

Welfarist Evaluations of Decision Rules under Interstate Utility Dependencies

Claus Beisbart¹ and Stephan Hartmann²

May 6, 2009, to appear in *Social Choice and Welfare*

¹TU Dortmund, Institute for Philosophy and Political Science, D-44225 Dortmund, Germany

²Center for Logic and Philosophy of Science, Tilburg University, PO Box 90153, 5000 LE Tilburg, The Netherlands

Abstract

We provide welfarist evaluations of decision rules for federations of states and consider models, under which the interests of people from different states are stochastically dependent. We concentrate on two welfarist standards, viz. that the expected average utility for a person in the federation be maximized or that the expected utilities for the different people be equal. We discuss an analytical result that characterizes the decision rule with maximum expected average utility, set up a class of models that display interstate dependencies and run simulations for different dependency scenarios in the European Union. We find that the results that Beisbart and Bovens (SCW 29, p. 581, 2007) established for two types of models without interstate dependencies are fairly stable, if interstate dependencies are switched on. There are exceptions, though: Sometimes the way in which alternative decision rules shape the welfare distribution is significantly affected by such dependencies. These exceptions particularly include cases in which the interests of people from different states are partly anti-correlated.

1 Introduction

Recent times have seen a growing interest in welfarist evaluations of decision rules (see Beisbart et al. 2005, Coelho 2005, Barberà & Jackson 2006, Beisbart & Bovens 2007 and Fleurbaey 2009, for instance). Welfarist approaches evaluate alternative decision rules for federations of states (constituencies, ...) by comparing the resulting welfare distributions. The aim of this paper is to extend the welfarist approach by relaxing crucial assumptions on which much work so far is based. We consider federations for which the interests of the people from different states are stochastically dependent.

In order to understand the welfarist approach, consider the following problem: A *federation* of states has a *decision board* where *representatives* vote on proposals drafted by some committee. Each representative represents a state. The decisions of the board are taken according to a *decision rule*. The problem is: Which decision rule can be recommended for the federation – which decision rule is best?

This problem has many real world applications. Recently, the Council of Ministers in the European Union (CM, for short) has received much attention in political as well as in academic circles (see Felsenthal & Machover 2000, e.g.). In the CM, representatives

of the states from the European Union (EU) vote on proposals drafted by the European Commission. The problem now is: What decision rule is best in this case?

In order to judge a decision rule better than another, we need a standard of evaluation. *Welfarist* evaluations of decision rules focus on the welfare distribution that results from a decision rule. They are consequentialist, and ask the question: cui bono – to whose benefit?

People can benefit from decision rules in the following way (cf. Beisbart et al. 2005, Sec. 1): If a particular proposal passes, some people will take profits, whereas other people will have to take some costs. For instance, if import tariffs are imposed on sugar beets in the EU, sugar farmers in the EU will take profits, whereas sugar consumers in the EU will have to pay. Now, whether the proposal about import tariffs passes, depends on the decision rule. The crucial question is whether states that vote in favor of the tariffs have sufficiently many votes to push the proposal through. In this way, the decision rule shapes the welfare distribution for the people in the EU.

Of course, it doesn't make sense to assess a decision rule on the basis of a particular proposal. Rather, the question is, what welfare distribution is produced in the long run on the basis of many proposals that are not yet known in detail. We will thus evaluate alternative decision rules assuming a probability model for proposals, and we will quantify the welfare of a person using an expected utility.

In this paper we will consider two welfarist *standards* for decision rules (cf. Beisbart & Bovens 2007, p. 582):

\mathcal{U} Decision rule R is pro tanto better than rule S, if the *expected average utility* in the federation is *larger* (see Dancy 2004, Ch. 2 for the notion of “pro tanto”). Here, averaging means to average over the people in the federation. If the populations are not affected by the proposals, maximizing the expected average utility is equivalent to maximizing the expected net utility in the federation.

\mathcal{E} Decision rule R is pro tanto better than rule S, if the distribution of the expected utilities for different people is closer to an equal distribution.

Whereas, under \mathcal{U} , the best decision rule maximizes expected utility, under \mathcal{E} , the best decision rule minimizes inequality.

Very recently, decision rules have been evaluated using these standards. Beisbart et al. (2005) compare a few alternative decision rules regarding \mathcal{U} for the CM. Barberà & Jackson (2006) generally specify the decision rule that maximizes expected utility. Beisbart & Bovens (2007) consider both \mathcal{U} and \mathcal{E} . However, all of these works rely on the assumption that the utilities from proposals for people from different states are stochastically independent. It is thus excluded that people from Spain, say, are typically proposed similar utilities as people from Portugal are. More colloquially spoken, the interests of the people from different states are assumed to be independent.

Assuming this interstate independence was certainly justified for developing the welfarist approach. But the assumption is not very plausible for many real world applications; and for advancing welfarist evaluations of decision rules further, it is crucial to give up the assumption. We will therefore do exactly this. Our approach is exploratory, i.e. we will not search for an empirically adequate model of some specific federation. Rather, we will set up a general framework for modeling dependencies, consider various examples of *dependency scenarios* and check how the rankings of alternative decision rules are affected. Beisbart & Hartmann (2006) also introduce dependencies, but they only compare the performances of very few decision rules. Also, they work with a different selection of models for proposals.

We will focus on the CM in the European Union as an example. However, the methods in this paper can be applied to any decision board in any federation. Also, our results indicate effects to be expected for other boards in other federations as well.

An alternative approach for assessing decision rules focuses on *voting power*. The voting power of a voter is the extent to which her vote can affect the outcome of a decision. It is often quantified in terms of the Banzhaf measure of voting power (see Felsenthal & Machover 1998 for an introduction, Laruelle & Valenciano 2005 for clarification and Felsenthal & Machover 2000 for an application to the CM). The Banzhaf measure for a voter equals the probability that her vote is pivotal. For calculating the Banzhaf measure one assumes that all votes are stochastically independent. Thus, a stochastic independence assumption is crucial for the voting power approach as well. This assumption is not completely unrelated to the assumption of independence in the welfarist approach, since the Banzhaf measure for a voter is a linear transform of the expected utility for that voter, if certain assumptions on the utilities from the proposals are made (Beisbart et al. 2005, Appendix and Beisbart & Bovens 2007). For probability models on votes with dependent votes see Good & Mayer (1975), Chamberlain & Rothschild (1981) and Gelman et al. (2004).

The plan of the paper is as follows: In Sec. 2, we briefly outline the welfarist framework. In Subsec. 2.2, a general analytical result regarding the expected average utility is stated. This result, however, is not useful in practice. We thus turn to models in Sec. 3. The expected utilities under these models cannot be calculated analytically any more, so we carry out simulations (Sec. 4). Our results are described in Secs. 5 and 6. We draw our conclusions in Sec. 7.

2 The welfarist framework

In order to evaluate decision rules on welfarist grounds, we need to describe how decision rules effect welfare distributions (cf. Schweizer 1990 and Beisbart et al. 2005 for the framework).

2.1 Basic definitions

Assume that the *federation* has n people. Label the *people* in the federation with numbers ranging from 1 to n . Let there be N *states* in the federation and label them from 1 to N . In the following, we will use lowercase letters for labels of people, whereas uppercase letters will be used for label of states. For $I = 1, \dots, N$, let K_I be the set of labels for people from the I th state. Set $n_I \equiv \#K_I$. Of course, $\sum_I n_I = n$.

Consider now a single proposal. The proposal can either be accepted or rejected. Without loss of generality, we assume that, for every $i \in \{1, \dots, n\}$, citizen i receives utility v_i , if the proposal is accepted, and zero utility otherwise. We assume interpersonal comparability of utility scales and form the average utility for citizens from state I from the proposal:

$$\bar{v}_I = \frac{1}{n_I} \sum_{i \in K_I} v_i. \quad (1)$$

The average utility from the proposal across the federation, call it \bar{v} , is then

$$\bar{v} = \frac{1}{n} \sum_{I=1}^N n_I \bar{v}_I. \quad (2)$$

We model proposals in a probabilistic way and think of the v_i as values of random variables V_i . Similarly, \bar{v}_I is the value of a random variable \bar{V}_I for $I = 1, \dots, N$, and \bar{v} is the value of a random variable \bar{V} . The random variables V_i have $p(v_1, \dots, v_n)$ as their joint probability density. In our analysis we take p to be exogenously fixed. p induces a joint probability density over the random variables \bar{V}_I ($I = 1, \dots, N$), call it $p(\bar{V}_1, \dots, \bar{V}_N)$.

The vote of state I 's representative is modeled as another random variable Λ_I . Λ_I takes the value $\lambda_I = 1$, if the representative votes in favour of the proposal (votes yes, for short). It takes the value $\lambda_I = 0$, if the representative votes against the proposal (votes no, for short). Vectors $(\lambda_1, \dots, \lambda_N)$ are called voting profiles.

How do the representatives vote and what is their strategy? In this paper we assume that a representative will vote yes, iff the proposed average utility for her state I , \bar{v}_I is positive. That is, $\lambda_I = \theta(\bar{v}_I)$, where θ is a variation of the well-known Heaviside step function:¹

$$\theta(x) \equiv \begin{cases} 0 & \text{if } x \leq 0 \\ 1 & \text{if } x > 0. \end{cases} \quad (3)$$

Under certain conditions, this choice of strategies constitutes a Nash equilibrium (Beisbart & Bovens 2007, Sec. 2).²

A decision rule R can be thought of as a function D^R that maps voting profiles into $\{0, 1\}$: $D^R : \{0, 1\}^N \rightarrow \{0, 1\}$. We set D^R at 1 for acceptance, and at 0 for rejection. There are $2^{(2^N)}$ possible decision rules.

We introduce other sets of random variables U_i^R , \bar{U}_I^R and \bar{U}^R . They denote a utility that is received, once a decision has been taken following rule R . For instance, U_i^R is the utility for person i , after a decision has been taken. In more detail,

$$U_i^R = D^R \times V_i, \quad \bar{U}_I^R = D^R \times \bar{V}_I, \quad \bar{U}^R = D^R \times \bar{V}. \quad (4)$$

In the following, we will often suppress the upper label R .

We can now calculate the expected utilities that a person i derives from a decision rule R under the probability distribution $p(v_1, \dots, v_n)$ over proposals, call it $E[U_i^R]$. It reads:

$$E[U_i^R] = \int dv_1 \cdots dv_n p(v_1, \dots, v_n) v_i D^R(\theta(\bar{v}_1), \dots, \theta(\bar{v}_N)), \quad (5)$$

where the \bar{v}_I depend on the v_j as specified in Eq. (1). In a similar way, we obtain for $E[\bar{U}_I^R]$

$$E[\bar{U}_I^R] = \int d\bar{v}_1 \cdots d\bar{v}_N p(\bar{v}_1, \dots, \bar{v}_N) \bar{v}_I D^R(\theta(\bar{v}_1), \dots, \theta(\bar{v}_N)) \quad (6)$$

and for $E[\bar{U}^R]$:

$$E[\bar{U}^R] = \frac{1}{n} \sum_{I=1}^N n_I E[\bar{U}_I^R]. \quad (7)$$

¹We thus assume that the representative of state I votes no, if $\bar{v}_I = 0$. Of course, this assumption is in a way arbitrary, since voting yes in this case would not make a difference for state I . Nevertheless, our choice does not make a difference, if $\bar{v}_I = 0$ has zero probability measure. That this is so is a very natural assumption.

²Barberà & Jackson (2006) are more general at this point by not assuming any specific strategy.

2.2 An analytical result for expected utility

Let us now focus on the expected average utility, i.e. the expectation value of \bar{U} . The best thing one could do for maximizing the expected average utility is this: Proposals with $\bar{v} > 0$ are accepted and proposals with $\bar{v} < 0$ are rejected. The expected average utility is then

$$E [\bar{V}\theta(\bar{V})] , \quad (8)$$

where the expectation value is calculated on the basis of $p(\bar{v}_1, \dots, \bar{v}_n)$. This is the maximum utility that one can draw from the proposals (Barberà & Jackson 2006, p. 325).

However, if one proceeds this way, one does not vote according to a decision rule. A decision rule takes votes and not utilities as inputs. Decisions according to a rule lose information present in the utilities (ib.). The expected utility for a decision rule will therefore in general be smaller than $E [\bar{V}\theta(\bar{V})]$.

This can be illustrated as follows. For an arbitrary rule R, we calculate $E [\bar{U}^R]$ by conditioning on all possible voting profiles $(\lambda_1, \dots, \lambda_N)$:

$$E [\bar{U}^R] = \sum_{\lambda_1, \dots, \lambda_N} D^R(\lambda_1, \dots, \lambda_N) E [\bar{V} | \theta(\bar{V}_1) = \lambda_1, \dots, \theta(\bar{V}_N) = \lambda_N] p(\lambda_1, \dots, \lambda_N) , \quad (9)$$

where $p(\lambda_1, \dots, \lambda_N)$ is the probability for a particular voting profile $(\lambda_1, \dots, \lambda_N)$. It can be calculated from the joint probability distribution over the \bar{V}_I . If \bar{V} can take positive and negative values with finite probabilities, each, the conditional expectation value $E [\bar{V} | \theta(\bar{V}_1) = \lambda_1, \dots, \theta(\bar{V}_N) = \lambda_N]$ in Eq. (9) can be expressed as

$$E [\bar{V} | \theta(\bar{V}_1) = \lambda_1, \dots, (\bar{V}_N) = \lambda_N \wedge \bar{V} > 0] p(\bar{V} > 0 | \lambda_1, \dots, \lambda_N) \\ + E [\bar{V} | \theta(\bar{V}_1) = \lambda_1, \dots, (\bar{V}_N) = \lambda_N \wedge \bar{V} < 0] p(\bar{V} < 0 | \lambda_1, \dots, \lambda_N) . \quad (10)$$

Here, the first addend is positive, whereas the second is negative. Looking at the value of \bar{V} itself always picks the positive term for every voting profile. But a decision rule has to yield either acceptance or rejection for each particular voting profile. If the rule yields acceptance for a voting profile $(\lambda_1, \dots, \lambda_N)$, then the second, negative addend in Eq. (10) is bought in as well. If the rule yields rejection, then the positive contribution from the first addend is missed. Thus, in either case, utilities are lost.

Obviously, the best *decision rule* R that one can think of yields acceptance if $E [\bar{V} | \lambda_1, \dots, \lambda_N]$ (or Eq. 10) is positive, and rejection if this term is negative. This follows from the proof that Barberà & Jackson (2006) give for their Theorem 1.

As Barberà & Jackson (2006) show, the decision rule that maximizes expected average utility can be described in a very simple way, if the \bar{V}_I are stochastically independent – it turns out to be a weighted rule (cf. Def. 2.3.14 in Felsenthal & Machover 1998, pp. 29–30). Call this rule the BJ solution. In the case of dependent utilities \bar{V}_I , the decision rule that maximizes expected average utility can in general not be cast in such simple terms any more. Fleurbaey (2009) is in fact able to characterize that rule in terms of weights, but the weights do now depend on the voting profile – i.e. if a different subset of representatives votes yes, a different set of weights is to be used. But such weights are not very useful any more. In the worst case, the optimal decision rule may thus just be a list that assigns each voting profile 1 or 0 according to what we have said. Call this the generalized BJ solution.

It can be doubted whether the generalized BJ solution is helpful in practice. First, it can be very difficult to obtain sufficient data for estimating the $E [\bar{U} | \lambda_1, \dots, \lambda_N]$. Second, if the generalized BJ solution cannot be cast in simple terms, voting itself will become difficult.

There will be no way to keep in mind the voting rule, so one can only look up the result from a list, once the votes are taken. The result might come as a surprise for the representatives themselves. If the generalized BJ solution is not monotonic – and this is very well possible – then $\lambda_I = \theta(\bar{v}_I)$ might not be a Nash equilibrium any more, and the representatives might have to reconsider their strategies.

For these reasons, we will not consider this analytical result further. Instead, we confine ourselves to a family of voting rules that can be cast in simple terms. In this paper we consider weighted voting rules (cf. the notion of a weighted voting game, Def. 2.3.14 in Felsenthal & Machover 1998, pp. 29–30). Under a weighted voting rule, the representative of state I has a weight w_I for $I = 1, \dots, N$, and a proposal is accepted, iff the sum of weights associated with the yes-votes exceeds a certain threshold of acceptance t . A weighted voting rule R can be represented as $D^R = \theta(\sum_I \lambda_I w_I - t)$. In the following, we restrict ourselves to weights that only depend on population size n_I . Moreover, it is assumed that the weights depend on population in the following way (Bovens & Hartmann 2007, Eq. 1):

$$w_I \propto (n_I)^\alpha . \quad (11)$$

Here, the exponent α is a measure of degressive proportionality. $\alpha = 0$ yields equal weights, $\alpha = 1$ yields proportional weights, and $\alpha \in (0, 1)$ parameterizes weightings in between. Without loss of generality, we will also assume that the voting weights are normalized: $\sum_I w_I = 1$. Our question then is: Which pairs of numbers (α, t) do best wrt to \mathcal{U} and \mathcal{E} , respectively? And is there an (α, t) -rule that one can recommend on both counts?³

3 Models for proposals

In this work, we confine ourselves to a particular class of probabilistic models for proposals. They are inspired by Crain et al. (1993).⁴

For setting up the class of models, we partition the federation into N_G groups of states (abbreviated as "G", labeled with uppercase Greek letters Γ, Δ). That is, each state is a member of exactly one group. The groups need not be organized in any way. Rather the idea is that people from the same group have dependent utilities, even if they are from different states.

The utilities from proposals for person i , i.e. the values v_i of the random variables V_i are now determined as a sum of stochastically independent addends:

$$v_i = \chi_\Gamma^G + \epsilon_I^S + \epsilon_i^P . \quad (12)$$

Here, the first addend, χ_Γ^G , is common to the people from the same group, the second addend, ϵ_I^S , is common to the people from the same state, and the third addend, ϵ_i^P , is specific to the individual person i . In Eq. (12) we assume that person i lives in state I ($i \in K_I$) and that state I is in group Γ . Let us now fix the details.

³Note that two pairs $(\alpha, t) \neq (\alpha', t')$ may represent the same decision rule. This is so, if they induce the same function D (cf. Barberà & Jackson 2006, p. 324). Nevertheless, for simplicity, pairs (α, t) will sometimes be called rules.

⁴The model by Crain et al. (1993) is about fractions of votes. We are concerned with utilities instead. This is of significant advantage. For the model by Crain et al. (1993) leads to finite (but minimal) probabilities that the fraction of votes for one candidate is larger than 1. The reason is that Gaussian probability densities for fractions of votes are assumed. Our model, which is about utilities, does not suffer from such a problem.

- χ_{Γ}^G is a *group utility tendency*. We assume a joint probability distribution over the $\chi_{\Gamma}^G, p^G(\chi_1^G, \chi_2^G, \dots, \chi_{N_G}^G)$. This probability distribution will always be assumed to be a multivariate Gaussian. Under this multivariate Gaussian, each χ_{Γ}^G has zero mean. The standard deviations for the different χ_{Γ}^G are identically set at σ^G , and the covariances are identically set at ρ^G . This leaves us with a covariance matrix with $(\sigma^G)^2$ as diagonal elements, and ρ^G as off-diagonal elements. Since the covariance matrix has to be positively semidefinite, we obtain the following constraints on ρ^G : $\rho^G / (\sigma^G)^2 \in [-1/(N_G - 1), 1]$.
- $\chi_{\Gamma}^G + \epsilon_I^S$ is a *national utility tendency*. Here ϵ_I^S is randomly drawn from a Gaussian with zero mean and a standard deviation σ_I^S . It is independent from χ_{Γ}^G . The national utility tendencies are thus scattered around the group utility tendencies χ_{Γ}^G . For simplicity, we will assume that the ϵ_I^S follow the same probability density, i.e. the σ_I^S are identical for all states I : $\sigma_I^S = \sigma^S$.
- ϵ_i^P encodes a utility contribution that is completely personal. The utilities for citizens from state I are scattered around the national utility tendency of their state. ϵ_i^P is supposed to be the value of a random variable that is independent from any other random variable. For simplicity, we assume that ϵ_i^P is normally distributed with zero mean and a standard deviation $\sigma_i^P = \sigma^P$ that is constant across the federation.

Note that, because of these assumptions, each random variable V_i has zero mean. This means that the benefits and the costs that are proposed for person i balance each other out, if the mean is taken.

Beisbart & Bovens (2007) focus on models with no interstate dependencies ($\sigma^G = \rho^G = 0$) and consider two types of federations. Each state in a federation might be an *aggregate*, i.e. the utilities of people from the same state are independent. Alternatively, each state in a federation might be an *interest group*, i.e. the utilities of people from the same state are perfectly correlated. Beisbart & Bovens (2007) do not consider interstate dependencies. Thus their *aggregate model* is a special model in our class, where $p^G(\chi_1^G, \dots)$ factorizes into $\delta(\chi_1^G) \times \delta(\chi_2^G) \dots$ and where $\sigma^P \neq 0$ and $\sigma^S = 0$. Their *interest group model* corresponds to $p^G(\chi_1^G, \dots) = \delta(\chi_1^G) \times \delta(\chi_2^G) \dots$, $\sigma^S \neq 0$ and $\sigma^P = 0$. Clearly, for these simple models, the welfare distribution that a decision rule produces is a function of the population sizes of the states only.

In this paper, we will start with the aggregate and the interest group models and we will then switch on interstate dependencies. In this way modifications of the aggregate/interest group model arise. For clarity, the aggregate/interest group model defined above will often be called *default aggregate/interest group model* henceforth.

Each specific model that we consider can be fixed by taking three decisions:

- Chose $\sigma^S = 0, \sigma^P = 1$ (for modifications of the default aggregate model) or $\sigma^S = 1, \sigma^P = 0$ (for modifications of the default interest group model).⁵
- Fix a partition of states.
- Fix the correlation matrix for $p^G(\chi_1^G, \dots)$, i.e. fix σ^G and ρ^G .

It is easy to see that, for each set of parameter values, the random variables U_i follow a multivariate normal distribution.

⁵Here, σ^P (σ^S) can without loss of generality be set at 1 (see p. 8).

Note that, under our model, the marginal probability densities for the personal utilities V_i are identical. The underlying idea is this: In fixing the probability model for proposals, there is no reason to assume that the proposals are biased towards any person in the federation.

For some real world federations, unbiasedness may not be a realistic assumption. However, in examining generic scenarios with dependencies, it is reasonable to start with the assumption of unbiasedness.

Symmetry. Start with a rule (α, t) and suppose that the threshold can not be exactly met by a coalition. That is, there is no subset $\mathcal{I} \subseteq \{1, \dots, N\}$ with $\sum_{I \in \mathcal{I}} w_I = t$. Then the rule $(\alpha, 1 - t)$ yields exactly the same welfare distribution as (α, t) . This result generalizes Proposition 1 in Beisbart & Bovens (2007), and the proof parallels their proof.

Suppose now that $N > 1$, $t \neq 0, 1$ and $\alpha \geq 0$. Consider the set of (n_1, \dots, n_N) -vectors with the following property: $n_I \neq 0$ for $I = 1, \dots, N$ and there is a subset $\mathcal{I} \subseteq \{1, \dots, N\}$ with $\sum_{I \in \mathcal{I}} (n_I)^\alpha = t$ (i.e. there is a coalition that exactly meets the threshold). It can be shown that this set has zero measure in \mathbb{R}^N . As a consequence, we do not expect that an arbitrary threshold $t \neq 0, 1$ can be exactly met by a coalition in a given federation. We therefore expect that t and $(1 - t)$ produce the same welfare distributions for a given α and a particular model for most federations, unless $t = 0, 1$. We will therefore only consider thresholds $t \geq .5$.

Scaling. Our class of models have the following scale invariance. If we multiply all model parameters (i.e., σ^G , ρ^G, σ^S and σ^P) by some real number $\kappa > 0$, our results scale linearly with κ as well. Thus effectively one model parameter drops out.

4 How results are obtained

In order to compare alternative decision rules regarding the standards \mathcal{U} and \mathcal{E} , we need to calculate the $E[U_i^R]$. Unfortunately, we are not aware of any analytical method to evaluate the integral from Eq. (5). We therefore carry out simulations. Utility vectors for proposals are randomly drawn, the voting profile is determined, and the collective decision is taken. If there is acceptance, the proposed utilities are distributed; if there is rejection, they are not so distributed. For saving computation time, we do not simulate the utilities of different persons i , i.e. the v_i . Rather, the \bar{v}_I are drawn randomly. Given values of the group utility tendencies, χ_1^G, \dots , the random variables \bar{V}_I are independent. The conditional probability density for values \bar{v}_I of \bar{V}_I , $p(\bar{v}_I | \chi_1^G, \chi_2^G, \dots)$ is a Gaussian with mean χ_1^G and with a standard variation $\sqrt{(\sigma^S)^2 + \frac{1}{n_i^S} (\sigma^P)^2}$.

In our simulations, we obtain the $E[\bar{U}_I]$, from which the expected average utility, $E[\bar{U}]$, is determined. For applying our egalitarian standard, we need a measure of inequality. So far, many measures of inequality have been proposed, but no consensus has yet been reached (see Sen 1997, e.g.). We will therefore work with a very straight-forward measure, viz. the standard deviation of the expected utilities for the people, viz. the $E[u_i]$. If f_I is the population fraction of state I , then our measure of inequality In can be calculated via

$$\text{In}^2 = \sum_I (f_I E[\bar{U}_I]^2) - E[\bar{U}]^2 . \quad (13)$$

This measure can be defended by uniquely characterizing it in terms of plausible desiderata

(Chakravarty 2001).⁶ Regarding our investigation, that measure does not only seem reasonable because of its straight forwardness, but also because it is well-defined even if some expected utilities are negative, which they are sometimes (ibid., p. 82). We will later (at the end of Sec. 6) also consider a different measure of inequality.

We checked the reliability of our simulations by comparing to analytical results for a federation of 15 states under the default aggregate and interest group models. For such a federation, it is still possible to obtain analytical results by summing up 2^{15} addends (Beisbart et al. 2005, Appendix). We find that 2 million realizations do reasonably well for 15 states. We thus run 2 million realizations of votes for our results.

For each model considered, we scan the (α, t) space. We often work with a spacing of .02 along the α - and .02 along the t -axis. We never consider $\alpha > 1$ or $\alpha < 0$. Obviously, our results hold strictly speaking only on our grid, and it is possible that a rule that has an $\alpha > 1$ is optimal wrt \mathcal{U} or \mathcal{E} . Moreover, errors due to the finite number of simulations might affect our results.

In the following, results for $E[\bar{U}]$ and the measure of inequality are always normalized by

$$E[\bar{V}\theta(\bar{V})] . \quad (14)$$

As we have shown in Subsec. 2.2, this quantity is an upper limit on $E[\bar{U}]$. We call the normalized expected utility efficiency, and the normalized measure of inequality normalized spread (of utilities).

5 Results for models without interstate dependencies

We will now present results. We start with very simple models and move on to more complicated models. Let us first mention results for the default aggregate/interest group models (Beisbart & Bovens 2007). Models of this kind do not yet display interstate dependencies ($\sigma^G = \varrho^G = 0$). Interstate dependencies are discussed in the next section.

5.1 The default aggregate model

Under the default aggregate model, utilities from proposals are independent for any two persons, regardless whether they are from the same state or not. Thus, $\sigma^S = 0$, whereas $\sigma^P = 1$. The default aggregate model is the fixed-size-block model with one person per block and no bias in Barberà & Jackson (2006), pp. 326, 328–9.

Under the default aggregate model, the distributions for the average utilities \bar{V}_I from proposals are different for the different states I . They are Gaussians that peak at 0 and that have a standard deviation of $\sigma^P/\sqrt{n_i}$. Thus, smaller states have a larger spread in their average utility from proposals.

Results are shown in Fig. 1. The efficiency (left panel) and the normalized spread (right panel) are shown for various decision rules in the (α, t) -space. We plot the efficiency and the normalized spread as functions of α . Different point styles designate different thresholds.

The rule $(\alpha = .5, t = .5)$, also known as Penrose 50, yields maximum expected utility for the federation. Penrose 50 does not only maximize expected utility in the constrained class of (α, t) -rules, but also in the class of all decision rules (Barberà & Jackson 2006, pp. 327 and 329 f.). For a qualitative understanding, note that the quantities $E[\bar{V}_I|\bar{V}_I > 0]$ and

⁶To be more precise, Chakravarty (2001) defends the variance, but for minimization, it does not make a difference whether we consider the variance or the standard deviation.




Figure 1: The welfare distribution for the default aggregate model for the EU (27 members). Efficiency (left panel) and normalized spread (right panel) are shown as functions of the α -parameter for different thresholds: triangles: $t = .5$; squares: $t = .6$; hexagons: $t = .8$. The results are based upon simulations with 2 million realizations.

$E[\bar{V}_I | \bar{V}_I < 0]$ are proportional to the standard deviation of \bar{V}_I . As a consequence, people from smaller states tend to benefit more than people from larger states, given that they benefit. Equally, they suffer more than larger ones, given they suffer. Consequently, what citizens from the smaller states get, makes a larger difference to the expected utility $E[U]$ than under the default interest group model, e.g. Therefore, a decision rule has to strike a fair compromise between larger and smaller states, if it is to maximize expected utility. For two simple mathematical arguments why Penrose 50 maximizes expected utility for the default aggregate model, see Beisbart & Bovens (2007), Sec. 5.

Our measure of inequality minimizes around $\alpha = .5$ for thresholds $t = .5, .6$. The reason why there is not perfect equality for equal weights ($\alpha = 0$) is again that people from smaller states tend to benefit more, given that they benefit, etc. Thus, under equal weights, the $E[\bar{U}_I]$ of smaller states are larger than those of larger states, and we have some inequality. Thresholds that are very close to 0 or 1 produce almost flat inequality curves close to zero.

Remarkably, under the default aggregate model, one rule – Penrose 50 – is almost optimal on both of our welfarist standards: It yields maximum expected utility and does astonishingly well in terms of equality, given the large range of In-values in the right panel of Fig. 1. Note, however, that, on our scans and simulations, Penrose 50 is not strictly the best rule regarding equality. Curves for high thresholds yield lower minima on our simulations.⁷

5.2 The default interest group model

Under the interest group model, on a particular proposal, people from the same state have exactly the same utility. Utilities for people from different states are independent and follow the same probability density. The interest group model originates, if $\sigma^P = 0$ and σ^S is set at one. Note that, under the interest group model and all of its modifications in this paper,

⁷Admittedly, this and other assessments that are based on both of the standards, are a bit problematic, unless the standards are related in a precise way – say by lexicographically ordering them or by weighing them in a particular way. What we are really looking for, if we assess decision rules on the basis of both standards, are rules that do fairly well regarding efficiency comparing the range of values that the efficiency takes in the related viewgraph, and the same regarding normalized spread.




Figure 2: The default interest group model. For point styles see the caption of Fig. 1.

the \bar{V}_I follow the same probability density. Our default interest group model is the fixed-number-of-blocks model with one block per state and no bias in Barberà & Jackson (2006), pp. 326, 328–9.

Results are shown in Fig. 2. For the thresholds shown, the expected utility mostly goes up, as α increases (left panel).⁸ The reason is that, under the default interest group model, each person has the same benefits from proposals, if she benefits. Since there are more people in larger states, the net utility is higher, if larger states have higher weights, as long as $\alpha \leq 1$. Among the thresholds shown in this Figure, $t = .5$ does best. The maximum expected utility is obtained at for a very high α and a threshold $t = .5$. The rule ($\alpha = 1, t = .5$) indeed represents the BJ solution.

Regarding our measure of inequality (right panel), $\alpha = 0$ is optimal for any threshold t – at this point the states have the same weights and thus according to the default interest group model the same expected utilities. As α increases, the differences between the different states’ weights become larger, and thus their expected utilities do so as well (cf. Proposition 1 on p. 331 in Barberà & Jackson 2006). For $t \geq .5$, larger thresholds yield less inequality – the higher the threshold, the more states have to agree with the outcome of the decision, and thus the smaller are the differences between them.

Altogether, for the default interest group model, our welfarist standards pull into different directions. Maximizing expected utility requires one to adopt high values of α and thresholds close to .5. Minimizing inequality requires one to adopt low values of α and thresholds close to 1 or 0.

5.3 The transition between the default aggregate and the default interest group model

Clearly, the default aggregate and interest group models are idealizations. Under realistic circumstances, citizens from one state will have some similarities in interests, such that the utilities from proposals are dependent. But we do not expect full correlations; rather the

⁸Here and at similar occurrences, “mostly” or “for most parts” means: in a large fraction of the $\alpha \in [0, 1]$ range, the function under consideration displays a certain feature, as we move from one α to the next in our spacing. There might be exceptions, but that there are is not implied. Note that really the curves for efficiencies and normalized spread are step functions. It is also possible that the functions jump back and forth on small scales, if they are sufficiently resolved.

partition	$\alpha = 0$	$\alpha = .5$	$\alpha = 1$
P0	1	1	1
P1	.85/.15	.66/.34	.46/.54
P2	.48/.52	.60/.40	.73/.26

Table 1: Basic characteristics of the three partitions P0–P2 that we investigate below. The entries in the table are the fractions of weights that a group has for a given α . Particularly, in the first column, the fractions of the numbers of states are given. In the last column, the population fractions of the groups are given. For P1, the first number refers to the group of small states. For P2, the first number refers to the group of western states.

utilities from proposals will fluctuate within one state. Thus, a model between the default aggregate and interest group models seems more realistic.⁹

Within our class of models, a transition between the default aggregate and the interest group models can be effected. Either we start with the default interest group model ($\sigma^P = 0$) and let intrastate fluctuations grow ($\sigma^P > 0$). As σ^P/σ^S approaches infinity, we end up with the aggregate model (in different units of utilities). Alternatively, we can start with the default aggregate model ($\sigma^S = 0$), and let the σ^S parameter grow. This transition is covered by the transition that Beisbart & Bovens (2007) study, although they use different terms.

Suppose now, we set both σ^P and σ^S at non-zero values and consider a fixed voting rule. We know that the $E[\bar{U}_I]$ only depend on the probability density for the \bar{V}_I . It can be shown that, in the moments of this probability density, terms with a σ^P are damped with a factor of $\frac{1}{\sqrt{n_I}^\nu}$, where $\nu \geq 1$ is a natural number. Thus, unless $\sigma^P/\sqrt{n_I} \gtrsim \sigma^S$, the results are pretty much the same as for the default interest group model. Put differently, the aggregate model is unstable; already a σ^S that is significantly smaller than σ^P can push the results to something like the results for the interest group model. We also confirmed this in terms of simulations.

We thus follow Beisbart & Bovens (2007) in concluding that the interest group model is the generic model to look at. Nevertheless, in the following, we will consider modifications of both types of models. But we will not investigate in detail how the transitions between the modifications of the aggregate and the interest group models look like. That is, we either set σ^S or σ^P at zero.

6 Results for models with interstate dependencies

We now switch on dependencies between the utilities that are proposed for citizens from different states. For this, we partition the EU into groups. On each proposal, citizens from states of the same group get the same group utility tendency χ_I^G .

For simplicity, under each partition, we will consider only one value of σ^G . For each partition and each σ^G , three values of ϱ^G are investigated, viz. $\varrho^G = 0$, the maximal and the minimal possible value of ϱ^G . We consider three different partitions. Under the partition P0, the whole federation forms one big group. Under P1, the four largest states form one group, the other states form another group. Under P2, there are two groups, which are not defined in terms of population, but rather comprise the more western/eastern states (the

⁹Apart from the default aggregate and interest group models and transitions between them, there is a third model defined by $\sigma^P = \sigma^S = 0$. But such a model is not interesting: Every person in the federation is always proposed the same utility $u_i = 0$.

agg, default, \mathcal{U}	$\alpha = .5, t = .5$		
agg, default, \mathcal{E}	high t		
agg with $\sigma^G = 2 \times 10^{-4}$	$\varrho^G = 0$	$\varrho^G = (\sigma^G)^2$	$\varrho^G = -(\sigma^G)^2$
P0, \mathcal{U}	$\alpha = .52, t = .5$	–	–
P0, \mathcal{E}	high t	–	–
P1, \mathcal{U}	$\alpha = .7, t = .5$	$\alpha = .52, t = .5$	$\alpha = .9, t = .5$
P1, \mathcal{E}	high t	high t	high t
P2, \mathcal{U}	$\alpha = .5, t = .5$	$\alpha = .52, t = .5$	$\alpha = .78, t = .5$
P2, \mathcal{E}	high t	high t	high t
<hr/>			
int, default, \mathcal{U}	$\alpha = 1, t = .5$		
int, default, \mathcal{E}	high t		
int with $\sigma^G = .4$	$\varrho^G = 0$	$\varrho^G = (\sigma^G)^2$	$\varrho^G = -(\sigma^G)^2$
P0, \mathcal{U}	$\alpha = .66, t = 0.5$	–	–
P0, \mathcal{E}	high t	–	–
P1, \mathcal{U}	$\alpha = .88, t = .5$	$\alpha = .62, t = .5$	$\alpha = 1, t = .5$
P1, \mathcal{E}	$\alpha = .76, t = .92$	high t	high t
P2, \mathcal{U}	$\alpha = .74, t = .5$	$\alpha = .62, t = .5$	$\alpha = 1, t = .5$
P2, \mathcal{E}	high t	high t	high t

Table 2: A summary of the decision rules that do best on our standards \mathcal{U} and \mathcal{E} , respectively. Upper half: default aggregate model and its modifications. Lower half: default interest group model and its modifications.

western/eastern states, for short). The eastern states encompass Cyprus, Estonia, Slovenia, Latvia, Lithuania, Finland, Slovakia, Bulgaria, Austria, Hungary, the Czech Republic, Greece, Romania and Poland.

We list important characteristics of the groups under the partitions in Table 1. Results for the decision rules that do best in terms of \mathcal{U} and \mathcal{E} , respectively, are summarized in Table 2, where the different partitions and models are considered. The results will of course be explained in more detail later in the paper. The results in the table rely on simulations that scan the (α, t) space, except for the default models and the standard \mathcal{U} , for which the analytic BJ solution is shown. Only thresholds up to $t = .98$ were considered. For high t , the results for alternative values of α and t are very close and probably not accurate enough as to allow for a meaningful comparison. Thus no optimal (α, t) pair is given. Note also that thresholds t and $(1 - t)$ almost always produce the same results; thus, if (α, t) is optimal wrt to a particular standard, $(\alpha, 1 - t)$ is very likely to be so as well.

6.1 Modifications of the aggregate model with dependencies

We first consider modifications of the aggregate model. Qualitatively, the picture is as follows: The utilities from proposals v_i have a group-specific and a personal component, each. States do not have any significance regarding the utilities proposed. They can be thought of as random parts of groups.

In the following models, σ^G is always set at 2×10^{-4} . The idea is that the utility contribution from the group utility tendency χ_I^G fluctuates significantly less than the contribution for individual persons. For Ireland, the standard deviation in group utility tendency is 40% of the standard deviation of \bar{V}_I under the default aggregate model. For some of the models investigated below, we also changed the value of σ and found qualitatively similar results.




Figure 3: Results for a modification of the aggregate model under partition P0 (the states in the EU form one big group). Efficiencies (left panel) and normalized spread (right panel) are shown as functions of α .

6.1.1 Partition P0

Under P0, there is only one group. There are thus no off diagonal elements ϱ^G that may be varied.

Under P0, proposals have a utility socket that changes from proposal to proposal, but is the same for each person from each state for a specific proposal. Results are shown in Fig. 3. If we compare to our default aggregate model, we observe that the decision rules are more efficient for all rules shown. The high efficiencies can be explained using Eq. (10): If the utilities \bar{V}_I for states are positively correlated, voting profiles under which most states cast the same vote carry large probabilistic weight. Profiles with many states voting yes yield acceptance under most of the rules considered here, and indicate a $\bar{v} > 0$. Thus, losses due to the decision rules are small. Likewise, profiles with many states voting no carry large probabilistic weight, yield rejection under most of the rules considered here, and indicate a $\bar{v} < 0$. Thus, losses due to missed positive utilities are small.

For many rules, the normalized spread in utilities is smaller than under the default aggregate model (right panel). Part of the explanation is that the normalizing $E[\bar{V}\theta(\bar{V})]$ is larger now, but this does not fully account for the new results for the normalized spread.

As far as the ranking of rules is concerned, there is not much change. A rule close to Penrose 50 maximizes expected utility, and, for thresholds $t = .5, .6$, the normalized spread minimizes around $\alpha = .5$ in the α -range considered. We conclude that the recommendations for the default aggregate model prove stable under the modification that we have considered. In particular, a rule close to $(\alpha = .5, t = .5)$ does very well on both welfarist standards.

6.1.2 Partition P1

Under P1, there are two groups, where one group comprises the four largest states (large states, for short, in this subsection and in Subsec. 6.2.2). This group has more population, but less states than the other group with the small states (see Table 1 for details).

In the first row of Fig. 4, we set σ^G at 2×10^{-4} and ϱ^G at 0. Utilities for people from different groups are thus independent. Again, the effects are not too dramatic. The efficiencies go up in comparison to the default aggregate model for all rules shown. Decision




Figure 4: Results for modifications of the aggregate model under partition P1 (two groups with large and small states, respectively).

rules with very high thresholds still yield significantly lower expected utilities than rules with a threshold of .5. The expected utility curve for $t = .5$ has still a maximum, which is now at $\alpha = .7$. This is also the overall maximum in the (α, t) -space.

The normalized spreads cover approximately the same range of values as for $\sigma^G = 0$. There is still a significant local minimum in all curves shown. The minima are shifted, comparing to the default aggregate model. Since the peak in the expected utility and the minimum in the measure of inequality shift in the same direction, we can come up with a political recommendation that does very well on both welfarist standards: a rule in the vicinity of $(\alpha = .6, t = .5)$.

For the second row of Fig. 4, ϱ^G is set at its maximal possible value. The utilities for people from different groups are now positively correlated. The results resemble the results that we have shown for partition P0. This is no surprise, since in both cases, positive dependencies between utilities for people from different states abound.

The case of $\varrho^G < 0$ is more interesting, since the groups compete with each other: If people from one group are proposed some good, people from the other group will fairly often be proposed some costs. As a consequence, some utilities from proposals are negatively dependent or negatively correlated. Results for the minimal possible value of ϱ^G are shown in the third row of Fig. 4. Low values of α produce very small efficiencies. For higher values of σ^G and for $\varrho^G = -(\sigma^G)^2$, $E[\bar{U}]$ is even below zero for some rules. The reason is that, under P1, for small values of α , the sum of weights for the the group with fewer people (the group with the small states) is larger than the thresholds $t = .5, .6, .8$. Proposals that are accepted thus typically put benefits (costs) on the people from the small (large) states. Since, altogether, more people live in the group with the large states, the net expected utility, $E[\bar{U}]$, is negative for $t = .5, .6, .8$. For the same reason, the spread in the utilities is very large for rules with thresholds $t = .5, .6$ and $\alpha \approx 0$. As α increases, things mostly become better, though, for all thresholds shown and on both standards. We have a maximum expected utility around $(\alpha = .9, t = .5)$ (see Table 2). The measure of inequality for $t = .5$ has a very sharp minimum. Unfortunately, it is not at .9, which means that there is no rule that is very good regarding both welfarist standards.

In order to check the stability of our results we also consider variations of partition P1. If the group with the small states is split into two subgroups, where one subgroup contains states with less than 10m people, the results are similar as under P1. Under a second variation, call it P1', Italy is counted as a small state such that the group with the large states does not have the majority of the people any more. For the minimal possible ϱ^G , the results are qualitatively different between P1 and P1'. For instance, whereas, under P1, equal weights ($\alpha = 0$) yield less efficiency than proportional weights ($\alpha = 1$) for $t = .5, .6, .8$, it is exactly the other way round under P1'. We conclude that, if there are negative dependencies in the model and if group membership is defined in terms of population, it is very important to know which group has the majority of the population.

6.1.3 Partition P2

Under partition P2, group membership is not defined in terms of population, but rather of geography. As one can see from Tab. 1, the group with the western states has about 75% of the EU population, but less states than the other group.

The most noteworthy results can be seen from Fig. 5. In the first row, $\sigma^G = 2 \times 10^{-4}$ and $\varrho^G = 0$. The efficiencies are quite high. The curve for $t = .5$ is still larger than the other curves for any value of α shown. The overall maximum is still the Penrose 50 rule.




Figure 5: Results for modifications of the aggregate model under partition P2 (western vs. eastern states).

The measure of inequality is very different from everything we have observed for the aggregate model and its modifications so far. For small values of α , the normalized spread is small; indeed, for each t shown, it is smaller than under P1. Then, for most parts of the curves $t = .5, .6$, the inequality increases as a function of α .

Why don't the expected utilities $E[U_I]$ almost equalize around $\alpha = .5$ under this model for $t = .5, .6$ as they do under the default aggregate model? The reason is as follows: For each group, there is an α -value in the range $\alpha \in [.5, .7]$ at which the $E[U_I]$ come very close for states I *within the group*. But for these values of α , people from the western states are better off than those from the eastern ones, and the $E[U_I]$ -value at which the expected utilities for people from the western states come very close is significantly higher than it is for the people from the eastern states. Part of the reason should be that, for $\alpha \geq .5$, the western states have a significantly larger sum of weights than the eastern states do.

In order to obtain a low measure of inequality, small values of α are required – at least for $t = .5, .6$. Thus, for this partition and this model, it is very difficult to be very good on both welfarist standards at the same time.

If an additional $\rho^G > 0$ is introduced, the efficiencies are often higher, but there is not much qualitative change comparing to $\rho^G = 0$. However, if the ρ^G -parameter is sufficiently high, the measure of inequality changes qualitatively. The curves for $t = .5, .6$ minimize close to $.5$ again (second row in Fig. 5). This makes it easier to be very good on both welfarist standards at the same time.

The last row shows results for the minimal possible $\rho^G < 0$. In this case, the curves for efficiency start at comparably low values. As under P1, voting rules with a small α are not efficient, since, under P2, the group with the larger population has less states and therefore less weights, and the majority of people will often have to pay for the benefits of people from the eastern states. As soon as α is sufficiently high, this changes very quickly. However, as α increases, the inequality mostly increases as well, at least if $t = .5, .6$. So it is again not possible to be very good on both welfarist standards at the same time.

6.2 Modifications of the interest group model with dependencies

Let us now consider modifications of the interest group model. The qualitative picture is as follows: On each particular proposal, people from the same state are proposed the same utility. This utility has two parts that are specific of the group and the state, respectively. As an example, we will set $\sigma^G = .4$. Thus, σ^G is 40% of the standard deviation for \bar{V}_I for all states I under the default interest group model. For Ireland, the change in the standard deviation is exactly the same as under the modifications of the aggregate model. The idea is that the utility contributions from the group utility tendencies χ_I^G are smaller than the contributions for the national utility tendencies, but that both contributions are of the same order of magnitude. The standard deviation for the proposed utilities \bar{V}_I is then $\sqrt{1 + .4^2}\sigma^S \approx 1.08\sigma^S = 1.08$. For some of the models investigated below, we also changed the value of σ and found qualitatively similar results.

6.2.1 Partition P0

Results for the partition P0 are shown in Fig. 6. From the left panel we observe that the efficiencies are larger under this model than under the default interest group model for all rules shown. The explanation why the efficiencies are high is the same as under the P0 modification of the aggregate model that we have considered.




Figure 6: Results for a modification of the interest group model under P0 (the states in the EU form one big group).

As under our aggregate model modification P0, the threshold $t = .5$ yields always higher expected utility than the other thresholds shown. As functions of α , the $E[\bar{U}]$ curves are flatter for $t = .5, .6$. The overall maximum for the expected utility is around ($\alpha \approx 0.66, t = 0.5$) (and not at $\alpha = 1, t = .5$ any more), but there is not a pronounced peak.

The normalized spread for the most part becomes larger, as α increases. Because of the symmetry of the model, $\alpha = 0$ still yields minimum (and zero) inequality, independently from the value of t . All in all, a reasonable welfarist compromise may consist in ($\alpha = 0, t = .5$) for this model. This yields zero inequality and a relatively high expected utility.

6.2.2 Partition P1

We now turn to modifications of the interest group model, under which large and small states form a group, each (P1). States from the group with larger states are again called large states, etc. in this subsection. Results for the partition can be seen from Fig. 7.

For $\sigma^G = .4$ and $\rho^G = 0$ (first row), all voting rules shown are again more efficient than under the default interest group model. Also, for $t = .5, .6$ the α -dependence of the curves is not as strong any more as under the default interest group model. The location of the maximum efficiency has shifted from ($\alpha = 1, t = .5$) (default) to ($\alpha = .88, t = .5$).

The results for the normalized spread are completely different from the results under the default interest group model. Equality is not established at $\alpha = 0$ any more. Rather, each curve for $t \in \{.5, .6, .8\}$ has a local minimum around $\alpha = .5$. For higher values of σ^G and $\rho^G = 0$, the minimum shifts further to the right. As far as α is concerned, it is easier to do reasonably well on both standards than it is under the default interest group model, since the maxima in the $E[\bar{U}]$ vs. α curves and the minima in the \ln vs. α curves are closer to each other than under the default interest group model. Unfortunately, the standards pull into different directions regarding the threshold.

Let us partly explain some of these results for $t = .5, .6$: For small α -values, the sum of weights from the group with the small states exceeds the threshold. Thus, under the probability model under consideration, the people from small states are likely to get what is in their interest. The votes of the large states, on the opposite, will often not make a difference for the outcome of the vote, and people from the large states are not that likely to get what is in their interest. This should at least be part of the reason why people




Figure 7: Results for modifications of the interest group model under P1 (large vs. small states).

from the large states receive significantly less benefits, and the spread in expected utilities is comparably high. As α increases, the large states are assigned higher weights. As a consequence, as α increases, the benefits increase for people from the group with the large states, whereas they decrease for people from many of the small states, at least if α is sufficiently large. Since the large states have the majority of the population, this is good for the whole federation: $E[\bar{U}]$ grows, as α increases, and at the same time inequality goes down. But for very high values of α , the situation for $\alpha = 0$ is reversed: Any person from any large state takes more profits than any person from any small states does.

Let us now switch on a maximal possible $\varrho^G = .16$ (second row). The efficiency curves for $t = .5, .6$ are very flat. Efficiencies are higher for $\varrho^G = .16$ than for $\varrho^G = 0$ for each rule shown. As a function of α , the normalized spread starts with small values and mostly becomes larger for all t -values shown. Altogether, this situation is quite favorable to the federation, and the choice of a decision rule is not that critical any more because the flatness of some efficiency curves. Note that there is qualitative similarity with the results for the P0 modification of the interest group model that we have considered.

Things are different, if we keep σ^G at .4 and set ϱ^G at the minimal possible value. As the third row in Fig. 7 shows, the efficiencies are lower than under $\varrho^G = 0$ for all rules shown, whereas the normalized spread is often very high. For the low efficiencies at $\alpha = 0$, the same explanation as under the corresponding P1 modification of the aggregate model applies. Rules with $\alpha \approx .7$ do comparably well on both welfarist standards for many thresholds t . But the choice of a threshold is still dilemmatic.

Our results for P1 are again quite stable, if the group with the small states is split into two groups as before (cf. end of Subsec. 6.1.2). There is more change in the results, if we compare to a partition under which the group with the large states does not have the majority of the people (P1'), particularly for the minimal ϱ^G . One difference is efficiency, where the curves for $t = .5, .6$ are not that steeply increasing functions of α any more.

6.2.3 Partition P2

Consider now modifications of the interest group model under partition P2, where the western and the eastern states form groups, each (see Tab. 1 for the partition and Fig. 8 for results).

If the utilities for people from different groups are independent ($\varrho^G = 0$, first row), the efficiencies are higher than they are under the default interest group model for all rules shown. At the same time, the efficiency curves are flatter for $t = .5, .6$, and the maximum of the $t = .5$ curve has shifted to the left. The maximum efficiency is now found at $(\alpha = .74, t = .5)$. But since this curve is comparably flat, the efficiency is not that sensitive to α any more.

The shape of the In vs. α curves reminds one of the results for the default interest group model. However, there is not perfect equality at $\alpha = 0$ any more at least for $t = .5, .6$. The reason is as follows: For $\alpha = 0$, the group with the eastern states has a larger sum of weights in the assembly, for which reason people from the eastern states take slightly larger profits than people from the western states do – hence the inequality. For the most part, In is an increasing function of α for all thresholds shown. Thresholds around .5 are worst for inequality for most values of α .

If we consider both welfarist standards, the choice of a decision rule is a bit easier than it is under the default interest group model. Particularly, the efficiency under $t = .5$ is not that sensitive to the α -value any more. But one has still to weigh between the different standards.




Figure 8: Results for modifications of the interest group model under P2 (western vs. eastern states).

If positive correlations between utilities for different groups are introduced and if $\varrho^G > 0$ is set at its maximal possible value (second row), the curves for the efficiencies become even higher. Also, for $t = .5, .6$ they are very flat. The local minima of In for $t = .5, .6$ seem still to be at $\alpha \neq 0$, but they are less pronounced. Under this model, $t = .5$ and that α -value at which In has a minimum in the $t = .5$ curve, define a rule that does reasonably well on both standards: The expected utility is comparably high, and inequality is very small.

It is very different for a parameter choice in which the utilities from proposal are negatively dependent for people from different groups. In the third row in Fig. 8, we consider the minimal ϱ^G compatible with $\sigma^G = .4$. The curves for efficiency cover a broad range of efficiency values. For each t , they mostly increase as a function of α . The overall maximum in our scans is obtained at $(\alpha = 1, t = .5)$. For larger σ^G and a minimal ϱ^G , even negative expected utilities are possible. Unfortunately, for each t , the measure of inequality is also mostly an increasing function of α with a minimum close to $\alpha = 0$, so minimizing In and maximizing efficiency is not possible at the same time.

In order to scrutinize the transition between the modifications of the interest and the aggregate models, we set $\sigma^P = 1$ and $\sigma^S = .01$. Roughly following the procedure how we set σ^G before, we determine σ^G such that for Ireland, $\sigma^G \approx .4\sqrt{(\sigma^S)^2 + \frac{1}{n_I}(\sigma^P)^2}$. As expected, the results that we get are very similar to our results for the interest group model both under the partitions P1 and P2.

6.3 A different measure of inequality

So far inequality has been measured with the standard deviation in the expected utilities of the people. One may object that the standard deviation is not an appropriate measure of inequality. As a reply, we have also considered a different measure of inequality, viz. the often used Gini coefficient (see, e.g., Sen 1997, pp. 29–31). For all the results shown in the viewgraphs in this paper, we have calculated the Gini coefficient with the expected utilities of the different people as incomes. For that, Eq. (3) in Karagiannis & Kovacevic' (2005), p. 120 has been used. An example that is fairly typical is shown in Fig. 9. When we compare the Gini coefficient in the left panel to the normalized spread in the right panel, then there are certainly differences. In particular, the curves for the different thresholds come much closer. The main reason seems to be that the Gini coefficient is not an absolute measure as is standard deviation, but rather is scale-invariant. That is, it remains invariant, if all incomes (or expected utilities) are multiplied with a positive number. The standard deviation would instead be multiplied with this number under this transformation.¹⁰ Part of the reason why the standard deviation is comparably small for $t = .8$ is therefore simply that the expected average utility $E[\bar{U}]$ is comparably low at this threshold (see Fig. 7). Under the Gini coefficient, which measures the inequality relative to the expected average utility, this effect is compensated, and the inequality is not that low for $t = .8$ any more. Something like this could also be obtained by normalizing the standard deviation by the mean, i.e. $E[\bar{U}]$.

But the differences in the measures are not that relevant for our purposes. The reason is that the shapes of the curves for a fixed t are similar under both measures. Consequently, the ordering of the different rules with a fixed t is not much affected. Also, what has to

¹⁰The measure that we have considered in most part of the paper, the normalized spread, is in fact normalized in some way. But it is not normalized by $E[\bar{U}]$, but rather by $E[\bar{V}\theta(\bar{V})]$, which is the same for all rules under a particular model. The effect that we are considering here is a normalization by $E[\bar{U}]$, which affects the results for the different rules in a different way even under the same model.




Figure 9: Inequality for a modification of the interest group model under P1 (large vs. small states) for parameter values specified in the left panel. Left panel: Gini coefficient. Right panel: normalized spread as used so far.

be done for obtaining a comparably small inequality is approximately the same for both measures of inequality in Fig. 9. All this is not only so in Fig. 9, but also for a considerable fraction of the results in the other viewgraphs from this paper. This indicates that, more generally, the recommendations that we give in this paper are not very much affected, if inequality is measured by the Gini coefficient (or a measure of relative inequality, more generally). Of course, ultimately, if one is interested in inequality as quantified by the Gini coefficient for a particular model, one has to scan the full (α, t) -space by looking at that very measure (although extreme values of t close to 0 or 1 may be left out, because they do not constitute plausible choices of a threshold anyway). For the purposes of this paper, though, and in view of the fact that the Gini coefficient is controversial as well, it was sufficient to show that an analysis with the Gini coefficient can be carried out and that it does not point to fundamentally different results.¹¹

7 Conclusions

In this paper, we have considered federations in which representatives of states vote on proposals that affect the people in the states. As an example, we have focused on the Council of Ministers in the European Union. We have evaluated alternative decision rules from a welfarist perspective. The idea is that a decision rule is better than another, if the resulting welfare distribution is better.

We have considered two welfarist standards, abbreviated as \mathcal{U} and \mathcal{E} . \mathcal{U} demands that the expected average utility for the people in the federation be maximum. \mathcal{E} demands that the inequality in the welfare distribution be minimal.

The welfare distribution that a decision rule produces given a particular model for the proposals has been calculated using the probabilistic framework set up by Beisbart et al. (2005) and Barberà & Jackson (2006). Differently from these works, however, we have not

¹¹One has to be careful in interpreting the Gini coefficient for our welfare distributions, since we have sometimes negative expected utilities (incomes). Another problem with the Gini coefficient is that it is more sensitive to errors due to the finite number of realizations in the simulations.

assumed that the utilities that are proposed for people from different states are stochastically independent – interstate utility dependencies have been allowed for.

Following Barberà & Jackson (2006), we have identified the decision rule that does best regarding standard \mathcal{U} in analytical terms, even for the case the utilities for people from different states are dependent. However, in general, this decision rule cannot be cast in simple terms. It is thus not suitable for practical purposes, or so we have argued.

We have therefore restricted ourselves to a family of weighted rules and to a specific class of probability models. The family of weighted rules can be parameterized using the threshold of acceptance, t , and the degree of degressive proportionality, α . Each model in our model class assigns each person a utility that has three parts: One is common to all people in a group of states, one is common to all people in a state only, one is perfectly personal. In this way interstate dependencies between the utilities that are proposed for the people arise.

Results have been obtained by means of simulations. We have started from the (default) aggregate and interest group models, which do not display interstate utility correlations (Beisbart & Bovens 2007). Under the default aggregate model, all utilities are independent, and the rule ($\alpha = .5, t = .5$) (Penrose 50) is very good on both standards. Under the default interest group model, the utilities for people from the same state are perfectly correlated, and the standards pull into different directions: Maximizing expected utility requires $t = .5$ and $\alpha = 1$, whereas equalizing the expected utilities for the different states requires a very low α and favors thresholds far away from $.5$.

We have then modified the default aggregate and interest group models by switching on interstate utility dependencies. Our dependencies are such that the marginal probability densities are identical for all people in the federation. The idea is that the proposals are not biased towards any person – at least in a stochastic sense. Of course, the results for the modifications of the aggregate and the interest group models depend on the details of the model parameters, and the aim of this paper has not so much been to obtain results for a specific model rather than showing how such models can be treated. But on the basis of the examples that we have investigated a few general things can be said.

First, interstate utility dependencies are important. If they are switched on, the welfare distributions that alternative decision produce given a particular model can change qualitatively. For a clear example compare the default aggregate model (Fig. 1) and our P2 modification with a minimal ϱ^G (third row of Fig. 5). Whereas, under the P2 modification, α -values close to zero minimize inequality for the t -values shown, the curves for inequality have a minimum around $\alpha = .5$ on the default aggregate model.

Second, whereas under the default models the welfare distribution that a decision rule produces is a function of the population sizes of the federation’s states only, this is not so any more under models with interstate utility correlations. In particular, answers to the following questions can be important for a welfarist evaluation of decision rules.

- Is group membership strongly correlated to population size?
- If the federation is partitioned into two groups of states, does the group with the higher population have more weights than the threshold requires?
- Are there positive or negative correlations between people from different groups?

Third, regarding dependencies, there is the following trend: If there are enough positive dependencies (i.e. under P0, $\sigma^G > 0$ and under P1 or P2, $\sigma^G > 0$ and ϱ^G maximal), the welfare distribution is less sensitive to the choice of a decision rule in the following sense:

There is a broader range of rules that are close to optimal in terms of expected utility, particularly for modifications of the aggregate model. At the same time, the expected utilities and the efficiencies (i.e. suitable normalized expected utilities) are higher than under no dependencies. This makes the choice of a decision rules easy and in a way less important. The choice of a decision rule is often more important for negative correlations.

Fourth, one can ask about the stability of the recommendations that we have mentioned for the default models above. The results for the aggregate model are fairly stable for the partitions that we called P0 and P1 above. That is, one can always find a decision rule that does well in terms of both standards. The α -value of this rule is sometimes larger than .5, the threshold can be set at .5. However, the results are not stable under partition P2, where \mathcal{E} requires a very low α close to zero.

The recommendations for the interest group model also remain rather stable, which means that the choice of a decision rule is still dilemmatic: \mathcal{U} requires one to chose something like $t = .5$ and a high α , whereas \mathcal{E} demands a threshold far away from .5 and a small α . Exceptions are the P1 modifications with $\varrho^G \leq 0$ that we have considered.

In spite of these results, much work remains to be done within the welfarist approach to decision rules. One task is to find empirically adequate probability models for proposals in real world federations. Maybe such models can be constrained by data on past voting profiles. The models that we have proposed in this work may serve as fitting models in such contexts. Another issue are budget constraints. In this paper, we have assumed that the utilities from proposals have zero means. But we do not exclude proposals that put a very high utility on the federation. This may seem unrealistic, since the budget of the federation is constrained. Our models with negatively dependent interests for people from different groups come closest to a federation with a budget constraint. Note that, according to our results, under such models there may be a qualitative difference in the way the decision rules produce welfare distributions, if we compare to the default models. But budget constraints require further examination. Other restrictions of the welfarist approach so far are discussed in Beisbart et al. (2005), Sec. 5. Our hope is that these restrictions can be overcome in future research.

Acknowledgments

We are grateful for useful comments by two anonymous referees for “Public Choice” and by a member of the editorial board. CB thanks the German Academic Exchange Service (DAAD) and the Center for Philosophy of Science at the University of Pittsburgh for support.

References

- Barberà, S. & Jackson, M., *On the Weights of Nations: Assigning Voting Weights in a Heterogeneous Union*, *Journal of Political Economy* **114** (2006), 317–339.
- Beisbart, C. & Bovens, L., *Welfarist Evaluations of Decision Rules for Boards of Representatives*, *Social Choice and Welfare* **29** (2007), 581–608.
- Beisbart, C. & Hartmann, S., *Welfarism and the Assessment of Social Decision Rules*, in: *Proceedings of the 1st International Workshop on Computational Social Choice* (Endriss, U. & Lang, J., eds.), 2006, pp. 35–48.
- Beisbart, C., Bovens, L., & Hartmann, S., *A Utilitarian Assessment of Alternative Decision Rules in the Council of Ministers*, *European Union Politics* **6** (2005), 395 – 418, Appendix online under http://www.uni-konstanz.de/eup/iss_64.htm.

- Bovens, L. & Hartmann, S., *Welfare, Voting and the Constitution of a Federal Assembly*, in: *Reasoning, Rationality and Probability* (Galavotti, M. C., Scazzieri, R., & Suppes, P., eds.), CSLI Publications, Stanford, 2007.
- Chamberlain, G. & Rothschild, M., *A Note on the Probability of Casting a Decisive Vote*, *Journal of Economic Theory* **25** (1981), 152–162.
- Chakravarty, S. R., *The Variance as a Subgroup Decomposable Measure of Inequality*, *Social Indicators Research* **53** (2001), 79–95.
- Coelho, D., *Maximin Choice of Voting Rules for Committees*, *Economics of Governance* **6** (2005), 159–175.
- Crain, W. M., Messenheimer, H. C., & Tollison, R. D., *The Probability of Being President*, *The Review of Economics and Statistics* **75** (1993), 683 – 689.
- Dancy, J., *Ethics without Principles*, Oxford University Press, Oxford, 2004.
- Fleurbaey, M., *One Stake One Vote*, working paper (2009).
- Felsenthal, D. S. & Machover, M., *The Measurement of Voting Power: Theory and Practice, Problems and Paradoxes*, Edward Elgar, Cheltenham, 1998.
- Felsenthal, D. S. & Machover, M., *Enlargement of the EU and Weighted Voting in its Council of Ministers*, Voting Power Report 01/00, London School of Economics and Political Science, Centre for Philosophy of Natural and Social Science, London; downloadable from <http://www.lse.ac.uk/vp>, 2000.
- Gelman, A., Katz, J. N., & Bafumi, J., *Standard Voting Power Indexes Don't Work: An Empirical Analysis*, *British Journal of Political Science* **34** (2004), 657 – 74.
- Good, I. J. & Mayer, L. S., *Estimating the Efficacy of a Vote*, *Behavioral Science* **20** (1975), 25–33.
- Karagiannis, Elias & Kovacevic', M., *A Method to Calculate the Jackknife Variance Estimator for the Gini Coefficient*, *Oxford Bulletin of Economics and Statistics* **62** (2000), 119 – 122.
- Laruelle, A. & Valenciano, F., *Assessing Success and Decisiveness in Voting Situations*, *Social Choice and Welfare* **24** (2005), 171 – 197.
- Schweizer, U., *Calculus of Consent: A Game-theoretic Perspective*, *Journal of Institutional and Theoretical Economics* **146** (1990), 28–54.
- Sen, A. *On Economic Inequality*. Expanded edition with a substantial annexe by James E. Forster and A. Sen, Clarendon Press, Oxford, 1997.