

A Characterization of Voting Power for Discrete Weight Distributions

Yoram Bachrach
 Microsoft Research, UK
 yobach@microsoft.com

Yuval Filmus
 Technion, Israel
 yuvalfi@cs.technion.ac.il

Joel Oren
 Univ. of Toronto, Canada
 oren@cs.toronto.edu

Yair Zick
 Carnegie Mellon Univ., USA
 yairzick@cs.cmu.edu

Abstract

Weighted voting games model decision-making bodies where decisions are made by a majority vote. In such games, each agent has a weight, and a coalition of agents wins the game if the sum of the weights of its members exceeds a certain quota. The Shapley value is used as an index for the true power held by the agents in such games.

Earlier work has studied the implications of setting the value of the quota on the agents' power under the assumption that the game is given with a fixed set of agent weights. We focus on a model where the agent weights originate from a stochastic process, resulting in weight uncertainty. We analyze the expected effect of the quota on voting power given the weight generating process. We examine two extreme cases of the balls and bins model: uniform and exponentially decaying probabilities. We show that the choice of a quota may have a large influence on the power disparity of the agents, even when the governing distribution is likely to result in highly similar weights for the agents. We characterize various interesting repetitive fluctuation patterns in agents' power as a function of the quota.

1 Introduction

Group decision making systems are prevalent in AI settings; expert systems and prediction markets have been commonly modeled using *weighted voting games*. The weighted voting game model, historically used by political scientists and economists to model political parties, has been the object of intense research by the AI/COMSOC community; in this setting, each agent (thought of as a political party) has a weight w_i . A group of agents is said to be winning (i.e., can form a government, or pass a bill) if its total weight exceeds a given quota q .

One key observation is that the true decision power of an agent does not necessarily correspond to its weight. For example, consider a parliament that has three parties, two with 49 seats, and one with 2 seats. Assuming that a majority of the votes (i.e., 50 votes) is required in order to pass a bill, all

three parties are equally critical to the legislative process: no single party can pass a bill on its own, whereas any two parties can. Thus, although one party has a far smaller weight than the other two, it holds the same electoral power. Various *power indices* have been proposed to measure the actual influence of an agent in such settings. One of the most prominent power indices is the Shapley power index (also referred to as the Shapley–Shubik power index). This index has played a central role in the analysis of real-life voting systems, such as the US electoral college, the EU council of members, and the UN security council. Empirical studies of weighted voting present an interesting phenomenon: changes to the quota of weights required to win the game can dramatically affect the agents' voting power. This phenomenon has been observed in real parliaments (e.g., the EU council of members) and in artificial WVGs [Leech and Machover, 2003; Zuckerman *et al.*, 2012].

Prior work (see Section 1.2) studied the following question: for a *given* set of weights, how does voting power change as a function of the quota? In this scenario, a central authority wants to enforce a certain property on the agents' voting power. For example, the central authority may be interested in minimizing or maximizing a certain party's electoral power, ensuring that all agents have equal power, or that agents have power proportional to their weight. To achieve its goal, the central authority may control the quota prior to the formation of the electorate; for example, quota manipulation has been considered in order to ensure power/weight proportionality in the EU council of members.

Earlier work examining the impact of control over the quota has shown relatively weak upper bounds on the potential difference in power caused by changes to the quota. For example, Zuckerman *et al.* [2012] provide an upper bound of $1/(n-i+1)$ on the additive difference in the Shapley value of the agent with the i -th smallest weight, as a result of changing the quota, as well as a simple criterion for the equivalence of the Shapley values under two different quota values. While some worst case analysis is shown to be tight, its conclusions seem unsatisfactory: these results merely tell us that for *some* weight vectors, power can vary widely, even when changes to the quota are small. Indeed, Zick *et al.* [2011] show that agent i may have minimal electoral power if the threshold is

$w_i + 1$, while its voting power is maximized when $q = w_i$.

Even though worst case results do not seem promising, there exists virtually no work analyzing the *average* case; this is where our work comes in. We focus on the following question:

What is the likely effect of quota changes on the Shapley value, when weights are sampled from some known prior?

In contrast to the models underlying previous work, many real world settings exhibit *uncertainty* about the agent weights. As a motivating example, consider the case where one must decide on the number of votes required to pass a bill. If we assume that changes like these are not common (as is the case for the EU council of members), one must consider not only voting power in the current parliament, but in future ones. Since we have no data on the composition of a future parliament, the expected effects of our choice of quota must be assessed based on some probabilistic model of vote distribution, obtained from polling data or a stochastic model.

1.1 Our Contributions

Using natural weight generation processes, we analyze the expected behavior of the Shapley value as a function of the quota. Our results show that voting power can behave in a rather unusual manner; for example, we show that even when weights are likely to be very similar, some quota choices are likely to cause significant differences in voting power. The current body of work represents a significant advance in our understanding of weighted voting games; via careful probabilistic analysis, we strongly generalize previous known results, and inform the design of both real and randomly generated voting systems.

Our work focuses on the *Balls and Bins* model [Mitzenmacher *et al.*, 2000; Mitzenmacher and Upfal, 2005; Raab and Steger, 1998]—a model that has received considerable attention in the computer science community. Briefly, m balls are independently thrown into n bins, where each ball lands in the i -th bin with probability p_i .

In the voting context, this process takes on an intuitive meaning: each ball represents a single voter, who votes for party i with probability p_i ; thus, the weight of a party after m tosses corresponds to the number of votes it received.

In Section 3, we study a simple model, where each ball lands in one of the n bins uniformly at random. We identify a repetitive fluctuation pattern in the Shapley values, with cycles of length $\frac{m}{n}$. We show that if the quota is sufficiently bounded away from the borders of its length- $\frac{m}{n}$ cycle, then the Shapley values of all agents are likely to be very close to each other. On the other hand, we show that due to noise effects, when the quota is situated close enough to small multiples of $\frac{m}{n}$, the highest Shapley value can be roughly double that of the smallest one.

In Section 4 we consider the case in which the probabilities decay exponentially, with a decay factor smaller than $1/2$. We show that analyzing this case essentially boils down to characterizing the Shapley values in a game where weights are a super-increasing sequence, which has already received attention (e.g., [Oren *et al.*, 2014; Aziz and Paterson, 2008; Zuckerman *et al.*, 2012]).

1.2 Related Work

Weighted voting games are a fundamental class of cooperative games, modeling several phenomena in AI domains (voting, threshold logics, task completion and others); moreover, their structure nicely lends itself to computational analysis and extensions (for more on WVGs in AI see [Chalkiadakis *et al.*, 2011; Chalkiadakis and Wooldridge, 2016]). Measuring agent influence in WVGs is commonly done using the Shapley value [Shapley, 1953]; this is because the Shapley value has several appealing properties, and is, in fact, the *only* function that has these properties (see Peleg and Sudhölter [2007] for details).

Computing the Shapley value is known to be computationally intractable [Chalkiadakis *et al.*, 2011], but easily approximable via random sampling [Bachrach *et al.*, 2010; Fatima *et al.*, 2008]. Approximation techniques exploit the inherent probabilistic nature of power indices, rather than assuming independent randomness in the weighted voting game itself, as we do here.

If one makes no assumptions on weight distributions then very little can be said about the effects of the quota on WVGs. Indeed, power measures are highly sensitive to varying quota values [Zick *et al.*, 2011; Zick, 2013], though predicting the effect of quota variation is computationally hard [Zuckerman *et al.*, 2012]. Our work takes a more principled approach to the matter.

The assumption of a prior on voter preferences is also a well-established practice. Many models of probabilistic voting have been discussed in the economic literature (see e.g., [Coughlin, 1992; Calvert, 1985; Enelow and Hinich, 1989]). More broadly, social choice theory has seen a recent surge in models that incorporate uncertainty on agent preferences (e.g., [Caragiannis *et al.*, 2013; 2014; Lu and Boutilier, 2011; Oren *et al.*, 2013]).

Several works have studied the effects of randomization on weighted voting games from a theoretical perspective [Tausman and Jelnov, 2012; Lindner, 2004; Penrose, 1946; Zick, 2013]. Finally, the effects of changes to the quota have also been studied empirically, mostly in the context of the EU council of members [Leech and Machover, 2003; Leech, 2002; Słomczyński and Życzkowski, 2006].

2 Preliminaries

General notation Given a vector $\mathbf{x} \in \mathbb{R}^n$ and a set $S \subseteq \{1, \dots, n\}$, let $x(S) = \sum_{i \in S} x_i$. For a random variable X , we let $\mathbb{E}[X]$ be its expectation, and $\text{Var}[X]$ be its variance. For a set S , we denote by $\binom{S}{k}$ the collection of subsets of S of cardinality k . The notation $T \in_R \binom{S}{k}$ means that the set T is chosen uniformly at random from $\binom{S}{k}$. We let $B(n, p)$ denote the binomial distribution with n trials and success probability p .

We let $O_p(\cdot)$ denote the usual big-O notation, conditioned on a fixed value of p . In other words, having $f(n) = O_p(g(n))$ means that there exist functions $K(\cdot), N(\cdot)$, such that for $n \geq N(p)$, $f(n) \leq K(p) \cdot g(n)$.

Finally, for a distribution D over \mathbb{R} and some event \mathcal{E} , we simplify our notation by letting $\Pr[\mathcal{E}(D)] = \Pr_{x \sim D}[\mathcal{E}(x)]$.

For example, for $a > 0$, we can write $\Pr[B(n, p) \leq a] = \Pr_{x \sim B(n, p)}[x \leq a]$.

Weighted voting games A *weighted voting game* (WVG) is given by a set of agents $N = \{1, \dots, n\}$, where each agent $i \in N$ has a non-negative weight w_i , and a *quota* q . Unless otherwise specified, we assume that the weights are arranged in non-decreasing order, $w_1 \leq \dots \leq w_n$.

A subset of agents $S \subseteq N$ is called *winning* (has value 1) if $w(S) \geq q$ and is called *losing* (has value 0) otherwise.

The Shapley value Let Sym_n be the set of all permutations of N . Given some permutation $\sigma \in \text{Sym}_n$ and an agent $i \in N$, we let $P_i(\sigma) = \{j \in N : \sigma(j) < \sigma(i)\}$; $P_i(\sigma)$ is called the set of i 's predecessors in σ . We say that i is *pivotal* for a set $S \subseteq N$ if S is losing but $S \cup \{i\}$ is winning. Similarly, i is pivotal for a permutation σ if the set $P_i(\sigma)$ is losing, but the set $P_i(\sigma) \cup \{i\}$ is winning. In other words, i is pivotal for a permutation σ if its predecessors have a total weight lower than q , but $w(P_i(\sigma)) + w_i \geq q$.

The Shapley–Shubik power index (often referred to as the Shapley value in the context of WVG's) is simply the probability that i is pivotal for a permutation $\sigma \in \text{Sym}_n$ selected uniformly at random. More explicitly,

$$\varphi_i(\mathbf{w}; q) = \frac{1}{n!} \sum_{\sigma \in \text{Sym}_n} \mathbb{I}(q - w_i \leq w(P_i(\sigma)) < q).$$

Here, $\mathbb{I}(\cdot)$ is the indicator function. Since $\sigma^{-1}(i)$ is distributed uniformly when σ is chosen at random from Sym_n , we also have the alternative formula

$$\varphi_i(\mathbf{w}; q) = \frac{1}{n} \sum_{\ell=0}^{n-1} \sum_{S \in \mathcal{R} \binom{N \setminus \{i\}}{\ell}} \mathbb{I}(w(S) < q \wedge w(S) + w_i \geq q). \quad (1)$$

When the WVG $\langle \mathbf{w}; q \rangle$ is clear from the context, we will simply write φ_i ; however, we often wish to emphasize the role of the quota q , and refer to the Shapley value of i as $\varphi_i(q)$. Finally, it is easy to show that $w_i \leq w_j$ implies $\varphi_i \leq \varphi_j$, and so if $w_1 \leq \dots \leq w_n$, then $\varphi_1 \leq \dots \leq \varphi_n$. Another useful property that follows immediately from the definitions is that $\sum_{i \in N} \varphi_i = 1$, assuming $0 < q \leq \sum_{i \in N} w_i$.

The Balls and Bins Distribution In its general form, the Balls and Bins distribution is derived as follows: given a set of n bins and a distribution represented by a vector $\mathbf{p} \in [0, 1]^n$ such that $\sum_{i=1}^n p_i = 1$, the process unfolds in m steps. At each step, a ball is thrown into one of the bins based on the probability vector \mathbf{p} . The resulting weights are then sorted in non-decreasing order $w_1 \leq \dots \leq w_n$.

3 The Balls and Bins Distribution: the Uniform Case

We begin our study of the balls and bins process by considering the most commonly studied version of the balls and bins model, in which each ball is thrown into one of the bins with equal probability, i.e., $p_i = 1/n$, for all $i \in N$.




Figure 1: The Shapley values of agents 1, 10, 20 and 30 in a 30-agent WVG where weights were drawn from a balls and bins distribution with $m = 10,000$ balls. The vertical lines are placed at integer multiples of $\frac{m}{n} = 333\frac{1}{3}$.

As Figure 1 shows for the case of $n = 30$, the behavior of the Shapley values demonstrates an almost perfect cyclic pattern, with intervals of length m/n in Figure 1). As can be seen in the figure, for quota values that are sufficiently distant from the interval endpoints, all of the Shapley values tend to be equal (as the Shapley values of the highest and lowest agents are equal in these regions). As the number of balls grows, all of the bins tend to have nearly the same number of balls in them; however, low weight discrepancy does not immediately translate to low power discrepancy: we can guarantee nearly equal voting power for some quotas, but not for others. Example 3.1 provides some intuition to this behavior, though the formal proofs turn out to be much more involved.

Example 3.1. Let us assume $m = 300$ and $n = 3$; under a uniform balls and bins distribution, we are very likely to obtain three weights $w_1 = W - \alpha$, $w_2 = W$, $w_3 = W + \beta$, where $W \simeq 100$, and with very small α and β . If we choose a quota well away from multiples of $\frac{m}{n} = 100$, say $q = 150$, then it is extremely likely that all agents are equally influential (any two parties form a winning coalition). However, when we choose a quota that is close to multiples of 100 (say $q = 100$), it is likely to place a lot of power in the hands of the largest party: it can always form a winning coalition on its own.

From a legislative perspective, our theorems provide some interesting guarantees on voting power, under the assumption that voter behavior follows a uniform Balls and Bins distribution. If one wishes to ensure that all parties have similar voting power, it suffices to set a quota that is sufficiently far from $\frac{m}{n}$ (and is close enough to 50%); if one wishes to ensure power disparity, setting a quota closer to $\frac{m}{n}$ is desirable.

We begin by providing a formula for the differences between two Shapley values.

Lemma 3.2. For all agents $i, j \in N$, suppose that $w_i \leq w_j$;

then

$$|\varphi_j - \varphi_i| = \frac{1}{n-1} \cdot \sum_{\ell=0}^{n-2} \Pr_{S \in \mathcal{R}[\binom{N \setminus \{i,j\}}{\ell}]} \left[q - w_j \leq w(S) < q - w_i \right].$$

We now give a theoretical justification for the near-identity of Shapley values for quotas that are well away from integer multiples of $\frac{m}{n}$. In this section, we do not assume that the weights w_1, \dots, w_n are ordered, in order to maintain the fact that the weights are random variables.

The idea of the proof is to use the following criterion, which is a consequence of Lemma 3.2 (proof omitted):

Proposition 3.3. *Suppose that for all agents $i, j \in N$ and for all subsets $S \subseteq N \setminus \{i, j\}$, we have $q \notin (w(S \cup \{i\}), w(S \cup \{j\}))$. Then all Shapley values are equal to $1/n$.*

Next, we show that the weights $w(S)$ are concentrated around points of the form $\ell \frac{m}{n}$.

Lemma 3.4. *Suppose that $m > 3n^2$. With probability $1 - 2(\frac{2}{e})^n$, the following holds: for all $S \subseteq N$, $|w(S) - \frac{|S|m}{n}| \leq \sqrt{3nm}$.*

Proof. The proof uses a straightforward Chernoff bound [Mitzenmacher and Upfal, 2005]. We can assume that $S \neq \emptyset$ (as otherwise the bound is trivial). For each non-empty set $S \subseteq N$, the distribution of $w(S)$ is $B(m, \frac{|S|}{n})$. Therefore for $0 < \delta < 1$,

$$\Pr \left[\left| w(S) - \frac{|S|m}{n} \right| > \delta \frac{|S|m}{n} \right] \leq 2e^{-\frac{\delta^2 |S|m}{3n}}.$$

Choosing $\delta = \sqrt{\frac{3n^2}{|S|m}} < 1$, we obtain a bound of $2e^{-n}$. Since there are 2^n possible sets S , a union bound implies that $|w(S) - \frac{|S|m}{n}| \leq \sqrt{3nm}$ with probability at least $1 - 2(\frac{2}{e})^n$. \square

Finally, we require the following simple property of quotas.

Proposition 3.5. *Let $n \leq m$ be two integers, then for any $q \in (0, m]$, there exists some $\ell \leq n$ such that $|q - \ell m/n| \geq m/n$.*

Theorem 3.6 is an immediate corollary of the above claims, as we now show.

Theorem 3.6. *Let $M = \frac{m}{3n^3}$. Suppose that $|q - \frac{\ell m}{n}| > \frac{1}{\sqrt{M}} \frac{m}{n}$ for all integers ℓ . Then with probability $\geq 1 - 2(\frac{2}{e})^n$, all Shapley values are equal to $\frac{1}{n}$.*

Proof. First, note that $M > 1$, as otherwise, it would imply that for all $\ell = 1, \dots, n$, $|q - \ell m/n| > \frac{1}{\sqrt{M}} \frac{m}{n} \geq \frac{m}{n}$. This is impossible according to Proposition 3.5. Thus, $M > 1$, i.e., $m > 3n^3 \geq 3n^2$.

Lemma 3.4 shows that when $m > 3n^3$, with probability $1 - 2(\frac{2}{e})^n$, for all sets S we have $|w(S) - \frac{|S|m}{n}| \leq \sqrt{3nm}$; we assume this condition on $w(S)$ henceforth. Suppose, for the sake of obtaining a contradiction, that $\varphi_i < \varphi_j$ for some agents i, j . Then Proposition 3.3 shows that there must exist

some $S \subseteq N \setminus \{i, j\}$ such that $q \in (w(S \cup \{i\}), w(S \cup \{j\}))$. Since both $w(S \cup \{i\})$ and $w(S \cup \{j\})$ are $\sqrt{3nm}$ -close to $\frac{(|S|+1)m}{n}$, this implies that $|q - \frac{(|S|+1)m}{n}| \leq \sqrt{3nm}$. However,

$$\sqrt{3nm} = \frac{m}{n} \sqrt{\frac{3n^3 m}{m^2}} = \sqrt{\frac{3n^3}{m}} \cdot \frac{m}{n} = \frac{1}{\sqrt{M}} \cdot \frac{m}{n},$$

contradicting our assumption that $|q - \frac{\ell m}{n}| \geq \frac{1}{\sqrt{M}} \cdot \frac{m}{n}$ for all ℓ . We conclude that with probability at least $1 - 2(\frac{2}{e})^n$, all agents have the same Shapley value $1/n$. \square

Theorem 3.6 implies that if the voter population is much larger than the number of candidates, and votes are assumed to be cast uniformly at random, then choosing a quota that is well away from a multiple of $\frac{m}{n}$ will most probably lead to an even distribution of power among the elected representatives.

3.1 How Weak Can the Weakest Agent Get?

As Theorem 3.6 demonstrates, if the quota is sufficiently bounded away from any integral multiple of $\frac{m}{n}$, then the distribution of power tends to be even among the agents. When the quota is close to an integer multiple of $\frac{m}{n}$, the resulting weighted voting game may not display such an even distribution of power; this is a result of weight perturbations originating in the intrinsic ‘‘noise’’ involved in the process.

Motivated by these observations, we now proceed to study the expected Shapley value of the weakest agent, φ_1 (recall that we assume that the weights are given in non-decreasing order).

We present two contrasting results. Let $q = \ell \cdot \frac{m}{n}$, for an integer ℓ . When $\ell = o(\log n)$, we show that the expected minimal Shapley value is roughly $\frac{1}{2n}$, and so it is at least half the maximal Shapley value, in expectation.

Theorem 3.7. *Let $q = \ell \cdot \frac{m}{n}$ for some integer $\ell = o(\log n)$. For $m = \Omega(n^3 \log n)$, $\mathbb{E}[\varphi_1] = 1/(2n) \pm o(1/n)$.*

In contrast, when $\ell = \Omega(n)$, this effect disappears.

Theorem 3.8. *Let $q = \ell \cdot \frac{m}{n}$ for $\ell \in \{1, \dots, n\}$ such that $\gamma \leq \frac{\ell}{n} \leq 1 - \gamma$ for some constant $\gamma > 0$. Then for $m = \Omega(n^3)$, $\mathbb{E}[\varphi_1] \geq 1/n - O_\gamma(\sqrt{\log n/n^3})$.*

The idea behind the proof of both theorems is the formula for φ_1 given in Lemma 3.9 (the full proof is replaced by a sketch due to space constraints). In this formula and in the rest of the section, the probabilities are taken over both the displayed variables and the choice of weights.

Lemma 3.9. *Let $q = \ell \cdot \frac{m}{n}$, where $\ell \in \{1, \dots, n-1\}$. For $m = \Omega(n^3 \log n)$, $\mathbb{E}[\varphi_1]$ equals*

$$\frac{\left(\frac{1}{2} - \frac{\ell}{n} + \Pr_{A \in \mathcal{R}[\binom{N \setminus \{1\}}{\ell-1}]} [w(A \cup \{1\}) \geq q] \right)}{n - \ell} \pm O\left(\frac{1}{n^2}\right)$$

Proof Sketch. Let $p_k = \Pr_{A \in \mathcal{R}[\binom{N \setminus \{k\}}{\ell-1}]} [q - w_1 \leq w(A) < q]$; Formula (1) shows that $\mathbb{E}[\varphi_1] = \frac{1}{n} \sum_{k=0}^{n-1} p_k$. We then consider three cases, corresponding to possible sizes of the set A in the formula for p_k ; each of these cases will contribute a

term in expression of the lemma. Since $w(A) \approx \frac{|A|m}{n}$, when $|A| \geq \ell + 1$ it is highly unlikely that $w(A) < q$. Similarly, since $w(A) + w_1 \approx \frac{(|A|+1)m}{n}$, when $|A| \leq \ell - 2$ it is highly unlikely that $w(A) \geq q - w_1$. So roughly speaking, $\mathbb{E}[\varphi_1] \approx \frac{p_{\ell-1} + p_\ell}{n}$. Furthermore, when $|A| = \ell - 1$, it is very likely that $w(A) < q$, and when $|A| = \ell$, it is very likely that $w(A) \geq q - w_1$. So $\mathbb{E}[\varphi_1]$ equals approximately

$$\frac{1}{n} \Pr_{A \in R[\binom{N}{\ell-1}]} [w(A) + w_1 \geq q] + \frac{1}{n} \Pr_{A \in R[\binom{N}{\ell}]} [w(A) < q].$$

The trick now is to relate the two terms; some manipulation shows that $\Pr_{A \in R[\binom{N}{\ell}]} [w(A) < q]$ equals

$$\frac{\left(n \Pr_{A \in R[\binom{N}{\ell}]} [w(A) < q] - \ell \Pr_{A \in R[\binom{N}{\ell-1}]} [w(A \cup \{1\}) < q] \right)}{n - \ell}.$$

To address the first term in the above expression, note that when $|A| = \ell$, $\mathbb{E}[w(A)] = q$, and so the first probability is roughly $1/2$. Therefore $\Pr_{A \in R[\binom{N}{\ell}]} [w(A) < q]$ is approximately

$$\frac{n}{2(n - \ell)} - \frac{\ell}{n - \ell} + \frac{\ell}{n - \ell} \Pr_{A \in R[\binom{N}{\ell-1}]} [w(A) + w_1 \geq q].$$

Substituting this in our estimate for $\mathbb{E}[\varphi_1]$, we obtain the desired result. \square

In order to estimate the expression $\Pr_{A \in R[\binom{N}{\ell-1}]} [w(A) + w_1 \geq q]$, we need a good estimate for w_1 . Such an estimate is given by the following lemma.

Lemma 3.10. *With probability $1 - 2/n$, we have that $\sqrt{m \log n / (3n)} \leq \frac{m}{n} - w_1 \leq \sqrt{4m \log n / n}$.*

We obtain this bound by applying the *Poisson approximation technique* to the Balls and Bins process, which we now roughly describe. Consider the case of a random event, defined with respect to the weight distribution induced by the process. The probability of the event can be well-approximated by the probability of an analogous event, defined with respect to n *i.i.d.* Poisson random variables, assuming the event is monotone in the number of balls.

We can now prove Theorem 3.7.

Proof of Theorem 3.7. It is not hard to show that $\Pr_{A \in R[\binom{N}{\ell-1}]} [w(A) + w_1 \geq q]$ is at most

$$\frac{n}{n - \ell + 1} \Pr[B(m, \frac{\ell-1}{n}) \geq q - w_1].$$

The concentration bound on w_1 (Lemma 3.10) shows that with probability $1 - 2/n$, $q - w_1 \geq \frac{(\ell-1)m}{n} + \sqrt{\frac{m \log n}{3n}}$. Assuming this, a Chernoff bound gives

$$\Pr[B(m, \frac{\ell-1}{n}) \geq q - w_1] \leq \Pr[B(m, \frac{\ell-1}{n}) \geq \frac{(\ell-1)m}{n} + \sqrt{\frac{m \log n}{3n}}] \leq e^{-\frac{m \log n / (3n)}{3(\ell-1)m/n}},$$

which is $o(1)$ using $\ell = o(\log n)$. Accounting for possible failure of the bound on $q - w_1$, we obtain

$$\Pr_{A \in R[\binom{N}{\ell-1}]} [w(A) + w_1 \geq q] \leq \left(1 - \frac{2}{n}\right) \cdot o\left(\frac{n}{n - \ell}\right) + \frac{2}{n} \cdot 1,$$

which is $o(1)$ assuming $\ell = o(\log n)$. Lemma 3.9 therefore shows that

$$\mathbb{E}[\varphi_1] \leq \frac{1}{2(n - \ell)} + o\left(\frac{1}{n - \ell}\right) + O\left(\frac{1}{n^2}\right) = \frac{1}{2n} + o\left(\frac{1}{n}\right),$$

since $\ell = o(\log n)$ implies $\frac{1}{n - \ell} = \frac{1}{n} + \frac{\ell}{n(n - \ell)} = \frac{1}{n} + o\left(\frac{1}{n}\right)$. Lemma 3.9 also implies a matching lower bound:

$$\mathbb{E}[\varphi_1] \geq \frac{1}{2(n - \ell)} - \frac{\ell}{n(n - \ell)} - O\left(\frac{1}{n^2}\right) \geq \frac{1}{2n} - o\left(\frac{1}{n}\right). \quad \square$$

In the regime of ℓ addressed by Theorem 3.7, $\Pr_{A \in R[\binom{N}{\ell-1}]} [w(A) + w_1 \geq q]$ was negligible. In contrast, in the regime of ℓ addressed by Theorem 3.8, $\Pr_{A \in R[\binom{N}{\ell-1}]} [w(A) + w_1 \geq q] \approx 1/2$, as the following lemma, which is proved in the full version of the paper using the Berry-Esseen theorem [Berry, 1941], shows.

Lemma 3.11. *Suppose $q = \frac{\ell m}{n}$ for an integer ℓ satisfying $\gamma \leq \frac{\ell-1}{n} \leq 1 - \gamma$, and let*

$$t_\varepsilon = \Pr_{A \in R[\binom{N}{\ell-1}]} [w(A) + w_1 \geq q : w_1 = \frac{m}{n} - \varepsilon \sqrt{m \log n / n}].$$

Then for $m \geq 4n^3$, we have that $t_\varepsilon \geq \frac{1}{2} - \frac{\varepsilon}{2\pi\gamma} \sqrt{\log n / n} - 1/n$.

As Lemma 3.10 shows, $1/3 \leq \varepsilon \leq 4$ with probability $1 - 2/n$, which explains the usefulness of this bound. We can now prove Theorem 3.8.

Proof of Theorem 3.8. Lemma 3.10 shows that with probability $1 - 2/n$, $w_1 = \frac{m}{n} - \varepsilon \sqrt{\frac{m \log n}{n}}$ for some $1/3 \leq \varepsilon \leq 4$, in which regime Lemma 3.11 shows that $t_\varepsilon \geq \frac{1}{2} - \frac{2}{\pi\gamma} \sqrt{\frac{\log n}{n}} - \frac{1}{n}$. Accounting for the case in which ε is out of bounds,

$$\Pr_{A \in R[\binom{N}{\ell-1}]} [w(A) + w_1 \geq q] \geq \left(1 - \frac{2}{n}\right) \left(\frac{1}{2} - \frac{2}{\pi\gamma} \sqrt{\frac{\log n}{n}} - \frac{1}{n}\right) \geq \frac{1}{2} - \frac{2}{\pi\gamma} \sqrt{\frac{\log n}{n}} - \frac{3}{n}.$$

Finally, using Lemma 3.9, we can show that $\mathbb{E}[\varphi_1]$ is at least $\frac{1}{n} - O_\gamma\left(\sqrt{\frac{\log n}{n^3}}\right)$. \square

4 The Balls and Bins Distribution: the Exponential Case

In Section 3, we showed that even when the distribution is not inherently biased towards any agent, substantial inequalities may arise due to random noise. We now turn to study the case in which the distribution is strongly biased. Returning to our formal definition of the general balls and bins process,

we assume that the probabilities in the vector \mathbf{p} are ordered in increasing order and $p_i/p_{i+1} = \rho$, for some $\rho < 1/2$; we henceforth refer to this distribution as ρ -exponential. We observe that as m approaches ∞ , the weight vector follows a power law with probability 1, where for each $i = 1, \dots, n-1$, $w_i/w_{i+1} = \rho$. Super-increasing weight vectors [Zuckerman *et al.*, 2012] turn out to naturally arise under this distribution. A series of positive weights $\mathbf{w} = (w_1, \dots, w_n)$ is said to be super-increasing (SI) if for every $i = 1, \dots, n$, $\sum_{j=1}^{i-1} w_j < w_i$.

The following three results (Lemma 4.1, Lemma 4.2 and Theorem 4.3) show that for a sufficiently large value of m , estimating the Shapley values in WVG's where the weights are sampled from an exponential distribution can be reduced to the study of Shapley values in a game with a prescribed (fixed) SI weight vector. This insight is useful, as voting power in SI WVGs is well understood (e.g., [Aziz and Paterson, 2008; Zuckerman *et al.*, 2012; Oren *et al.*, 2014]).

The following lemma shows that when the weights are sampled from an exponential distribution, given enough voters, it is highly probable that the resulting weights (i.e. the number of balls in each bin) are super-increasing.

Lemma 4.1. *Assume that m voters submit the votes according to the ρ -exponential distribution for some $0 < \rho < 1/2$. There is a (universal) constant $C > 0$ such that if $m \geq C\rho^{-n}(1 - 2\rho)^{-2} \log n$ then the resulting weight vector is super-increasing with probability $1 - O(\frac{1}{n})$. Also, as $m \rightarrow \infty$, the probability approaches 1.*

The proof of the lemma uses a standard concentration bound, and is omitted due to space constraints.

Before we proceed, it would be helpful to provide some intuition about the behavior of the Shapley values. Assuming that agent weights are given by an increasing sequence \mathbf{w} of n reals, consider the set of all distinct subset sums of the weights $\mathcal{S}(\mathbf{w}) = \{s : \exists S \subseteq N \text{ s.t. } s = w(S)\}$. We assume that $\mathcal{S}(\mathbf{w})$ is ordered; i.e., $\mathcal{S}(\mathbf{w}) = \{s_j\}_{j=1}^t$, such that $s_j < s_{j+1}$. It is easy to show that $\varphi_i(q)$ is constant within the interval $(s_j, s_{j+1}]$ for all $s_j, s_{j+1} \in \mathcal{S}(\mathbf{w})$. Given a set $S \subseteq N$, we let S_+ be the subsequent set in the ordering induced by $\mathcal{S}(\mathbf{w})$; thus, there exists a unique interval $I^{\mathbf{w}}(S) = (s_j, s_{j+1}]$ such that $w(S) = s_j, w(S_+) = s_{j+1}$.

Suppose that \mathbf{w} is generated using a Balls and Bins process with probabilities \mathbf{p} , where \mathbf{p} is a SI sequence; then it stands to reason that if a sufficiently large number of balls is tossed (i.e., m is large enough), then the voting power distribution under \mathbf{w} will be very close to the power distribution under the weight vector \mathbf{p} . This intuition is captured in the following lemma, whose proof is omitted.

Lemma 4.2. *Suppose that $\mathbf{p} = (p_1, \dots, p_n)$ is a SI sequence summing to 1, and let w_1, \dots, w_n be obtained by sampling m balls from the distribution \mathbf{p} .*

Given some $\tau \in \mathcal{IP}(S)$ for some $S \subseteq N$, such that $|\tau - w(S)|, |\tau - w(S_+)| \geq \sqrt{\log(nm)/n}$. If \mathbf{w} is SI, then

$$\Pr[\forall i \in N, \varphi_i(\mathbf{w}; m\tau) = \varphi_i(\mathbf{p}; \tau)] \geq 1 - \frac{2}{(nm)^2}.$$

Combining both lemmas, we obtain our main result on the exponential case of the Balls and Bins distribution.

Theorem 4.3. *Assume that $m \geq C\rho^{-n}(1 - 2\rho)^{-2} \log n$ voters submit their votes according to the ρ -exponential distribution, for some $0 < \rho < 1/2$. Given $\tau \in \mathcal{IP}(S)$*

such that $|\tau - w(S)|, |\tau - w(S_+)| \geq \sqrt{\frac{\log(nm)}{n}}$, then $\Pr[\forall i \in N, \varphi_i(\mathbf{w}; mT) = \varphi_i(\mathbf{p}; T)] \geq 1 - O(\frac{1}{n})$. Furthermore, for all but finitely many values of $T \in (0, 1]$, $\lim_{m \rightarrow \infty} \Pr[\varphi_i(\mathbf{w}; mT) = \varphi_i(\mathbf{p}; T)] = 1$.

Proof. Lemma 4.1 gives a constant $C > 0$ such that if $m \geq C\rho^{-n}(2\rho - 1)^{-2} \log n$ then \mathbf{w} is SI with probability $1 - O(1/n)$. Hence the first part of the theorem follows from Lemma 4.2.

For the second part, Lemma 4.1 shows that as $m \rightarrow \infty$, the probability that \mathbf{w} is SI approaches 1. Suppose now that τ does not belong to $\mathcal{S}(\mathbf{w})$ (these are the finitely many exceptions). When m is large enough, the conditions of Lemma 4.2 are satisfied, and so as $m \rightarrow \infty$, the error probability in that lemma goes to 0. The second part of the theorem follows. \square

The theorem shows that in the case of the exponential distribution, if the number of balls is large enough then we can calculate with high probability the Shapley values of the resulting distribution based on the Shapley values of the original exponential distribution (without sampling).

5 Conclusions and Future Work

We study the Shapley value as a function of the quota under a number of natural weight distributions, allowing us to reason rigorously about the effect of quota changes. We were also able to show that when the probability vectors constitute SI sequences, these probabilities can be treated as if they were the deterministic agent weights. The take-home message from our work is that changes to the quota matter, even when weights are nearly identical. Given the relative success of this analysis, it would be interesting to study other natural weight distributions (the case of i.i.d. weights is studied by Filmus *et al.* [2016]). Moreover, our results show that employing probabilistic approaches to cooperative games (beyond the case of WVG's) may be a useful research avenue. Other classes of cooperative games with a natural notion of resources are a promising direction for future work. For example, consider *threshold task games* [Chalkiadakis *et al.*, 2010]: these are natural generalizations of WVGs to a setting with more than one task: every task t has a quota $q(t)$ and a value $V(t)$; a value of a set S is the value of the best task it can complete. Assuming that agent weights are sampled from some distribution, it would be interesting to study the expected effects on agent values, based on the choice of tasks. Similarly, in *multicommodity network flow games* [Markakis and Saberi, 2005], each agent controls an edge in a directed graph; the value of a set $S \subseteq N$ is the maximal multicommodity flow it can transfer through a graph. Again, the effects of choices of commodities on agents' values are a promising direction for future study.

Acknowledgments We thank the IJCAI reviewers for their helpful comments. This research was partially conducted while Oren and Zick were at MSR; Zick was supported by A*STAR Singapore.

References

- [Aziz and Paterson, 2008] H. Aziz and M. Paterson. Computing voting power in easy weighted voting games. *CoRR*, abs/0811.2497, 2008.
- [Bachrach *et al.*, 2010] Y. Bachrach, E. Markakis, E. Resnick, A.D. Procaccia, J.S. Rosenschein, and A. Saberi. Approximating power indices: theoretical and empirical analysis. *AAMAS'10*, pages 105–122, 2010.
- [Berry, 1941] A. C. Berry. The accuracy of the gaussian approximation to the sum of independent variates. *Transactions of the American Mathematical Society*, 49(1):122–136, 1941.
- [Calvert, 1985] R.L. Calvert. Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence. *American Journal of Political Science*, pages 69–95, 1985.
- [Caragiannis *et al.*, 2013] I. Caragiannis, A.D. Procaccia, and N. Shah. When do noisy votes reveal the truth? In *EC'13*, pages 143–160, 2013.
- [Caragiannis *et al.*, 2014] I. Caragiannis, A.D. Procaccia, and N. Shah. Modal ranking: a uniquely robust voting rule. In *AAAI'14*, pages 616–622, 2014.
- [Chalkiadakis and Wooldridge, 2016] G. Chalkiadakis and M.J. Wooldridge. Weighted voting games. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 16. Cambridge University Press, 2016.
- [Chalkiadakis *et al.*, 2010] G. Chalkiadakis, E. Elkind, E. Markakis, M. Polukarov, and N.R. Jennings. Cooperative games with overlapping coalitions. *Journal of Artificial Intelligence Research*, 39(1):179–216, 2010.
- [Chalkiadakis *et al.*, 2011] G. Chalkiadakis, E. Elkind, and M.J. Wooldridge. *Computational Aspects of Cooperative Game Theory*. Morgan and Claypool, 2011.
- [Coughlin, 1992] P.J. Coughlin. *Probabilistic voting theory*. Cambridge University Press, 1992.
- [Enelow and Hinich, 1989] J.M. Enelow and M.J. Hinich. A general probabilistic spatial theory of elections. *Public Choice*, 61(2):101–113, 1989.
- [Fatima *et al.*, 2008] S.S. Fatima, M. Wooldridge, and N.R. Jennings. A linear approximation method for the shapley value. *Artificial Intelligence*, 172(14):1673–1699, 2008.
- [Filmus *et al.*, 2016] Y. Filmus, J. Oren, and K. Soundararajan. Shapley values in weighted voting games with random weights. *Submitted*, 2016.
- [Leech and Machover, 2003] D. Leech and M. Machover. Qualified majority voting: the effect of the quota. *LSE Research Online*, 2003.
- [Leech, 2002] D. Leech. Designing the voting system for the council of the european union. *Public Choice*, 113:437–464, 2002.
- [Lindner, 2004] I. Lindner. *Power measures in large weighted voting games*. PhD thesis, Univ. of Hamburg, 2004.
- [Lu and Boutilier, 2011] T. Lu and C. Boutilier. Vote elicitation with probabilistic preference models: Empirical estimation and cost tradeoffs. In *ADT'11*, pages 135–149, 2011.
- [Markakis and Saberi, 2005] E. Markakis and A. Saberi. On the core of the multicommodity flow game. *Decision support systems*, 39(1):3–10, 2005.
- [Mitzenmacher and Upfal, 2005] M. Mitzenmacher and E. Upfal. *Probability and computing - randomized algorithms and probabilistic analysis*. Cambridge University Press, 2005.
- [Mitzenmacher *et al.*, 2000] M. Mitzenmacher, A. W. Richa, and R. Sitaraman. The power of two random choices: A survey of techniques and results. In *Handbook of Randomized Computing*, pages 255–312. Kluwer, 2000.
- [Oren *et al.*, 2013] Joel Oren, Yuval Filmus, and Craig Boutilier. Efficient vote elicitation under candidate uncertainty. In *IJCAI'2013*, pages 309–316, 2013.
- [Oren *et al.*, 2014] J. Oren, Y. Filmus, Y. Zick, and Y. Bachrach. On the effects of priors in weighted voting games. In *COMSOC'14*, 2014.
- [Peleg and Sudhölter, 2007] B. Peleg and P. Sudhölter. *Introduction to the Theory of Cooperative Games*, volume 34. Springer, 2007.
- [Penrose, 1946] L.S. Penrose. The elementary statistics of majority voting. *Journal of the Royal Statistical Society*, 109(1):53–57, 1946.
- [Raab and Steger, 1998] M. Raab and A. Steger. “balls into bins” - a simple and tight analysis. In *RANDOM'98*, pages 159–170, 1998.
- [Shapley, 1953] L.S. Shapley. A value for n -person games. In *Contributions to the Theory of Games, vol. 2*, Annals of Mathematics Studies, no. 28, pages 307–317. Princeton University Press, Princeton, N. J., 1953.
- [Słomczyński and Życzkowski, 2006] W. Słomczyński and K. Życzkowski. Penrose voting system and optimal quota. *Acta Physica Polonica B*, 37(11):3133–3143, 2006.
- [Tauman and Jelnov, 2012] Y. Tauman and A. Jelnov. Voting power and proportional representation of voters. Dept. of Econs. Working Papers 12–04, Stony Brook Univ., 2012.
- [Zick *et al.*, 2011] Y. Zick, A. Skopalik, and E. Elkind. The shapley value as a function of the quota in weighted voting games. In *IJCAI'11*, pages 490–495, 2011.
- [Zick, 2013] Y. Zick. On random quotas and proportional representation in weighted voting games. In *IJCAI'13*, pages 432–438, 2013.
- [Zuckerman *et al.*, 2012] M. Zuckerman, P. Faliszewski, Y. Bachrach, and E. Elkind. Manipulating the quota in weighted voting games. *Artificial Intelligence*, 180–181:1–19, 2012.