player a *dummy*, i.e., reducing her power to 0), or ensuring that all players have different power-index values (or, on a more local scale, ensuring that two given players have either different or equal power-index values). In this paper, we study these issues from both the worst-case and the algorithmic perspective. We give matching upper and lower bounds on the worst-case relative and absolute effects that a change of the quota may have on a given player's power. As in several applications the ranking of the agents is more important than the exact

*email: michez@cs.huji.ac.il.
†email: pfali@cs.rochester.edu. Supported in part by grant NSF-CCF-0426761.
‡email: yorambac@gmail.com.
§email: eelkind@gmail.com.

[1] In voting theory literature, dishonest behavior by the central authority is usually referred to as "control", while the term "manipulation" is reserved for voters' dishonest behavior. However, in this paper we will use both terms interchangeably.

power they possess, we also study the problem of setting the quota so as to guarantee a particular relation (equality or inequality) between two players' power-index values. A related issue that we consider is that of selecting the quota value to ensure that all players with different weights have different power-index values. Finally, we investigate the quota manipulation problem from computational perspective. We describe a polynomial-time algorithm for testing whether there is a quota value that makes a given player a dummy, and we show that the problem of deciding which of the two quotas is better for a particular player is complete for the complexity class PP, which is believed to be more powerful than NP.

Related work The Shapley value originated in a seminal paper (Shapley 1953) which considered how to fairly allocate the utility gained by the grand coalition in cooperative games. A subsequent paper (Shapley and Shubik 1954) applies the Shapley value to weighted voting games, so this value is referred to as the *Shapley–Shubik power index* in this context. A first version of the Banzhaf power index was introduced in (Banzhaf 1965); a more natural definition was later proposed in (Dubey and Shapley 1979). Both power indices have been well studied (Roth 1988). Their practical applications include analyzing the voting structures of the European Union Council of Ministers and the IMF (Machover and Felsenthal 2001; Leech 2002). Computational complexity of power indices is also quite well understood: while computing both indices is #P-complete (Garey and Johnson 1979; Prasad and Kelly 1990), they can be computed in polynomial time when all weights are at most polynomial in the number of players (Matsui and Matsui 2000), and several papers (e.g., (Fatima, Wooldridge, and Jennings 2007; Bachrach et al. 2008)) discuss ways to *approximate* them. Some of these algorithms work well in practice and thus justify the use of power indices as payoff distribution schemes.

Computational aspects of various forms of dishonest behavior in voting with m alternatives received a lot of attention in recent years (Endriss and Lang 2006). Specifically, this research considers *manipulation* (dishonest behavior by voters), *control* (dishonest behavior by the election authority), and *bribery* (dishonest behavior by an outside party). This stream of work, and, in particular, the papers devoted to control, provides motivation for our research, but results for the model with m alternatives cannot be directly applied to our setting. Several papers deal with manipulations aimed at increasing the Shapley value of an agent in various domains (Conitzer and Sandholm 2004; Yokoo et al. 2005). Perhaps the closest in spirit to our work is (Bachrach and Elkind 2008), which considers manipulation by *voters* in weighted voting games. However, to the best of our knowledge, manipulation by the *center* in the context of weighted voting games has not been studied before.

Preliminaries and Notation

Weighted Voting Games A *weighted voting game* $G = [I; \mathbf{w}; q]$ is given by a set of players $I = \{1, \ldots, n\}$, a vector of players' *weights* $\mathbf{w} = (w_1, \ldots, w_n)$ and a *quota* q. A *coalition* is a subset of players $J \subseteq I$. A coalition J is *winning* if its total weight meets or exceeds the quota, i.e., $\sum_{j \in J} w_j \geq q$ and is *losing* otherwise.

We write $v(J) = 1$ if J wins and $v(J) = 0$ if J loses. We say that an agent $i \in J$ is *pivotal* to coalition J if $v(J) = 1$ and $v(J \setminus \{i\}) = 0$; similarly, i *contributes* to J if $v(J) = 0$, $v(J \cup \{i\}) = 1$. A player i is called a *dummy* if he does not contribute to any coalition, i.e., for any $J \subseteq I$ we have $v(J \cup \{i\}) = v(J)$. We denote by $w(J)$ the total weight of a coalition J, i.e., $w(J) = \sum_{i \in J} w_i$. For the purposes of this paper, we can assume without loss of generality that $0 < w_1 \leq \cdots \leq w_n$ and that $0 < q \leq w(I)$. Therefore, we will make these assumptions throughout the paper, unless explicitly specified otherwise.

Shapley–Shubik Index and Banzhaf Index Both Shapley–Shubik index and Banzhaf index measure an agent's marginal contribution to possible coalitions. However, they differ in the underlying coalition formation scenarios: while the Shapley–Shubik index implicitly assumes that the agents join a coalition in random order, the Banzhaf index is based on the assumption that each agent decides whether to join a coalition independently at random. Both of these measures can be defined for a much larger class of games than weighted voting games. However, in what follows we provide definitions that are specialized to our scenario.

Let Π be the set of all one-to-one mapping from I to I; an element of Π is denoted by π. Set $S_\pi(i) = \{j \mid \pi(j) < \pi(i)\}$: the set $S_\pi(i)$ consists of all predecessors of i in π. The *Shapley–Shubik index* of the ith agent in a game $G = [I; \mathbf{w}; q]$ is denoted by $\varphi_i(G)$ and is given by the following expression:

$$\varphi_i(G) = \frac{1}{n!} \sum_{\pi \in \Pi} [v(S_\pi(i) \cup \{i\}) - v(S_\pi(i))]. \qquad (1)$$

In words, the Shapley-Shubik power index counts the fraction of all orderings of the agents in which agent i is pivotal for the coalition formed by his predecessors and himself. We will occasionally abuse notation and say that an agent i is pivotal for a permutation π if it is pivotal for the coalition $S_\pi(i) \cup \{i\}$.

The Banzhaf index $\beta_i(G)$ of an agent i in a game $G = [I; \mathbf{w}; q]$ is computed as follows:

$$\beta_i(G) = \frac{1}{2^{n-1}} \sum_{S: i \notin S} [v(S \cup \{i\}) - v(S)]. \qquad (2)$$

This index simply counts the number of coalitions for which agent i is pivotal.

Both of these indices have several useful properties that make them very convenient to work with. In particular, both of them have the *dummy player* property, which states that the value of the index for a given player is 0 if and only if he does not contribute to any coalition, and the *symmetry* property, which states that if two players have equal weights, then their indices are equal. Also, Shapley–Shubik index (but not the Banzhaf index) has the *normalization property*, which claims that the sum of Shapley–Shubik indices of all players is equal to 1. All of these properties are easy to verify from the definitions.

To simplify notation, given a game $G = [I; \mathbf{w}; q]$, we will sometimes write $\varphi_i(q)$ and $\beta_i(q)$ instead of $\varphi_i(G)$ and $\beta_i(G)$ if I and $\mathbf{w}$ are clear from the context.

Upper and Lower Bounds for a Single Player

We will start this section by showing that the center can significantly change the players' Shapley–Shubik and Banzhaf index by manipulating the quota. We then proceed to quantify the worst case effects of this manipulation for all players. We will be interested both in the *ratios* of the player's powers for a given pair of quotas and in their *differences*.

Example 1. Consider a weighted voting game $G = [I; (1, 2, 3); 3]$. In this game, the player 3 is pivotal to three coalitions (namely, $\{3\}$, $\{1, 3\}$ and $\{2, 3\}$) and to four permutations (namely, 312, 321, 132 and 231), so we have $\beta_3(G) = 3/4$, $\varphi_3(G) = 2/3$. Now change the quota to 1. in the resulting game G, player 3 is only pivotal if it joins an empty coalition or appears first in a permutation, so we have $\beta_3(G) = 1/4$, $\varphi_3(G) = 1/3$.

A natural bound on manipulator's influence is the worst-case ratio between a given player's values of the index in the two games corresponding to two different values of the quota. Unfortunately, as we will now show, this ratio can only be bounded for the largest player; for all other players, it might be possible to turn them into dummies. Hence, at least in some weighted voting games, the center can change the agents' power by more than a constant factor.

Theorem 2. *Given a set of players I, $|I| = n$, there exists a weight vector $\mathbf{w}$, $0 < w_1 \le \dots \le w_n$ and $q, q' \le w(I)$ such that for $i = 1, \dots, n-1$, we have $\varphi_i(q') = \beta_i(q') = 0$, while $\varphi_i(q) \ne 0$, $\beta_i(q) \ne 0$. On the other hand, for any $\mathbf{w}$ such that $0 < w_1 \le \dots \le w_n$ and any $q, q' \le w(I)$, we have $\varphi_n(q)/\varphi_n(q') \le n$, $\beta_n(q)/\beta_n(q') \le 2^{n-1}$, and these bounds are tight.*

Proof. Set $\mathbf{w} = (1, \dots, 1, n)$. In the game $G = [I; \mathbf{w}; 1]$ all players have equal power, so by symmetry we have $\varphi_i(1) = 1/n$ for $i = 1, \dots, n$. Moreover, each player is pivotal for exactly one coalition, so we have $\beta_i(1) = 1/2^{n-1}$. On the other hand, in the game $G' = [I; \mathbf{w}; n]$, all the players except for the last one are dummies, so their Shapley–Shubik and Banzhaf indices are 0, and we have $\varphi_n(n) = 1$, $\beta_n(n) = 1$. Hence, $\varphi_n(n)/\varphi_n(1) = n$, $\beta_n(n)/\beta_n(1) = 2^{n-1}$.

To show that the ratio $\varphi_n(q')/\varphi_n(q)$ cannot exceed n, it is enough to note that for any n-player weighted voting game G it holds that $1/n \le \varphi_n(G) \le 1$, where both inequalities follow from the fact that for any i, $1 \le i < n$, $0 \le \varphi_i(G) \le \varphi_n(G)$ and $\sum_{k=1}^{n} \varphi_k(G) = 1$. Similarly, in any weighted voting game G we have $1/2^{n-1} \le \beta_n(G) \le 1$, so the ratio $\beta_n(q')/\beta_n(q)$ cannot exceed 2^{n-1}. □

By considering the weight vector $\mathbf{w} = (1, 2, 4, \dots, 2^{n-1})$ and quotas $q = 2^{k-1} - 1$, $k = 2, \dots, n-1$, we can show that the ratios $\varphi_i(q')/\varphi_i(q)$ cannot be bounded by a constant even if it is required that $\varphi_i(q) \ne 0$; we omit the details.

Since the previous approach yielded no meaningful bounds for the first $n-1$ players, we will now try to bound the worst-case *difference* between a given player's values in the corresponding games. We obtain tight bounds for this problem.

Theorem 3. *For a set of players I, $|I| = n$, any weight vector $\mathbf{w}$, $0 < w_1 \le \dots \le w_n$ and any $q, q' \le w(I)$, for $i = 1, \dots, n-1$ the difference $\varphi_i(q) - \varphi_i(q')$ can be at most $1/(n-i+1)$ and this bound is tight. For player n, the difference $\varphi_n(q) - \varphi_n(q')$ can be at most $1 - 1/n$, and this bound is tight.*

Proof. Set $\mathbf{w} = (1, 2, 4, \dots, 2^{n-1})$. In the game $[I; \mathbf{w}; 2^k]$, where $k \in \{1, \dots, n-1\}$, the first k players are dummies, and the last $n-k$ players have equal power, $1/(n-k)$. Hence, for $i = 1, \dots, n-1$, by changing the quota from 2^i to 2^{i-1}, we change the Shapley–Shubik index of the ith player from 0 to $1/(n-i+1)$, as required. To see that this bound is tight, consider an arbitrary weight vector $\mathbf{w}'$ that satisfies $0 < w'_1 \le \dots \le w'_n$, a player i, $1 \le i < n$, and a quota $q' \le w'(I)$. Naturally, $\varphi_i(I; \mathbf{w}'; q) \ge 0$ and monotonicity of the Shapley–Shubik index implies that for $j > i$ we have $\varphi_i(I; \mathbf{w}'; q) \le \varphi_j(I; \mathbf{w}'; q)$. As $\sum_{k=i}^{n} \varphi_k(I; \mathbf{w}'; q) \le 1$, we have $\varphi_i(I; \mathbf{w}'; q) \le 1/(n-i+1)$.

For player n and our weight vector $\mathbf{w}$, changing the quota from 2^{n-1} to 1 changes n's Shapley–Shubik index from 1 to $1/n$, yielding the difference $1 - \frac{1}{n}$. Since in any n-player voting game G we have $\frac{1}{n} \le \varphi_n(G) \le 1$, this gives a tight bound for player n. □

Theorem 4. *For a set of players I, $|I| = n$, any weight vector $\mathbf{w}$, $0 < w_1 \le \dots \le w_n$ and any $q, q' \le w(I)$, for $i = 1, \dots, n-1$ the difference $\beta_i(q) - \beta_i(q')$ can be at most $\binom{n-i}{\lfloor \frac{n-i}{2} \rfloor} \cdot 2^{i-n}$ and this bound is tight. For player n, we have $\beta_n(q) - \beta_n(q') \le 1 - 1/2^{n-1}$ and this bound is tight.*

The proof of this statement is considerably more technical than that of the previous results, and is omitted due to lack of space.

Separating the Players

In the previous section we focused on the effects that a change of quota can have on the value of the index for a single player, both in absolute and in relative terms. These results are important in the situation where we are interested in the power of that player, irrespective of the effects it may have on other players. Another motivation for changing the quota could be affecting the relative power of two players i and j. For instance, suppose that $w_i < w_j$, and the center prefers player i to player j. From the monotonicity properties of both Shapley–Shubik index and Banzhaf index, it follows that for any value of the quota q both $\varphi_i(q) \le \varphi_j(q)$ and $\beta_i(q) \le \beta_j(q)$. Hence, the best that the center may hope for is to find the value of the quota q that satisfies $\varphi_i(q) = \varphi_j(q)$ (or $\beta_i(q) = \beta_j(q)$). Conversely, if the center prefers player j to player i, it may try to choose the quota so that $\varphi_j(q)$ is strictly greater than $\varphi_i(q)$ (respectively, $\beta_j(q)$ is strictly greater than $\beta_i(q)$). In what follows, we show that both of these tasks are easy to achieve.

Before we present these proofs, note that by symmetry, if $w_i = w_j$, then $\varphi_i(q) = \varphi_j(q)$ and $\beta_i(q) = \beta_j(q)$. That is, if the weights of two players are equal, changing the quota will not change the fact that their powers (under both Shapley–Shubik index and Banzhaf index) are equal.

Theorem 5. *Consider a set of players $I = \{1, \dots, n\}$ and a vector of weights $\mathbf{w} = (w_1, \dots, w_n)$ that satisfies $w_1 \leq \dots \leq w_n$. For each player j there is a quota value q such that for each player i with $w_i < w_j$ it holds that $\varphi_i(q) < \varphi_j(q)$ and $\beta_i(q) < \beta_j(q)$. Also, there is a quota value q' such that for each two players i and j it holds that $\varphi_i(q') = \varphi_j(q')$ and $\beta_i(q') = \beta_j(q')$.*

Proof. To prove the first part of the theorem, let us fix a player j and a quota $q = w_j$. Consider any i with $w_i < w_j$ and any permutation π in which i is pivotal. It is easy to see that j is pivotal for the permutation π' obtained from π by transposing i and j (we have to consider two cases, namely, $\pi(i) < \pi(j)$ and $\pi(i) > \pi(j)$, in both cases the statement is obvious). On the other hand, there are also permutations where j is pivotal, but i would not be pivotal in a permutation obtained by transposing j and i: just consider permutations that start with j. Hence, under the quota $q = w_j$, the number of permutations where j is pivotal is strictly greater than the number of permutations where i is pivotal and so $\varphi_j(q) > \varphi_i(q)$. The proof for the Banzhaf index is similar.

To prove the second part of the theorem, set $q' = w_1$. Then each player i is pivotal for exactly $(n-1)!$ permutations (the ones where he or she appears first), and to exactly one coalition ($\{i\}$). Hence, the Shapley–Shubik indices of all players, as well as their Banzhaf indices, are equal. □

The center may also be interested in finding a quota that ensures that all players have different Shapley–Shubik or Banzhaf indices. This choice can be motivated by fairness, i.e., a desire that a player with a larger weight has strictly more influence than a player with a smaller weight. Unfortunately, it turns out that this is not always possible.

Definition 6. A sequence of positive numbers $(w_1, \dots, w_n)$ is called *super-increasing* if for each $2 \leq k \leq n$, we have $\sum_{j=1}^{k-1} w_j < w_k$.

We will now prove that for any super-increasing weight vector of length at least 3, there is no separating quota.

First we need the following definition.

Definition 7. Given a weighted voting game $G = [I; \mathbf{w}; q]$, players i and j are called *interchangeable* if for every permutation π on I such that i is pivotal for π, transposing i and j makes j pivotal; and for every permutation π on I such that j is pivotal, transposing i and j makes i pivotal.

It is easy to see that if two players are interchangeable, then their Shapley–Shubik indices, as well as their Banzhaf indices, are equal.

Lemma 8. *For any game $G = [I; \mathbf{w}; q]$ with $|I| \geq 3$ and a super-increasing vector of weights $\mathbf{w} = (w_1, \dots, w_n)$, it holds that either players 1 and 2 are interchangeable, or players 2 and 3 are interchangeable.*

Lemma 8 immediately implies the following result.

Theorem 9. *For any game $G = [I; \mathbf{w}; q]$ with $|I| \geq 3$ and a super-increasing vector of weights $\mathbf{w} = (w_1, \dots, w_n)$, either $\varphi_1(q) = \varphi_2(q)$ and $\beta_1(q) = \beta_2(q)$, or $\varphi_2(q) = \varphi_3(q)$ and $\beta_2(q) = \beta_3(q)$. Consequently, there is no separating quota for $\mathbf{w}$.*

The proof of Lemma 8 is omitted from the current version of the paper due to lack of space.

Setting the Quota: Algorithmic Results

In this section, we focus on computational complexity aspects of quota-related problems. These issues are important from practical perspective, as in reality the center may be computationally bounded, and therefore not able to use approaches that require superpolynomial computation time. In what follows, we assume that, unless specified otherwise, the players' weights are given in binary. Hence, we are interested in algorithms whose running time is polynomial in the number of players n and the input description size $\log w(I)$.

The first problem we will study is that of making a given player a dummy. This is a very natural goal for a central authority that strongly dislikes a particular agent: e.g., an election authority that wants to ensure that a particular extremist party has no influence in the parliament. In what follows, we describe a polynomial-time algorithm for this problem.

Definition 10. Given a weight vector $\mathbf{w} = (w_1, \dots, w_n)$ such that $0 < w_1 \leq w_2 \leq \dots \leq w_n$ and a weight w, we say that w is *essential* for $\mathbf{w}$ if for all $1 \leq t \leq n$, $\sum_{i=1}^{t-1} w_i \geq w_t - w$.

The next theorem justifies using the term *essential* in Definition 10: A player whose weight is essential for the vector of weights of the remaining players is never a dummy, irrespective of the choice of the quota value for the game.

Theorem 11. *Let $\mathbf{w} = (w_1, \dots, w_n)$ be a vector of weights such that $0 < w_1 \leq w_2 \leq \dots \leq w_n$. A weight w is essential for $\mathbf{w}$ if and only if there is no quota q, $0 < q \leq w + \sum_{i=1}^{n} w_i$, such that $n+1$ is a dummy in a game $G(q) = [\{1, \dots, n, n+1\}; (w_1, \dots, w_n, w); q]$.*

Proof. Let w, $\mathbf{w}$, and G be as in the statement of the theorem. We first show that if w is not essential for $\mathbf{w}$ then there is a quota q such that $n+1$ is a dummy in $G(q)$. By Definition 10, if $\mathbf{w}$ is not essential then there is an integer t, $1 \leq t \leq n$, such that $w + \sum_{i=1}^{t-1} w_i < w_t$. However, this means that a coalition is successful in $G(w_t)$ if and only if it contains at least one player from the set $\{t, t+1, \dots, n\}$. Thus, adding player $n+1$ to a coalition can never push it from being a losing one to being a winning one and so $n+1$ is a dummy in $G(w_t)$. This completes the first part of the proof.

Let us now assume that w is essential for $\mathbf{w}$. We will show that in this case there is no quota q such that $n+1$ is a dummy in $G(q)$. We need to show that the distance between the adjacent sums of subsets of $\{w_1, \dots, w_n\}$ is no bigger than w. Formally, we will prove that for all integers t, $1 \leq t \leq n$, for all x s.t. $0 \leq x \leq \sum_{i=1}^{t} w_i$ there exists x' such that $0 \leq x' < w$ and $x + x'$ is a sum of some subset of $\{w_1, \dots, w_t\}$.

Our proof follows by induction on t. For the basis, $t = 1$, let x be a real number such that $0 \leq x \leq w_1$. If $x = 0$, define $x' = 0$, and $x + x' = 0$ is a sum of empty subset of $\{w_1\}$. If $0 < x \leq w_1$, define $x' = w_1 - x$. Then $0 \leq x' < w$, and $x + x' = w_1$. For the inductive step we assume that the claim holds for some integer $t - 1$, and we show that this implies our claim for t. Let x be a real number such that $0 \leq x \leq \sum_{i=1}^{t} w_i$. We consider 3 cases:

1. If $x \leq w_t - w$ then, since w is essential for $\mathbf{w}$, $0 \leq x \leq \sum_{i=1}^{t-1} w_i$ and from the inductive assumption there exists $0 \leq x' < w$ s.t. $x + x'$ is a sum of subset of $w_1, \dots, w_{t-1}$.
2. If $w_t - w < x \leq w_t$, then set $x' = w_t - x$, and then $0 \leq x' < w$ and $x + x' = w_t$.
3. If $w_t < x \leq \sum_{i=1}^{t} w_i$ then $0 < x - w_t \leq \sum_{i=1}^{t-1} w_i$, and by the inductive assumption there exists x', $0 \leq x' < w$ such that $x - w_t + x'$ is a sum of subset of $w_1, \dots, w_{t-1}$, therefore $x + x'$ is a sum of subset of $w_1, \dots, w_t$.

This shows that the difference between two adjacent sums of subsets of $\{w_1, \dots, w_n\}$ is at most w. Since for any quota q, $0 < q \leq w + \sum_{i=1}^{n} w_i$, it holds that $\emptyset$ is a losing coalition for $G(q)$ and $\{1, \dots, n+1\}$ is a winning coalition for $G(q)$, there is at least one coalition for which $n + 1$ is pivotal. $\square$

Theorem 11 yields a simple algorithm for testing whether there exists a quota making a specific agent a dummy: it suffices to check whether the weight of that player is essential for the vector of the other players' weights (sorted in nondecreasing order), and this can be done using $O(n)$ additions and comparisons. Using this algorithm, we can now check what quota minimizes the Banzhaf index of an agent.

Theorem 12. *There exists a polynomial time algorithm that finds the value of the quota which minimizes the Banzhaf index of a given player.*

Proof. Use the algorithm described above to check if there is a quota that makes an agent a dummy player, and if so, return this quota. Otherwise, return quota $q = \min\{w_1, \dots, w_n\}$. Under q, the Banzhaf index of our agent is $1/2^{n-1}$, since the only coalition it contributes to is the empty set. $\square$

Comparing Two Values of the Quota

In the previous section we showed that when the center can choose any quota she likes, some of the associated computational problems (e.g., minimizing a player's Banzhaf index) become easy. However, in real-life scenarios, the center may be restricted in the choice of quota: for example, the center might be able to modify the quota only very slightly or have a choice of only several quota values. We show that deciding which of two given quotas favors a player is computationally hard, even if the quotas differ only by 1.

The notion of hardness that we will make use of is PP-hardness, which is believed to be considerably stronger than NP-hardness: any PP-hard problem is NP-hard, but not vice versa. We show that this problem is *in* PP, i.e., that it is PP-complete, thus pinpointing its exact complexity.

We omit the full proof of our hardness result from this version of the paper. In what follows, we formally define our computational problem, describe the problem class PP, state our complexity results, and briefly discuss the techniques used in the proof.

Definition 13. Let f be either the Shapley–Shubik index or the Banzhaf index. In the Quota_f problem we are given a set of players I, $|I| = n$, a vector of weights $\mathbf{w} = (w_1, \dots, w_n)$, two quota values, q' and q'', and an index $i \in I$. Let $G' = [I; \mathbf{w}; q']$, $G'' = [I; \mathbf{w}; q'']$. The task is to decide whether $f_i(G') > f_i(G'')$.

The class PP (see, e.g., (Papadimitriou 1994)) captures the notion of probabilistic polynomial-time computation. The idea is that one can look at nondeterministic computations in terms of a probabilistic ones: An NP machine (a nondeterministic polynomial-time Turing machine) at each computation step tosses a coin to choose the next move uniformly at random from the set of possible ones, as defined by its transition relation. Thus, we can naturally define the probability of an event that an NP machine N accepts a string x. Formally, we say that a language L belongs to PP if there exists an NP machine N such that: $x \in L$ if and only if the probability that N accepts x is at least $\frac{1}{2}$.

PP is a surprisingly powerful class. For example, $\mathrm{NP} \subseteq \mathrm{PP}$ and, in fact, it even holds that $\Theta_2^p \subseteq \mathrm{PP}$ (Beigel, Hemachandra, and Wechsung 1991). (Θ_2^p is the class of decision problems that can be solved via parallel access to NP, also known as $\mathrm{P}^{\mathrm{NP}[\log]}$.) Used as an oracle, PP is essentially as powerful as $\#\mathrm{P}$ (Balcázar, Book, and Schöning 1986); in fact, $\#\mathrm{P}$ can be viewed as a functional counterpart of PP.

There are many PP-complete problems. (Faliszewski and Hemaspaandra 2008) studied the following one.

Definition 14 ((Faliszewski and Hemaspaandra 2008))**.** Let f be either the Shapley-Shubik index or the Banzhaf index. Let PowerCompare$_f$ problem be the following: Given two weighted voting games, G' and G'', a player i in G', and a player j in G'', does it hold that $f_i(G') > f_j(G'')$.

Faliszewski and Hemaspaandra show that this problem is PP-complete both for the Shapley-Shubik index and for the Banzhaf index. They reduce from SAT-Compare, the problem that given two propositional formulas, x and y, asks if $\#\mathrm{SAT}(x) > \#\mathrm{SAT}(y)$, where $\#\mathrm{SAT}(x)$ is the function that takes as input a propositional formula x and returns the number of satisfying assignments for x.

As Quota_f is a special case of PowerCompare$_f$, the result of (Faliszewski and Hemaspaandra 2008) immediately implies that Quota_f is in PP both for $f = \varphi$ and $f = \beta$. To show that Quota_f is PP-hard, rather that using the result of (Faliszewski and Hemaspaandra 2008) as a black box, we make use of a technical lemma proved in that paper, which provides a reduction from SAT-Compare to $\mathrm{SubsetSum}$-Compare that has several useful properties. ($\mathrm{SubsetSum}$-Compare is defined similarly to SAT-Compare, i.e., it compares the number of solutions to two instances of the NP-complete Subset-Sum). We then show that an instance of $\mathrm{SubsetSum}$-Compare output by this reduction can be transformed into an instance of Quota_f for $f = \varphi, \beta$, so that a "yes"-instance of the former problem becomes a "yes"-instance of the latter problem and vice versa. We thus obtain the following theorem.

Theorem 15. Quota_φ *and* Quota_β *are* PP*-complete.*

One can strengthen Theorem 15 as follows.

Theorem 16. Quota_φ *and* Quota_β *remain* PP*-complete even if we restrict them to involve quotas that differ by* 1.

The proof of Theorem 16 is much more involved than the proofs of the previous results, and is therefore omitted.

Discussion Our hardness results show that computational complexity can be a barrier to manipulation by the central authority, as they imply that it will be difficult for the center to choose the quota so as to obtain the desired result. Moreover, as PP is a more powerful complexity class than NP, and our problems are complete for it, the manipulators will not be able to use the existing techniques for problems in NP. However, PP-hardness does not necessarily imply that the problem is hard *on average*; proving that manipulating the quota is hard in this sense is an interesting open problem. Furthermore, even though power indices themselves are hard to compute, a hardness of manipulation result is still significant: power indices reflect the distribution of power among the agents, and the center may want to manipulate this distribution even if it cannot compute it.

On the flip side, it is known (Matsui and Matsui 2000) that both Shapley-Shubik and Banzhaf indices are easy to compute if the weights are polynomially bounded (or, equivalently, given in unary). Clearly, these algorithms can be used to solve Quota_φ and Quota_β, as we can compute a player's power index for both quotas, and choose the one that gives us a better outcome. Hence, computational complexity alone does not provide adequate protection from this form of manipulation, and other approaches are needed.

Conclusion

We have considered quota control manipulations in weighted voting games, where the central authority sets the game's quota to suit its purposes. We have shown the central authority can affect the agents' power by choosing the proper quota, quantified the possible effect of such manipulations, discussed the problem of equalizing and unequalizing agents' power and discussed the computational complexity of finding the proper quota for various purposes. We gave a tractable procedure for testing whether there exists a quota that makes a given player a *dummy*, and shown that checking which of two possible quota values makes a certain agent more powerful is PP-complete.

Several directions remain open for further research. Since manipulations through quota control are possible in weighted voting games, what measures can be taken against such manipulations? Are there restricted domains where there is a polynomial algorithm for checking which quota makes a certain agent more powerful than another agent? Are there other interesting domains where such control manipulations are possible? Are there other payoff division schemes that are more resistant to such manipulations?

References

Bachrach, Y., and Elkind, E. 2008. Divide and conquer: False-name manipulations in weighted voting games. In *Proc. of AAMAS'08*.

Bachrach, Y.; Markakis, V.; Procaccia, A. D.; Rosenschein, J. S.; and Saberi, A. 2008. Approximating power indices. In *Proc. of AAMAS'08*.

Balcázar, J.; Book, R.; and Schöning, U. 1986. The polynomial-time hierarchy and sparse oracles. *Journal of the ACM* 33(3):603–617.

Banzhaf, J. F. 1965. Weighted voting doesn't work: a mathematical analysis. *Rutgers Law Review* 19:317–343.

Beigel, R.; Hemachandra, L.; and Wechsung, G. 1991. Probabilistic polynomial time is closed under parity reductions. *Information Processing Letters* 37(2):91–94.

Conitzer, V., and Sandholm, T. 2004. Computing Shapley values, manipulating value division schemes, and checking core membership in multi-issue domains. In *Proc. of AAAI'04*.

Dubey, P., and Shapley, L. 1979. Mathematical properties of the Banzhaf power index. *Mathematics of Operations Research* 4(2):99–131.

Endriss, U., and Lang, J., eds. 2006. *Proceedings of the 1st International Workshop on Computational Social Choice (COMSOC-2006)*. ILLC, University of Amsterdam.

Faliszewski, P., and Hemaspaandra, L. 2008. The complexity of power-index comparison. Technical Report TR-929, Department of Computer Science, University of Rochester, Rochester, NY.

Fatima, S. S.; Wooldridge, M.; and Jennings, N. R. 2007. A randomized method for the Shapley value for the voting game. In *Proc. of AAMAS'07*, 955–962.

Garey, M. R., and Johnson, D. S. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: W. H. Freeman and Company.

Leech, D. 2002. Voting power in the governance of the international monetary fund. *Annals of Operations Research* 109(1-4):375–397.

Machover, M., and Felsenthal, D. S. 2001. The treaty of Nice and qualified majority voting. *Social Choice and Welfare* 18(3):431–464.

Matsui, Y., and Matsui, T. 2000. A survey of algorithms for calculating power indices of weighted majority games. *Journal of the Operations Research Society of Japan* 43.

Papadimitriou, C. M. 1994. *Computational complexity*. Reading, Massachusetts: Addison-Wesley.

Prasad, K., and Kelly, J. S. 1990. NP-completeness of some problems concerning voting games. *Internat. J. Game Theory* 19(1):1–9.

Roth, A. E., ed. 1988. *The Shapley Value: Essays in Honor of Lloyd S. Shapley*. Cambridge University Press.

Shapley, L. S., and Shubik, M. 1954. A method for evaluating the distribution of power in a committee system. *American Political Science Review* 48:787–792.

Shapley, L. S. 1953. A value for n-person games. *Contributions to the Theory of Games* 31–40.

Yokoo, M.; Conitzer, V.; Sandholm, T.; Ohta, N.; and Iwasaki, A. 2005. Coalitional games in open anonymous environments. In *Proceedings of AAAI'05*.