The Mean Voter, the Median Voter, and Welfare-maximizing Voting Weights[†]

Nicola Maaser[*]
and
Stefan Napel[*,‡]

March 5, 2012

Abstract

Representatives from differently sized constituencies take political decisions by a weighted voting rule and adopt the ideal point of the weighted median amongst them. Preferences of each representative are supposed to coincide with the constituency's median voter. The paper investigates how each constituency's population size should be mapped to a voting weight for its delegate when the objective is to maximize the total expected utility generated by the collective decisions. Depending on the considered utility functions, this is equivalent to approximating the sample mean or median voter of the population by a weighted median of sub-sample medians. Monte Carlo simulations indicate that utilitarian welfare is maximized by a *square root rule* if the ideal points of voters are all independent and identically distributed. However, if citizens are risk-neutral and their preferences are sufficiently positively correlated within constituencies, i.e., if heterogeneity between constituencies dominates heterogeneity within, then a *linear rule* performs better.

Keywords: weighted voting, two-tier voting systems, square root rules, representative democracy, utilitarianism

1 Introduction

An important application of voting power analysis (see Felsenthal and Machover 1998 for a comprehensive overview) concerns the question of how voting weights should be assigned in *two-tier voting systems.* At the bottom tier, countries, states, districts, or other kinds of constituencies each elect a representative who will on their behalf cast a block vote in a top

[†]We are grateful for constructive discussion at the Leverhulme Trust's *Voting Power in Practice Symposium* 2011.

[*]Dept. of Economics, University of Bayreuth, Universitätsstr. 30, 95447 Bayreuth, Germany, and Public Choice Research Centre, Turku, Finland; phone: +49-921-552880; e-mail: nicola.maaser@uni-bayreuth.de, stefan.napel@uni-bayreuth.de

[‡]Corresponding author.

tier assembly or council. The Council of Ministers of the European Union (EU) is one of the most prominent examples of such a system, and much research on fair or optimal design of voting rules has been stimulated by successive EU enlargements. Other examples include the International Monetary Fund (see Leech and Leech 2009), the German Bundesrat and, with inessential qualifications, the US Electoral College.

In order to evaluate the design of two-tier voting systems a multitude of normative criteria can be brought to bear. Depending on the application at hand, desirable features include the responsiveness of collective decisions to individual preference changes, the capability to reach a decision, or equality of representation. From the perspective of mainstream economics, *utilitarian welfare* is a particularly prominent criterion (see Harsanyi 1955 and, e.g., Barberà and Jackson 2006). In particular, if the design of a two-tier voting system maximizes the total expected utility of the citizens, it is Pareto efficient: no other system can raise expected utility of some citizens without lowering it for others.

In this paper, we study the relationship between the allocation of block voting rights, i.e., the *voting weights* of constituency representatives, and the utilitarian welfare that is induced by the outcomes of a two-tier decision making process. We consider a model in which the feasible policy alternatives constitute a finite or infinite real interval. Voter preferences are assumed to be single-peaked, i.e., an individual's utility from a particular collective decision is strictly decreasing in distance to the respective voter's ideal point. These ideal points are conceived of as random variables with an identical continuous distribution for all citizens.

For ease of exposition, we suppose that each constituency comprises an odd number of voters. Then we assume, first, that the policy advocated by the single representative of any given constituency is congruent with the ideal point of the respective *constituency's median voter*. Second, the decision which is taken at the top tier is identified with the position of the *pivotal representative*. This representative is determined by the given allocation of voting weights and a 50% decision quota together with the policy positions of all delegates. It corresponds to the *weighted median amongst the delegates* and to the *core* of the spatial voting game in the assembly. Consideration of the respective, generically single-valued core provides a short-cut to the equilibrium outcome of various conceivable negotiation protocols, which might structure strategic bargaining at the council level.[1] As long as the weighted median of delegates who represent their constituencies' median voters is a reasonable approximation for the outcomes generated by the two-tier voting system, the actual processes of preference aggregation within the council and within the constituencies can remain unspecified. In particular, the latter could differ across constituencies.

We take a set of differently sized constituencies as given and seek to find the weight allocation rule which maximizes total expected utility. We presume that the preferences over policy outcomes have the same cardinal intensity across voters and distinguish between two utility specifications. Namely, each voter's cardinal utility function decreases either linearly or quadratically in the distance between the individual's ideal point and the collective policy outcome. The former specification corresponds to voters who are risk

[1]See, e.g., Cho and Duggan (2009).

neutral, i.e., who are indifferent between facing distances x and y to their ideal points with probabilities p and $1-p \in (0,1)$, or suffering the expected distance $px+(1-p)y$ for sure. The latter specification describes risk-averse individuals.

Standard results from statistics imply that the position of the *population's median voter* maximizes total expected utility for the linear specification, whereas the position of the *population's mean voter* maximizes it for the quadratic specification. It is, however, a non-trivial question how the best estimates for the sample median or sample mean, respectively, can be obtained by computing a weighted median of the medians of differently sized subsamples. We are not aware of – and have unfortunately neither been able to obtain – general analytical results on this issue. We, therefore, conduct extensive computer simulations.

The main finding of our Monte Carlo analysis is that a *square root rule* should be used in order to allocate voting weights if all citizens are a priori identical in a strong sense: namely, if their ideal points come from the same probability distribution and are *statistically independent* of each other. Note, however, that in this case there should be little objection to redrawing constituency boundaries. Obviously, the problem of maximizing total expected utility could then be readily resolved by creating constituencies of equal population size and giving each representative the same weight if the number of constituencies is fixed – or by creating an all-encompassing, single constituency if not. This observation motivates the consideration of citizens that are a priori identical in a weaker sense: their ideal points come from the same probability distribution but are *positively correlated* within the constituencies. For this scenario, a *degressively proportional rule* remains optimal for the quadratic utility specification, but the right degree of degressivity depends on the given vector of population sizes. And, importantly, total expected utility is maximized by a *linear rule* if voters are risk-neutral, i.e., if utility falls linearly in distance, and the degree of within-constituency similarity (or dissimilarity between constituencies) is sufficiently high.

The design of welfare-maximizing voting rules for two-tier systems of representative democracy has received formal mathematical consideration only quite recently.[2] Barberà and Jackson (2006) study the design of efficient voting rules in a fairly general setup for binary decisions. They derive a square root allocation rule for the so-called "fixed-size-of-blocks model", which assumes a great degree of independence between the preferences of members of the same constituency. By contrast, they show a directly proportional allocation of weights to be optimal in their "fixed-number-of-blocks model", which reflects strong preference alignments between individuals within the same constituency (and independence across constituencies). These results are corroborated by Beisbart and Bovens (2007). Closely related, Beisbart et al. (2005) evaluate total expected utility under dif-

[2]Historically, most attention has been devoted to giving each citizen an equally effective voice in elections (cf. *Reynolds v. Sims*, 377 U.S. 533, 1964). In two-tier voting systems, this calls for an a priori equal chance of each voter to indirectly determine the policy outcome. For binary policy spaces, Penrose (1946) has shown that individual powers are approximately equalized if voting weights of the representatives are chosen such that their Penrose-Banzhaf voting powers (Penrose 1946; Banzhaf 1965) are proportional to the square root of the corresponding population sizes. An extension to convex policy spaces is provided by Maaser and Napel (2007) and Kurz et al. (2011).

ferent decision rules for the Council of Ministers of the European Union and the premise that proposals always affect all individuals from a given country identically. Koriyama and Laslier (2011) argue in great generality that a utilitarian ideal requires vote allocation rules to be degressively proportional.

The considered objective of maximum total utility is intimately linked with achieving congruence between individual preferences and the collective policy. For binary decisions that are taken by the citizens directly (corresponding to the degenerate case of singleton constituencies and uniform weights), Rae (1969) has shown that the probability that the average citizen "has his way" (i.e., is in agreement with the voting outcome) is maximized by 50% majority rule.[3] But the outcome of indirect, two-tier decision processes can easily deviate from that of direct democracy: even under simple majority rule it is possible that the alternative adopted by the body of representatives is supported only by a minority of all citizens.

The degree of majoritarianism of a two-tier system decreases in the expected difference between the size of the popular majority camp and the number of citizens in favor of the assembly's decision. Felsenthal and Machover (1998, pp. 63–78; 1999) study this so-called *mean majority deficit* in a binary voting model. They find it to be minimal under a square root allocation of voting weights.[4] As shown by Felsenthal and Machover, minimization of the mean majority deficit can also be interpreted in a somewhat utilitarian vein, namely as maximizing the sum of citizens' indirect voting power as measured by the non-normalized Penrose-Banzhaf index. Kirsch (2007) considers optimal weights for a related notion of majoritarian deficit. Similarly, Feix et al. (2008) investigate the *probability* of situations where the decision taken by the representatives and a hypothetical referendum decision diverge.[5] All these investigations consider the case of binary alternatives. Moving to richer policy spaces, Maaser and Napel (2012) analyze the expected discrepancy between a two-tier and a direct-democratic single-tier system in a one-dimensional spatial voting model.

A dichotomous pattern has emerged from this literature: rules that relate voting weights to the square root of population sizes have been found to be optimal under various objective functions if citizens are assumed to be homogeneous in the sense of having *independent and identically distributed* (i.i.d.) preferences. But square root rules cease to be optimal, and often a linear rule replaces them, if dependence of some sort or another is introduced. Investigations that highlight the critical role played by the degree of similarity within constituencies as opposed to that between constituencies include Gelman et al. (2002), Barberà and Jackson (2006), Kirsch (2007), Beisbart and Bovens (2007), Feix et al. (2008), Kaniovski (2008), and Maaser and Napel (2012).[6] For example, extending the main result of Felsenthal and Machover (1999) from $\{0,1\}$-choices to the convex policy space $[0,1]$,

[3]Dubey and Shapley (1979) provide a generalization of this result to the domain of all simple games.

[4]Felsenthal and Machover refer to this allocation rule as the *second square root rule* in order to distinguish it from Penrose's (1946) (first) square root rule, which requires representatives' *voting powers* – rather than their weights – to be proportional to the square roots of their constituencies' population sizes.

[5]This situation is known in the social choice literature as a *referendum paradox* (see, e.g., Nurmi 1998).

[6]Also see Felsenthal and Machover (1998, pp. 70ff).

Maaser and Napel (2012) find that the *direct democracy deficit* is minimized when voting weights are allocated to representatives in proportion to the square root of constituency population sizes if ideal points are i.i.d. However, if sufficiently strong positive correlation of preferences within each constituency is introduced, then the best weight allocation rule is linear instead.

2 Model

We will consider a different objective function here than in Maaser and Napel (2012). But the baseline model of two-tier decision making is the same in both papers. The following description and overlapping parts of the analysis will draw directly on the presentation in Maaser and Napel (2012).

Consider the partition $\mathfrak{C} = \{\mathcal{C}_1, \dots, \mathcal{C}_r\}$ of a large voter population into r *constituencies* with $n_j = |\mathcal{C}_j| > 0$ members each. Let $n \equiv \sum_j n_j$ and all n_j be odd numbers for simplicity. The preferences of any voter $i \in \{1, \dots, n\} = \bigcup_j \mathcal{C}_j$ are assumed to be single-peaked with *ideal point* ν^i in a convex one-dimensional *policy space* $X \subset \mathbb{R}$, i.e., a finite or infinite interval. These ideal points are conceived of as realizations of random variables with an identical continuous a priori distribution; any given profile $(\nu^1, \dots, \nu^n)$ of ideal points is interpreted as reflecting voter preferences on a specific one-dimensional policy issue (a tax level, expenditure on a public good, extent of redistribution, boldness of pension reform, etc.).

A collective decision $x \in X$ on the issue at hand is taken by an assembly or council of representatives $\mathcal{R}$ which consists of one representative from each constituency. Without going into details, we assume that the preferences of $\mathcal{C}_j$'s representative are congruent with its median voter, i.e., representative j has ideal point

$$\lambda_j = \text{median}\{\nu^i \colon i \in \mathcal{C}_j\}.$$

This is clearly an idealizing abstraction because political agents can often exploit informational asymmetries in order to pursue their own rather than their principal's preferences (e.g., concerning their privileges – see Gerber and Lewis 2004 for empirical evidence on the effect of constituency heterogeneity on the alignment between representative and median voter).

In the top-tier assembly $\mathcal{R}$, each constituency $\mathcal{C}_j$ has *voting weight* $w_j \geq 0$. Any subset $S \subseteq \{1, \dots, r\}$ of representatives which achieves a combined weight $\sum_{j \in S} w_j$ above $q \equiv 0.5 \sum_{j=1}^{r} w_j$, i.e., comprises a *simple majority* of total weight, can implement a policy $x \in X$. So collective decisions are taken according to the weighted voting rule $[q; w_1, \dots, w_r]$.

Let $\lambda_{k:r}$ denote the k-th leftmost ideal point amongst the representatives (i. e., the k-th order statistic of $\lambda_1, \dots, \lambda_r$) and consider the random variable P defined by

$$P = \min\Big\{ l \in \{1, \dots, r\} \colon \sum_{k=1}^{l} w_{k:r} > q \Big\}.$$

For a generic weight vector $(w_1, \ldots, w_r)$, representative $P : r$'s ideal point, $\lambda_{P:r}$, is the unique policy that beats any alternative $x \in X$ in a pairwise vote, i. e., it constitutes the core of the voting game in $\mathcal{R}$ with weights $w_1, \ldots, w_r$ and a 50% quota. Without any formal analysis of decision procedures that might be applied in $\mathcal{R}$ (see Banks and Duggan 2000, or Cho and Duggan 2009), we assume that the policy agreed in the council coincides with the ideal point of *pivotal representative* $P\!:\!r$. In summary, the policy outcome produced by the two-tiered voting system is

$$x_{\mathcal{R}} = \lambda_{P:r}\,.$$

For ideal point profile $(\nu^1, \ldots, \nu^n)$ the total utility that the society receives from $x_{\mathcal{R}}$ is

$$\bar{U} = \sum_{i=1}^{n} -|\nu^i - x_{\mathcal{R}}|, \quad \text{or} \tag{1a}$$

$$\hat{U} = \sum_{i=1}^{n} -(\nu^i - x_{\mathcal{R}})^2 \tag{1b}$$

if for each voter utility decreases (a) linearly or (b) quadratically in the distance between his ideal point and the outcome.

Taking partition $\mathfrak{C}$ as given we would like to answer the following question: *Which allocation of voting weights maximizes the total expected utility of the two-tier voting system?* Or, more formally, we search for weight allocation rules W which approximately solve the problems

$$\max_{w_1,\ldots,w_r} \mathbf{E}\left[\bar{U}\right], \quad \text{and} \tag{2a}$$

$$\max_{w_1,\ldots,w_r} \mathbf{E}\left[\hat{U}\right], \tag{2b}$$

respectively, where by an "allocation rule" we mean a simple mapping W which assigns weights $(w_1, \ldots, w_r) = W(\mathcal{C}_1, \ldots, \mathcal{C}_r)$ to any given partition of a large population. Our criterion for acceptably "simple" mappings $W\colon \mathfrak{C} \mapsto (w_1, \ldots, w_r)$ will be that they are *power laws*, i.e., $w_j = n_j^\alpha$ for some constant $\alpha \in [0, 1]$. This class of mappings nests the square root and linear rules which have played prominent roles in the previous literature.[7]

3 Analysis

Under the model's assumptions, it can be shown that societal welfare $\bar{U}$ would be maximized if, for any realization of voter preferences, we had

$$x_{\mathcal{R}} = \text{median}\{\nu^1, \ldots, \nu^n\}.$$

[7] To be precise, Penrose's square root rule is nested only asymptotically, namely when $\mathfrak{C}$ involves a great number r of constituencies with a non-pathological size distribution. See Lindner and Machover (2004) and Chang et al. (2006) on the vanishing difference between voting weights and voting powers as $r \to \infty$.

That is, it is the unconstrained ideal to choose the preferred policy of the median individual in the union (see, e.g., Schwertman et al. 1990). This policy outcome would also be brought about by frictionless collective decision-making in a full assembly of all citizens under simple majority rule since it beats every alternative policy in a pairwise vote. Because of the loss of information that results from only aggregating the votes of the top-tier representatives, however, $x_{\mathcal{R}}$ does generally *not* coincide with the median ideal point in the population. Problem (2a) is thus equivalent to that of minimizing the expected value of $|x_{\mathcal{R}} - \text{median}\{\nu^1, \dots, \nu^n\}|$, which is referred to as the *direct democracy deficit* in Maaser and Napel (2012).[8]

While the median has the property of minimizing the sum of absolute distances, the sum of squared distances is minimized by the mean (see, e.g., Cramér 1946, Sect. 15.4). Thus, the ideal non-voting solution to problem (2b) would be to always implement the policy that corresponds to the *mean* of ideal points $\{\nu^1, \dots, \nu^n\}$. Our maximization problem can, therefore, be reframed in the case of quadratic utility functions as follows: by which simple weight allocation rule do we achieve a particularly "small" expected distance between $x_{\mathcal{R}}$ and the mean voter position? In principle, an estimate of the overall mean could be obtained by taking the n_j-weighted mean of $\lambda_1, \dots, \lambda_r$. If, however, representatives' positions are aggregated by voting under strategic interaction rather than being averaged (e.g., by a bureaucrat) then the outcome $x_{\mathcal{R}}$ at the top-tier will match one of the representatives' positions in the considered spatial voting model, namely their n_j^α-*weighted median* in our model. This will usually differ from the n_j-*weighted mean.* Optimal statistical aggregation *by averaging* does not really help in solving the problem of optimal aggregation *by voting.*

If the ideal points of voters $i \in \mathcal{C}_j$ are pairwise independent and come from an arbitrary identical distribution F with positive density f on X, then its median position λ_j asymptotically has a normal distribution with mean $\mu = F^{-1}(0.5)$ and standard deviation

$$\sigma_j = \frac{1}{2\, f(\mu)\sqrt{n_j}} \tag{3}$$

(see, e.g., Arnold et al. 1992, p. 223). The variance of the position of $\mathcal{C}_j$'s representative is the smaller, the greater the population size n_j.

This implies that even in the seemingly trivial case of uniform weights $w_1 = \ldots = w_r$, the top-tier decision $x_{\mathcal{R}} \in X$ has a rather non-trivial distribution when constituency sizes differ. Namely, $x_{\mathcal{R}}$ is then an order statistic of *differently* distributed random variables, for which relatively few limit results are known. For non-identical weights $w_1, \dots, w_r$, $x_{\mathcal{R}}$ is a combinatorial function of such order statistics. Therefore, it seems extremely hard – at least to us – to obtain or approximate solutions to (2a) and (2b) analytically. We will now briefly look at two special cases in order to develop some intuition, and then turn to computer simulations in Section 4.

[8]Note that even though total utility from the decisions which result from the considered two-tier process typically falls short of the global maximum achieved under a direct democracy, representative democracy has a number of advantages. These presumably also generate utility for citizens which is not considered in our model.

First, consider the trivial case of equipopulous constituencies. Any uniform weight allocation $w_1 = \ldots = w_r > 0$ then maximizes total expected utility and the optimal value of α remains undetermined. Under identical weights, the pivotal representative's ideal point is the (unweighted) median of $\lambda_1, \ldots, \lambda_r$. How close this comes to the population's sample median and mean, respectively, will depend on the number of symmetric constituencies in the partition.[9]

So far no assumptions have been made regarding how voter ideal points are jointly distributed. It can be convincingly argued that – for the kind of constitutional design problem that we are dealing with – specific knowledge about individual preferences should be ignored. From behind the constitutional "veil of ignorance" all citizens should be considered *identical* a priori. This corresponds to drawing every ideal point ν^i from the *same* marginal probability distribution F. However, such a constitutional a priori perspective does not necessarily entail that preferences of citizens must also be conceived of as *independent* of each other. It is true that the i.i.d. assumption for all ideal points ν^i with $i \in \bigcup_j \mathcal{C}_j$, i.e., consideration of the product distribution F^n, is a particularly compelling benchmark. Still, the partition $\mathfrak{C}$ may have reasons that need to be acknowledged behind the "veil of ignorance" (e.g., geographic barriers, ethnics, language, or religion). For these reasons voter preferences are likely to be more closely connected within constituencies than across them.

As a second case of interest, suppose that $\nu^i = \nu^h$ whenever $i, h \in \mathcal{C}_j$. This carries the notion that citizens have on average closer links with each other within constituencies than across constituencies to its extreme. Problem (2a) has a clear-cut solution in this situation: $\mathbf{E}[\bar{U}]$ is maximal if the linear weight allocation rule $w_j = n_j$ for $j = 1, \ldots, r$, i.e., $\alpha^* = 1$, is used. Perfect correlation within constituencies implies that the ordered ideal points of all citizens $i = 1, \ldots, n$,

$$\nu^{1:n} \leq \nu^{2:n} \leq \nu^{3:n} \leq \ldots \leq \nu^{n-1:n} \leq \nu^{n:n},$$

can be written as

$$\underbrace{\lambda_{1:r} = \ldots = \lambda_{1:r}}_{n_{1:r} \text{ times}} \leq \underbrace{\lambda_{2:r} = \ldots = \lambda_{2:r}}_{n_{2:r} \text{ times}} \leq \ldots \leq \underbrace{\lambda_{r:r} = \ldots = \lambda_{r:r}}_{n_{r:r} \text{ times}}.$$

Thus, weights proportional to population sizes make representative j pivotal in $\mathcal{R}$ if and only if his policy position (and thus that of all $\mathcal{C}_j$-citizens) is also the population median. In the non-degenerate case of high but not perfect correlation within constituencies this optimality of proportional weights can be expected to apply *approximately*. The simulations reported in Section 4 indeed confirm this intuition: with linear individual utility functions total expected utility is maximized by an essentially linear rule provided that the ideal points of the citizens vary noticeably more across than within constituencies.

The above extreme case is also instructive to appreciate that a linear rule cannot be optimal in general when individual utility decreases *quadratically* in the distance between

[9] See Beisbart and Bovens (2011) for a related investigation in a binary voting model. They ask the worst-case question: which *number* of equipopulous districts *maximizes* the mean majority deficit?

ideal point and outcome. When constituencies differ in population size, the overall frequency distribution of policy positions will typically not be symmetric. It will be skewed to the right if a majority of the large constituencies prefers a policy to the left of the center (see Figure 1), and it is skewed to the left if the large constituencies have ideal points on the right. But then the population median (which would result from $\alpha = 1$) does not provide a good estimate of the population mean: the sample median is necessarily located to the left of the sample mean if the distribution is skewed to the right (just like average income is larger than median income if there are many small incomes and a few very large ones). A value of $\alpha < 1$ then produces a smaller deviation between the pivotal representative's ideal point and the mean ideal point. Figure 1 illustrates this by taking EU27 members as an example. In the figure all citizens within each constituency have identical policy preferences drawn from a uniform distribution on $[-1, 1]$. For the depicted right-skewed realization $(\nu_1, \ldots, \nu_n)$ of ideal points, $\alpha = 0.71$ is best among the considered parameters $\alpha = \{0, 0.01, \ldots, 1\}$: the associated outcome $x_{\mathcal{R}}$ is as close as possible to the mean of all ideal points. The same degressivity parameter $\alpha = 0.71$ would be optimal for the ideal point realization $(\nu'_1, \ldots, \nu'_n)$ with $\nu'_i = -\nu_i$ for all $i = 1, \ldots, n$, which is skewed to the left. For realizations that give rise to an essentially symmetric frequency distribution, $\alpha = 0.71$ performs as well as any alternative value (such as $\alpha = 1$). We can, therefore, conclude that $\alpha = 1$ must be suboptimal if one averages over all possible frequency distributions, i.e., considers the expected value $\mathbf{E}[\hat{U}]$. The *optimal* value of α depends on the constituency configuration at hand as well as on the theoretical distribution of individual ideal points.[10] It might but need not be close to 0.5.

When we consider non-degenerate degrees of correlation between the ideal points within a given constituency, it is even more difficult to come up with a clear intuition for what the best degree of degressivity should be. In the benchmark case of ideal points that are all pairwise independent and drawn from the same symmetric distribution, computation of the n_j-weighted mean of $\lambda_1, \ldots, \lambda_r$ would be the theoretically best way to estimate both the location of the sample median and the sample mean. The n_j-weighted mean is sensitive to outliers amongst the representatives' ideal points. This rules out optimality of $\alpha = 0$ because uniform weights select the median representative's ideal point and hence disregard any information about outliers. But a too great value of α would enable representatives from large constituencies to implement their preferred policy even if they happen to be outliers. It is not obvious at the outset what "too great" means and which α strikes the right balance.

An admittedly crude intuitive argument in favor of $\alpha = 0.5$ runs as follows. First consider the linear utility specification, so that the theoretical ideal is to approximate the population's *median voter* as well as possible. If all voter ideal points ν^i are i.i.d. then each individual $i = 1, \ldots, n$ a priori has probability $1/n$ to be the population median. The latter is hence located in constituency $\mathcal{C}_j$ with probability n_j/n. This makes weights

[10] The problem of finding the optimal value of α bears some resemblance to choosing an appropriate power-law transformation in order to improve the symmetry of a skewed empirical distribution (see, e.g., Yeo and Johnson 2000).

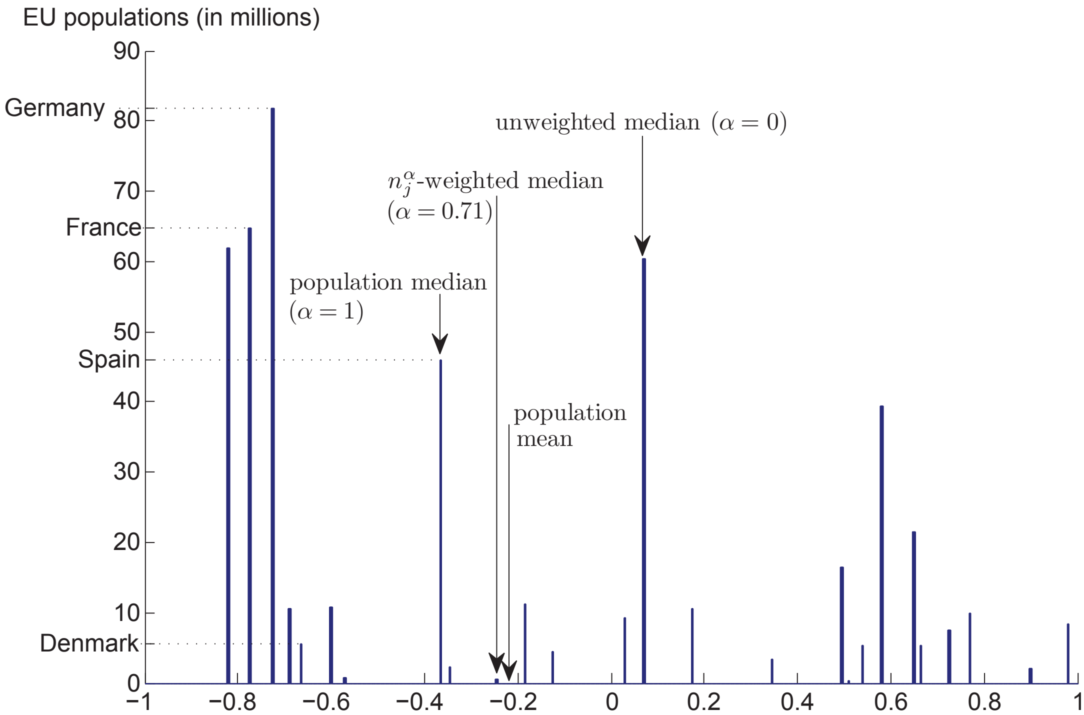

Figure 1: Sample frequency distribution of policy positions for EU27 member countries (2010 Eurostat population data).

which induce top-tier pivot probabilities proportional to the respective population sizes a particularly reasonable starting point. As Kurz et al. (2011) have shown, proportionality between the probability of the event $\{j = P\!:\!r\}$ and n_j can be achieved approximately by selecting weights w_j that are proportional to the *square root* of n_j in the i.i.d. case.

Second, for the quadratic utility specification, the goal is to approximate the population's *mean voter* by selecting a particular weighted median of the representatives. The mean voter is a virtual one who does not belong to any particular constituency. Notably, the mean ideal point will almost surely differ from those of *all* voters $i \in \{1, \ldots, n\}$ when the ideal point distribution has a density. Thus, the intuition provided for a square root rule in the case of linear utility does not apply directly to the case of quadratic utility functions. However, for the symmetric ideal point distributions which we focus on in this paper, the population mean and median will be very close to each other if all voters are pairwise independent. One may conjecture, therefore, that $\alpha = 0.5$ will work well under an i.i.d. assumption irrespective of the utility specification.

4 Simulations

Since we are unable to obtain more precise analytical insights – let alone any useful approximation of $\mathbf{E}[\bar{U}]$ or $\mathbf{E}[\hat{U}]$ as a function of α – we apply the *Monte-Carlo method.* It exploits that the empirical average of s independent realizations of $\bar{U} = \sum_{i=1}^{n} -|\nu^i - x_{\mathcal{R}}|$ and $\hat{U} = \sum_{i=1}^{n} -(\nu^i - x_{\mathcal{R}})^2$ converges to $\mathbf{E}[\bar{U}]$ and $\mathbf{E}[\hat{U}]$, respectively, as $s \to \infty$ by the law of large numbers.

In order to obtain realizations of $\bar{U}$ and $\hat{U}$ for the case of i.i.d. voter ideal points, we first draw n (pseudo-)random numbers from a given distribution F, giving rise to a list $\mathbf{v} = (\nu^1, \ldots, \nu^n)$.[11] Second, $\mathbf{v}$ is sorted within consecutive blocks of size $n_1, n_2, \ldots, n_r$ in order to obtain the corresponding realizations of the constituency medians $\lambda_1, \lambda_2, \ldots, \lambda_r$. We then infer the weighted median of these, using weights $w_j = n_j^{\alpha}$ for values of α which range from 0 to 1 in steps of 0.01, and thus obtain $x_{\mathcal{R}}$ for each value of α. The resulting values of $\bar{U}$ and $\hat{U}$ are recorded, and the procedure is repeated for one million iterations. Finally, we determine the values of α, denoted by $\bar{\alpha}^*$ and $\hat{\alpha}^*$ which produced the largest average total utility $\bar{U}$ and $\hat{U}$, respectively.

In our simulations we typically consider sets of $r = 25$ constituencies. Experience suggests that simulation results then do no longer exhibit strong dependence on the combinatorial peculiarities of the configuration at hand (this would be the case for significantly smaller numbers of constituencies). Most of the considered population configurations are artificial: sizes $n_1, \ldots, n_r$ are obtained by drawing random numbers from a specified distribution. The entry $\mathbf{U}(10^3, 3 \cdot 10^3)$ in Table 1, for instance, indicates that realizations of $\mathfrak{C}$

[11]Since the considered number of voters in each constituency $\mathcal{C}_j$ is large ($n_j \gg 50$), the respective population and constituency medians will approximately have normal distributions irrespective of the specific F which one considers. For the sake of completeness, let it still be mentioned that individual ideal points were drawn from a standard uniform distribution $\mathbf{U}(0, 1)$ in our simulations. The MATLAB source code is available upon e-mail request.

(a)

	$\bar{\alpha}^*$				
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
$\mathbf{U}(1000, 3000)$	**0.52**	**0.51**	**0.51**	**0.51**	**0.53**
$\mathbf{N}(2000, 200)$	**0.43**	**0.62**	**0.65**	**0.57**	**0.44**
$\mathbf{N}(2000, 400)$	**0.48**	**0.52**	**0.50**	**0.47**	**0.51**
$\mathbf{P}(1.0, 200)$	**0.52**	**0.52**	**0.52**	**0.53**	**0.52**

(b)

	$\hat{\alpha}^*$				
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
$\mathbf{U}(1000, 3000)$	**0.51**	**0.51**	**0.48**	**0.51**	**0.53**
$\mathbf{N}(2000, 200)$	**0.49**	**0.62**	**0.65**	**0.57**	**0.44**
$\mathbf{N}(2000, 400)$	**0.52**	**0.52**	**0.50**	**0.46**	**0.53**
$\mathbf{P}(1.0, 200)$	**0.52**	**0.52**	**0.51**	**0.53**	**0.52**

Table 1: Welfare-maximal α for i.i.d. voters and (a) linear utility or (b) quadratic utility

are considered for which each constituency size between 1 000 and 3 000 voters had *uniform* probability. Besides the uniform distribution, also truncated *normal* distributions $\mathbf{N}(\mu, \sigma)$ and *Pareto* distributions $\mathbf{P}(\kappa, \theta)$ with skewness parameter κ and threshold parameter θ have been employed in order to generate population configurations. For each "distribution type" of the population configuration, five independent realizations of $n_1, \ldots, n_r$ have been investigated. So Table 1 reports the respective optimal values $\bar{\alpha}^*$ (linear utility) and $\hat{\alpha}^*$ (quadratic utility) for altogether 20 different configurations.

The 95%-confidence intervals around the empirical mean of $\bar{U}$ and $\hat{U}$ are typically too wide to rule out that a neighbor of the reported best value of α produces a higher level of welfare. However, differences are significant when sufficiently distinct values like $\alpha = 0.5$ and $\alpha = 1$ are compared.[12] The obtained estimates of $\mathbf{E}[\Delta]$ are in most cases unimodal functions of α, i.e., increasing on $[0, \alpha^*)$ and decreasing on $(\alpha^*, 1]$. Overall, results in Table 1 are suggesting strongly that a square root allocation rule is close to being optimal (within the class of elementary power laws) if the ideal points of all voters are independent and identically distributed.

Concerning cases in which the ideal points of citizens are *not* independent and identically distributed, we focus on a special type of positive correlation within constituencies. In particular, we determine individual ideal points ν^i by a two-step random experiment:

[12] In particular, variation in population sizes $n_j \sim \mathbf{N}(2000, 200)$ is rather small. This results in an objective function that is essentially flat for a large range of values of α.

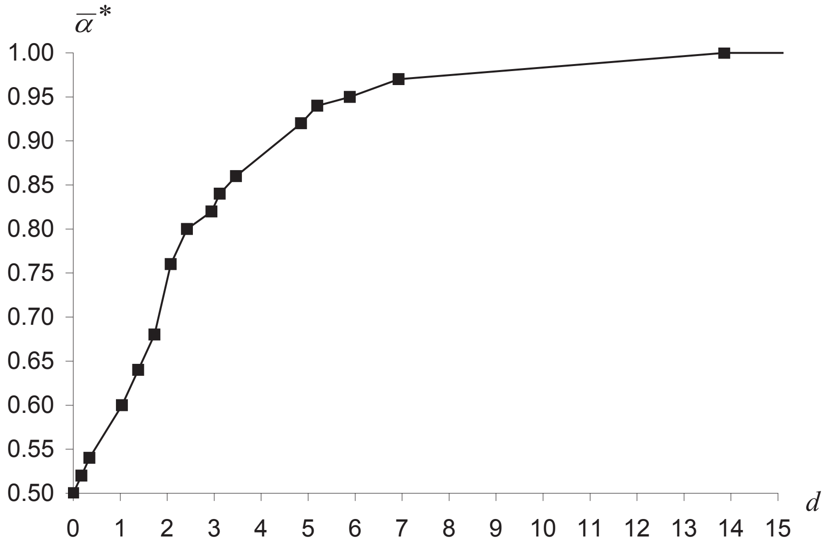

Figure 2: Welfare-maximal α for EU27 and linear utility as dissimilarity ratio d is varied

first, we draw a constituency-specific parameter μ_j independently for each $j = 1, \ldots, r$ from an identical distribution G with standard deviation σ_{ext}. Parameter σ_{ext} captures the degree of *external heterogeneity* between $\mathcal{C}_1, \ldots, \mathcal{C}_r$ for the policy issue at hand. The realization of parameter μ_j is taken to reflect the expected ideal point of citizens from $\mathcal{C}_j$ on a given policy issue. Each citizen $i \in \mathcal{C}_j$ is then assigned an individual ideal point ν^i from a distribution F_{μ_j} which has mean μ_j and is otherwise just a shifted version of some distribution $F \equiv F_0$ for each constituency $j = 1, \ldots, r$.[13] F's standard deviation σ_{int} is a measure of the *internal heterogeneity* in any constituency. It reflects opinion differences within any given $\mathcal{C}_j$. In summary, our second set of simulations has taken the ideal points of all citizens to be *identically distributed* with convolved a priori distribution $G * F$, but to involve *dependencies*: citizens in constituency $\mathcal{C}_j$ all experience the same shift μ_j, which is independent of μ_k for any $k \neq j$.

The ratio $\sigma_{ext}/\sigma_{int} =: d$ between external and internal heterogeneity provides a measure of the degree to which citizens are more similar within than between constituencies or, loosely speaking, the *preference dissimilarity* of the constituencies. In the i.i.d. case *no* dissimilarity exists between different constituencies, i.e., results in Table 1 are based on $d = 0$. Table 2 reports optimal values $\bar{\alpha}^*$ and $\hat{\alpha}^*$ for the same configurations as in Table 1 and two positive dissimilarity levels, namely $d = 8$ and the degenerate case of infinite dissimilarity ($\sigma_{int} = 0$). While results for i.i.d. ideal points did not significantly differ between the linear specification of individual utility functions and the quadratic one in Table 1, this is no longer the case when significant preference correlations exist.

The optimality of $\bar{\alpha}^* = 1$ as $d \to \infty$ for the linear specification has already been explained in our theoretical discussion in Section 3 (considering fixed $\sigma_{ext} > 0$ and $\sigma_{int} \to 0$).

[13]Specifically, we draw μ_j from a uniform distribution $\mathbf{U}(-a, a)$ with variance σ^2_{ext}, and then obtain $\nu^i = \mu_j + \varepsilon$ with $\varepsilon \sim \mathbf{U}(0, 1)$.

(a)

		$\bar{\alpha}^*$				
		(1)	**(2)**	**(3)**	**(4)**	**(5)**
$\mathbf{U}(1000, 3000)$	$d = 8$	**0.96**	**0.96**	**0.96**	**0.97**	**0.96**
	$d = \infty$	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**
$\mathbf{N}(2000, 200)$	$d = 8$	**0.96**	**0.95**	**0.97**	**0.96**	**0.97**
	$d = \infty$	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**
$\mathbf{N}(2000, 400)$	$d = 8$	**0.96**	**0.96**	**0.95**	**0.97**	**0.95**
	$d = \infty$	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**
$\mathbf{P}(1.0, 200)$	$d = 8$	**0.97**	**0.96**	**0.96**	**0.97**	**0.97**
	$d = \infty$	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**

(b)

		$\hat{\alpha}^*$				
		(1)	**(2)**	**(3)**	**(4)**	**(5)**
$\mathbf{U}(1000, 3000)$	$d = 8$	**0.48**	**0.49**	**0.52**	**0.53**	**0.48**
	$d = \infty$	**0.51**	**0.49**	**0.54**	**0.55**	**0.50**
$\mathbf{N}(2000, 200)$	$d = 8$	**0.47**	**0.51**	**0.65**	**0.53**	**0.62**
	$d = \infty$	**0.53**	**0.54**	**0.56**	**0.61**	**0.51**
$\mathbf{N}(2000, 400)$	$d = 8$	**0.50**	**0.54**	**0.47**	**0.54**	**0.52**
	$d = \infty$	**0.49**	**0.51**	**0.51**	**0.51**	**0.47**
$\mathbf{P}(1.0, 200)$	$d = 8$	**0.64**	**0.56**	**0.55**	**0.68**	**0.66**
	$d = \infty$	**0.64**	**0.56**	**0.55**	**0.68**	**0.66**

Table 2: Welfare-maximal α for two different preference dissimilarity ratios d and (a) linear utility or (b) quadratic utility

The findings reported in Table 2(a) indicate that this result extends in close approximation to more moderate levels of dissimilarity such as $d \geq 8$. Figure 2 demonstrates that a situation in which nearly linear voting weight allocations maximize $\mathbf{E}[\bar{U}]$ arises quickly as the preference dissimilarity which underlies the policy ideals of representatives in $\mathcal{R}$ increases. The figure considers $r = 27$ and a population configuration based on recent Eurostat data for members of the European Union.[14] The EU Council of Ministers is the predominant example of a two-tier voting system because its members officially represent national governments and, eventually, the citizenries of the member states. Note, however, that the current weighted voting rules for the Council, and also its future ones as codified in the Treaty of Lisbon, involve supermajority requirements in multiple dimensions, while Figure 2 is based on the assumption of a 50% decision quota. We leave an investigation of the effect of supermajority rules on the maximizer (and maximum) of utilitarian welfare in our spatial voting framework to future research.

The optimal levels of $\hat{\alpha}^*$ for a quadratic utility specification, displayed in Table 2(b), fail to show convergence to any specific rule as $d \to \infty$. In particular, it does not seem to make a significant difference whether dissimilarity is moderate or extreme. Moreover, the reported values of $\hat{\alpha}^*$ do not differ noticeably from their i.i.d. counterparts in Table 1 except for Pareto-distributed population configurations (where constituency sizes have a skewed distribution).

We argued in our discussion of Figure 1 that $\alpha = 1$ should not be expected to be optimal when individual utility functions are quadratic and $d \to \infty$, and that it is not clear which particular $\hat{\alpha}$ should be optimal. Table 2 suggests vaguely that a square root allocation might actually do best when constituency sizes are drawn from a symmetric distribution (uniform or normal). But certainly more weight needs to be given to large constituencies than under a square root law when the population distribution is skewed (Pareto).

Note that even if the distribution of *population sizes* $n_1, \ldots, n_r$ is symmetric, the realized frequency distributions of *ideal points* will be skewed more often than not. For instance, a frequency distribution like the one displayed in Figure 1 will still be common even if we have constituency sizes that range equidistantly from some smallest value $\underline{n}$ to a largest value $\overline{n}$ (mimicking a uniform distribution on $[\underline{n}, \overline{n}]$). So some degressively proportional weighting scheme raises total expected utility relative to a linear rule. We conjecture that, for symmetric distributions of constituency sizes, the average distance between the sample median and the sample mean is larger, the larger the variance of $n_1, \ldots, n_r$. Therefore, the greater the variance of $n_1, \ldots, n_r$, the smaller the optimal value $\hat{\alpha}^*$. This hypothesis is supported by our simulation data. In particular, Figure 3 displays the welfare-maximizing level $\hat{\alpha}^*$ in case of the quadratic utility specification and degenerate preference dissimilarity ($d = \infty$) for altogether 52 distinct population configurations that were drawn either from uniform and (truncated) normal distributions with $r = 25$ or $r = 35$. A higher standard deviation s of the population sizes $n_1, \ldots, n_r$ visibly translates into a smaller optimal value

[14] We have used 2010 population data measured in 1 000 individuals for computational reasons. This corresponds with the "block model" in Barberà and Jackson (2006), which supposes that a constituency can be subdivided into equally sized "blocks" whose members have perfectly correlated preferences within blocks, but are independent across blocks.

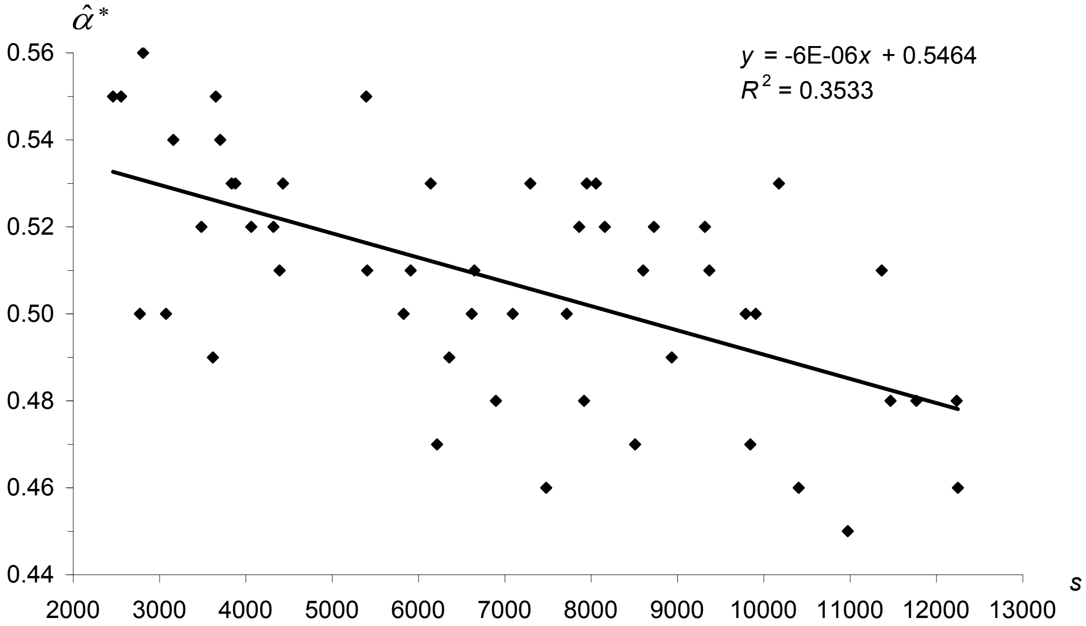

Figure 3: Optimal $\hat{\alpha}$ for 52 population configurations with sample standard deviations s.

$\hat{\alpha}^*$. The slope of the corresponding regression line is not very steep, but it is significantly different from zero. Still, a value of $\alpha = 0.5$ is never very far off. This is in line with the findings in Table 2 for a symmetric distribution of constituency sizes.

5 Concluding remarks

The findings of our investigation of utilitarian welfare or total expected utility of the citizens in a spatial voting model might be summarized – cum grano salis – as supporting the conclusions of the related literature on binary voting models (see Section 1). In particular, *if* the preferences of the voters are characterized by independent and identically distributed (i.i.d.) ideal points over a one-dimensional policy space, then using a *square root rule* for allocating voting weights performs best. This is irrespective of whether voters' utility decreases linearly or quadratically in distance from their policy ideal (corresponding to risk neutrality or a particular extent of risk aversion when facing uncertain collective decisions). Unfortunately, we could provide but a vague intuition for *why* a square root law obtains.

Our findings are also consistent with the binary voting literature in that optimality of a square root rule – be it elementary like $w_j = n_j^{0.5}$ or sophisticated like the seminal suggestion by Penrose (1946) – tends *not* to extend to situations in which the i.i.d. assumption is violated. It has increasingly come to be understood that when voters have a priori identical random preferences in the binary case or on some richer space, like the one considered here, and these preferences exhibit *positive correlation* within constituencies, then there is a potentially very rapid phase transition from $\alpha = 0.5$ to $\alpha = 1$ performing best.

However, the results shown in Table 2(b) and Figure 3 cast some doubt on this dichotomy between using a square root rule for similar constituencies and a linear rule for sufficiently dissimilar constituencies. Even though $\alpha = 0.5$ ceases to be welfare-maximizing in the considered class of allocation rules, especially when the distribution of population sizes is skewed, it tends to perform better than a linear rule when individuals have a quadratic utility function. This is surprising given that our assumptions such as single-peakedness of preferences in a one-dimensional policy space and the prominent role of the (weighted) median are rather straightforward generalizations from the realm of binary voting.

Note additionally that our utilitarian welfare investigation builds on the restrictive postulate that different individuals derive the same satisfaction or dissatisfaction when a policy at a certain distance from their ideal point is implemented. In other words, we conduct interpersonal comparisons of utility. These cannot be avoided by any utilitarian welfare analysis in economics or political science. And, here, they can be defended by the a prioristic nature of the investigation: they express the value judgment that all individuals should be treated as anonymous equals in constitutional analysis. Still, the fact that our findings differ for different specifications of voter utility – and rather distinct conclusions might be derived concerning the most desirable allocation of voting weights in, e.g., the EU's Council of Ministers – might be seen as weakening the appeal of total expected utility as a guide to the "best" weight allocation rule.

References

Arnold, B. C., N. Balakrishnan, and H. N. Nagaraja (1992). *A First Course in Order Statistics*. New York: John Wiley & Sons.

Banks, J. S. and J. Duggan (2000). A bargaining model of collective choice. *American Political Science Review 94*(1), 73–88.

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

Barberà, S. and M. O. Jackson (2006). On the weights of nations: Assigning voting weights in a heterogeneous union. *Journal of Political Economy 114*(2), 317–339.

Beisbart, C. and L. Bovens (2007). Welfarist evaluations of decision rules for boards of representatives. *Social Choice and Welfare 29*(4), 581–608.

Beisbart, C. and L. Bovens (2011, forthcoming). Minimizing the threat of a positive majority deficit in two-tier voting systems with equipopulous units. *Public Choice*. [DOI: 10.1007/s11127-011-9810-2].

Beisbart, C., L. Bovens, and S. Hartmann (2005). A utilitarian assessment of alternative decision rules in the Council of Ministers. *European Union Politics 6*(4), 395–419.

Chang, P.-L., V. C. Chua, and M. Machover (2006). L S Penrose's limit theorem: Tests by simulation. *Mathematical Social Sciences 51*(1), 90–106.

Cho, S. and J. Duggan (2009). Bargaining foundations of the median voter theorem. *Journal of Economic Theory 144*(2), 851–868.

Cramér, H. (1946). *Mathematical methods of statistics*. Princeton Landmarks in Mathematics. Princeton, NJ: Princeton University Press. Reprinted 1999.

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

Feix, M. R., D. Lepelley, V. Merlin, J.-L. Rouet, and L. Vidu (2008). Majority efficient representation of the citizens in a federal union. mimeo, Université de la Réunion, Université de Caen, and Université d'Orléans.

Felsenthal, D. and M. Machover (1998). *The Measurement of Voting Power – Theory and Practice, Problems and Paradoxes*. Cheltenham: Edward Elgar.

Felsenthal, D. and M. Machover (1999). Minimizing the mean majority deficit: The second square-root rule. *Mathematical Social Sciences 37*(1), 25–37.

Gelman, A., J. N. Katz, and F. Tuerlinckx (2002). The mathematics and statistics of voting power. *Statistical Science 17*(4), 420–435.

Gerber, E. R. and J. B. Lewis (2004). Beyond the median: Voter preferences, district heterogeneity, and political representation. *Journal of Political Economy 112*(6), 1364–1383.

Harsanyi, J. C. (1955). Cardinal welfare, individual ethics, and interpersonal comparison of utility. *Journal of Political Economy 63*(4), 309–321.

Kaniovski, S. (2008). The exact bias of the Banzhaf measure of power when votes are neither equiprobable nor independent. *Social Choice and Welfare 31*(2), 281–300.

Kirsch, W. (2007). On Penrose's square-root law and beyond. *Homo Oeconomicus 24*(3/4), 357–380.

Koriyama, Y. and J.-F. Laslier (2011). Optimal apportionment. mimeo, École Polytechnique.

Kurz, S., N. Maaser, and S. Napel (2011). On the egalitarian weights of nations. mimeo, University of Bayreuth.

Leech, D. and R. Leech (2009). Reforming IMF and World Bank governance: in search of simplicity, transparency and democratic legitimacy in the voting rules. Warwick Economic Research Papers 914.

Lindner, I. and M. Machover (2004). L. S. Penrose's limit theorem: Proof of some special cases. *Mathematical Social Sciences 47*(1), 37–49.

Maaser, N. and S. Napel (2007). Equal represenation in two-tier voting systems. *Social Choice and Welfare 28*(3), 401–420.

Maaser, N. and S. Napel (2012). A note on the direct democracy deficit in two-tier voting. *Mathematical Social Sciences 62*(2), 174–180.

Nurmi, H. (1998). *Rational Behaviour and Design of Institutions*. Cheltenham: Edward Elgar.

Penrose, L. S. (1946). The elementary statistics of majority voting. *Journal of the Royal Statistical Society 109*(1), 53–57.

Rae, D. W. (1969). Decision rules and individual values in constitutional choice. *American Political Science Review 63*(1), 40–56.

Schwertman, N. C., A. J. Gilks, and J. Cameron (1990). A simple noncalculus proof that the median minimizes the sum of the absolute deviations. *American Statistician 44*(1), 38–39.

Yeo, I.-K. and R. A. Johnson (2000). A new family of power transformations to improve normality or symmetry. *Biometrika 87*(4), 954–959.