Price of Fairness in Budget Division and Probabilistic Social Choice

Marcin Michorzewski[1] **Dominik Peters**[2] **Piotr Skowron**[1]
[1]University of Warsaw [2]Carnegie Mellon University
marcin.michorzewski@gmail.com, dominikp@cs.cmu.edu, p.skowron@mimuw.edu.pl

Abstract

A group of agents needs to divide a divisible common resource (such as a monetary budget) among several uses or projects. We assume that agents have approval preferences over projects, and their utility is the fraction of the budget spent on approved projects. If we maximize utilitarian social welfare, the entire budget will be spent on a single popular project, even if a substantial fraction of the agents disapprove it. This violates the individual fair share axiom (IFS) which requires that for each agent, at least $1/n$ of the budget is spent on approved projects. We study the price of imposing such fairness axioms on utilitarian social welfare. We show that no division rule satisfying IFS can guarantee to achieve more than an $O(1/\sqrt{m})$ fraction of maximum utilitarian welfare, in the worst case. However, imposing stronger group fairness conditions (such as the core) does not come with an increased price, since both the conditional utilitarian rule and the Nash rule match this bound and guarantee an $\Omega(1/\sqrt{m})$ fraction. The same guarantee is attained by the rule under which the spending on a project is proportional to its approval score. We also study a family of rules interpolating between the utilitarian and the Nash rule, quantifying a trade-off between welfare and group fairness. An experimental analysis by sampling using several probabilistic models shows that the conditional utilitarian rule achieves very high welfare on average.

1 Introduction

Suppose a group of agents needs to divide a common budget among different projects, and we wish to vote over this division. For example, a city might wish to allocate a part of its spending according to the preferences of its residents over different uses ("participatory budgeting"; Cabannes, 2004). Or participants of a workshop might decide how to divide the workshop time among different discussion topics. We are interested in settings where the common budget is perfectly divisible, and each project can receive an arbitrary fraction of the resource. Thus, an outcome of the voting process can be seen as a vector of non-negative numbers, one for each project, summing up to 1. Each agent is assumed to have preferences over these distributions. This framing suggests that we can use *probabilistic social choice functions* (Brandt 2017) to guide the voting process: a probabilistic social choice function takes as input a profile of preferences, and returns a probability distribution over projects, which we interpret as a division of the budget.

Arbitrary preferences over the set of distributions can be complicated and difficult to elicit, so the literature usually assumes that preferences come from some restricted class, such as von Neumann–Morgenstern preferences (i.e., additive and linear), or given by ordinal rankings over the projects which are then extended to preferences over distributions in some natural way. We will use *approval voting*: agents indicate, for each project, whether they approve or disapprove it, and we assume that the utility of an agent is equal to the fraction of the budget spent on approved projects.[1] This setting has been introduced by Bogomolnaia, Moulin, and Stong (2005). Approval ballots make it easy for agents to express their preferences, and they are easy to elicit. The approval-based setup is well-studied and many rules and fairness notions have been proposed.

Given agents' approval preferences, a natural approach is to pick the budget division that maximizes utilitarian social welfare (i.e., the sum of agents' utilities). Suppose there are ten agents, six of which approve $\{a\}$, and four approve $\{b\}$. Then spending 100% of the budget on a maximizes social welfare, since shifting spending from b to a increases welfare. In general, the utilitarian rule will always spend the entire budget on projects with the highest approval score. While this behavior may be desirable in some contexts, often it is undesirable to completely ignore minority interests. In the kinds of settings mentioned in the beginning, it seems more appropriate to spend 60% on a and 40% on b.

The focus of several early papers about aggregation rules for participatory budgeting was social welfare maximization (e.g., Goel et al., 2019; Benade et al., 2017). Indeed, it seems desirable for the budget division to have high utilitarian social welfare, so that voters have, on average, high satisfaction. But this may not be the only goal, and several recent papers have worked on formalizing notions of fairness and proportionality in this setting (e.g., Aziz, Lee, and Talmon, 2018; Fain, Munagala, and Shah, 2018; Talmon and Faliszewski, 2019;

[1]Thus, we do not allow decreasing or increasing returns, nor do we consider substitutes and complements.

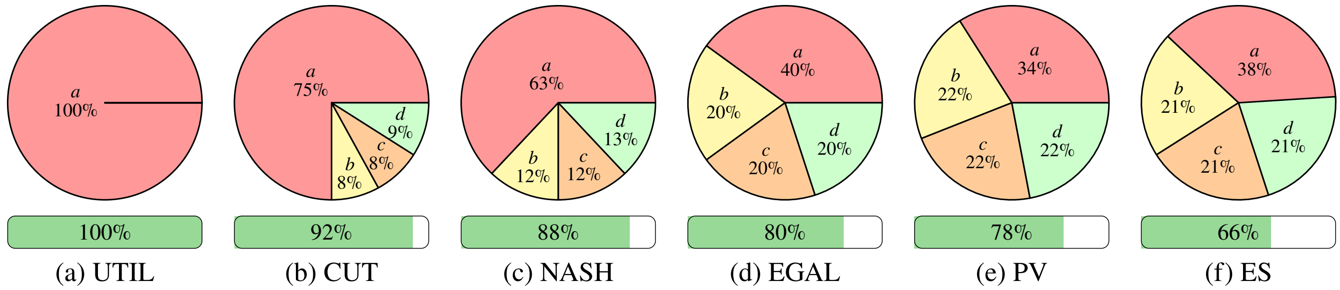

Figure 1: The voting rules we study evaluated on the profile $I = (ab, ac, ad, bcd)$, with the output distribution drawn as a pie chart. For each distribution, the green bar indicates its utilitarian social welfare as a fraction of the optimum. For $m = 4$, our theoretical guarantees predict that (b), (c), and (e) achieve at least 50%, and that (d) achieves at least 25% on every profile.

Airiau et al., 2019). In our divisible approval-based setting, fairness notions have been proposed by Bogomolnaia, Moulin, and Stong (2005), Duddy (2015), Fain, Goel, and Munagala (2016), Aziz, Bogomolnaia, and Moulin (2019) and Brandl et al. (2019). One example is *individual fair share*, which requires that every agent's utility is at least $1/n$, where n is the number of agents. Other notions require that groups are fairly represented: intuitively, a group of α% of the voters should control how an α fraction of the budget is spent.

Clearly, if a rule is to satisfy such fairness axioms, it cannot simultaneously maximize utilitarian welfare. But we can hope that there are rules which are fair while not losing too much in welfare terms. Figure 1 illustrates the outputs of five previously studied voting rules on an example profile with four voters and four projects, and calculates the utilitarian welfare achieved on this example as a fraction of the optimum. Duddy's (2015) *conditional utilitarian* rule (CUT) achieves 92%, while maximizing Nash welfare achieves 88%.

Our aim in this paper is to formally quantify the trade-off between fairness axioms and maximization of utilitarian social welfare. In particular, we ask what fraction of the optimum utilitarian welfare can be achieved by a rule satisfying, say, IFS, in the worst case over all profiles (related notions in other contexts are often referred to as the "price of fairness"; see, e.g., Bertsimas, Farias, and Trichakis, 2011; Caragiannis et al., 2012). We also study specific voting rules that have been proposed for this setting, and likewise ask whether these rules approximate utilitarian social welfare. In most cases, we obtain asymptotically tight bounds.

Our first result shows that a rule satisfying IFS can, at most, guarantee to provide a $\frac{2}{\sqrt{m}}$ fraction of optimum utilitarian welfare, where m is the total number of projects under consideration. This bound is obtained from analyzing a simple family of examples. Naively, one would expect that imposing stronger fairness axioms (such as group fair share, or the core) would further limit the obtainable utilitarian welfare, but this is not the case: We show that both the CUT rule and the Nash rule asymptotically match the bound, and provide at least a $\frac{2}{\sqrt{m}} - \frac{2}{m}$ fraction of the optimum on every profile, even though both rules satisfy fairness axioms much stronger than IFS. For the Nash rule, the proof of our positive result in fact uses its fairness properties: we show that any rule satisfying either the core or the average fair share axiom (Aziz, Bogomolnaia, and Moulin 2019) guarantees an $\Omega(\frac{1}{\sqrt{m}})$ fraction of optimum welfare. To obtain our positive result for the CUT rule, we use Brandl et al.'s (2019) notion of *implementability* to prove that, on every profile, the utilitarian welfare of the CUT rule exceeds that of the Nash rule.

We also study some other rules. For two rules (the egalitarian rule and the uncoordinated equal shares rule) we find that their guarantees are worse than that of CUT and the Nash rule. For the rule that funds each project in proportion to its approval score, we again obtain a guarantee of $\frac{2}{\sqrt{m}}$, which is somewhat surprising since this rule, in contrast to the others, does not attempt to be fair to voters. Instead, it can be seen as being fair to candidates.

We then study a family of rules that interpolate between the utilitarian rule and the Nash rule. Members of this family can guarantee a higher fraction of the optimum social welfare. While Nash satisfies the strongest fairness properties, there are other rules in the family which guarantee high average welfare to large groups. Thus, in this family, we quantify a trade-off between utilitarian welfare and group fairness.

We complement our theoretical results with an experimental analysis, using random preference models based on impartial culture, mixtures of Mallow's models, and on a spatial (Euclidean) model. We find that, in the average case, all rules outperform their worst-case guarantee. The CUT rule achieves a remarkably high utilitarian welfare, exceeding 95% for many parameter values.

2 Preliminaries

An *instance* or *profile* is a triple $I = (N, O, A)$, where $N = \{1, \dots, n\}$ is a set of voters, $O = \{o_1, \dots, o_m\}$ is a set of projects, and $A = (A_1, \dots, A_n)$ is a list of *approval sets*, i.e., subsets of O, one for each voter. The approval set A_i of voter $i \in N$ contains those projects that i approves or finds acceptable; we assume it is non-empty. Let $\mathcal{I}_m$ denote the set of all instances with $|O| = m$. For a project o_j, we write $N(o_j) = \{i \in N : o_j \in A_i\}$ for the set of voters who approve o_j. The *approval score* of o_j is $|N(o_j)|$.

A *distribution* over O is a vector $\boldsymbol{x} \in [0, 1]^m$ with $\sum_{j=1}^{m} x_j = 1$. We write $\Delta(O)$ for the set of distributions. We can interpret $\boldsymbol{x}$ as either a lottery or as a division of a fixed-size budget among the projects. Given an instance I and a distribution $\boldsymbol{x}$, the *utility* of voter $i \in N$ for $\boldsymbol{x}$ is $u_i(\boldsymbol{x}) = \sum_{o_j \in A_i} x_j$, i.e., the total fraction that $\boldsymbol{x}$ spends on projects approved by i. The *(utilitarian) social welfare* of $\boldsymbol{x}$ is $\mathrm{sw}(I, \boldsymbol{x}) = \sum_{i \in N} u_i(I, \boldsymbol{x})$. It is useful to note that the maximum feasible social welfare $\mathrm{sw}^*(I) = \max_{\boldsymbol{x} \in \Delta(O)} \mathrm{sw}(I, \boldsymbol{x})$ is attained by the distribution placing 100% on the project with highest approval score, so that $\mathrm{sw}^*(I) = \max_{o_j \in O} |N(O_j)|$. Finally, let us define the *normalized social welfare* of $\boldsymbol{x}$ as $\widehat{\mathrm{sw}}(I, \boldsymbol{x}) = \mathrm{sw}(I, \boldsymbol{x}) / \mathrm{sw}^*(I)$, the fraction of the optimum social welfare achieved by $\boldsymbol{x}$.

Voting rules A (probabilistic) voting rule f is a function which assigns to every instance I a set of distributions over O (usually the output contains only one distribution, but several could be tied). The voting rules discussed in this paper are:

- The *utilitarian rule* (UTIL) which selects all distributions $\boldsymbol{x}$ maximizing $\mathrm{sw}(I, \boldsymbol{x})$.
- The *conditional utilitarian rule* (CUT) which selects the distribution $\frac{1}{n} \sum_{i \in N} \boldsymbol{x}_i$, where for each $i \in N$, $\boldsymbol{x}_i$ is the uniform distribution over the projects in A_i that have the highest approval score.
- The *Nash rule* (NASH) which selects all distributions $\boldsymbol{x}$ maximizing $\prod_{i \in N} u_i(\boldsymbol{x})$, or equivalently $\sum_{i \in N} \log u_i(\boldsymbol{x})$.
- The *egalitarian rule* (EGAL) which selects all $\boldsymbol{x}$ maximizing $\min_{i \in N} u_i(\boldsymbol{x})$ (possibly breaking ties using leximin).
- The *point voting rule* (PV) which selects $\boldsymbol{x}$ where x_j is proportional to the approval score $|N(o_j)|$, for each $o_j \in O$.
- The *uncoordinated equal shares rule* (ES) which selects the distribution $\frac{1}{n} \sum_{i \in N} \boldsymbol{x}_i$, where for each $i \in N$, $\boldsymbol{x}_i$ is the uniform distribution over A_i.

These rules were introduced and studied previously (Aziz, Bogomolnaia, and Moulin 2019; Duddy 2015; Brandl et al. 2019). Figure 1 shows the rules evaluated on an example.

3 Fairness Axioms

Let us begin our study of the impact of imposing fairness axioms on utilitarian welfare by stating several axioms that have been proposed. Fix an instance I with n voters and a distribution $\boldsymbol{x} \in \Delta(O)$. Then $\boldsymbol{x}$ satisfies

- *individual fair share* (IFS) if $u_i(\boldsymbol{x}) \geqslant 1/n$ for all $i \in N$;
- *group fair share* (GFS) if for every $S \subseteq N$, we have $\sum_{o_j \in \bigcup_{i \in S} A_i} x_j \geqslant |S|/n$;
- *implementability* if we can write $\boldsymbol{x} = \frac{1}{n} \sum_{i \in N} \boldsymbol{x}_i$ for some distributions $(\boldsymbol{x}_i)_{i \in N}$ such that $x_{i,j} > 0$ only if $o_j \in A_i$;
- *average fair share* (AFS) if for every $S \subseteq N$ such that $\bigcap_{i \in S} A_i \neq \emptyset$, we have $\frac{1}{|S|} \sum_{i \in S} u_i(\boldsymbol{x}) \geqslant |S|/n$;
- the *core* if for every $S \subseteq N$, there is no vector $\mathbf{z} \in [0,1]^m$ with $\sum_{j=1}^m z_j = |S|/n$ such that $u_i(\mathbf{z}) > u_i(\boldsymbol{x})$ for all $i \in S$.

Intuitively, an implementable distribution is obtained by splitting the budget into pieces of size $1/n$, and letting each voter spend their piece on approved projects. AFS requires that *cohesive* coalitions (whose members approve at least one project in common) have high average welfare. The core requires that no group could spend their fair share of the budget in a way that each group member prefers to $\boldsymbol{x}$.

	UTIL	CUT	NASH	EGAL	PV	ES
IFS		+	+	+		+
GFS		+	+			+
impl.		+	+			+
AFS			+			
core			+			

Table 1: Voting rules, and the fairness axioms they satisfy

A voting rule f satisfies one of these notions if $f(I)$ satisfies it for all instances I. Table 1 shows the rules introduced in Section 2 and which of these fairness axioms are satisfied by them. The lower three axioms are pairwise logically independent; each of them implies GFS, and GFS implies IFS. For proofs and further discussion, see Aziz, Bogomolnaia, and Moulin (2019). For a discussion of implementability see Brandl et al. (2019) and Guerdjikova and Nehring (2014); CUT and ES satisfy implementability by definition, and the first-order conditions for NASH show it satisfies it too.

Our first result shows that imposing the weakest of the above fairness axioms leads, in the worst case over profiles, to a substantial loss of utilitarian social welfare.

Theorem 1. *For each $m \geqslant 2$, there exists an instance I on m projects such that for every distribution $\boldsymbol{x}$ that satisfies IFS, we have $\widehat{\mathrm{sw}}(I, \boldsymbol{x}) \leqslant \frac{2}{\sqrt{m}}$.*

Proof. Fix some $m \geqslant 2$. Write $m = k^2 + r$ for some $k \geqslant 1$ and $1 \leqslant r \leqslant k + 1$. Construct an instance I as follows. Take $n = k^2 + k$ voters. Voters $i = 1, 2, \ldots, k$ approve only o_1; for each $i \in \{k+1, k+2, \ldots, k+k^2\}$ voter i approves only o_{i-k+1}. Thus, o_1 is approved by k voters and k^2 other projects are each approved by a single voter.

Suppose $\boldsymbol{x}$ satisfies IFS for this instance. Then for each $j = 1, \ldots, k^2 + 1$, we must have $x_j \geqslant \frac{1}{n}$. It follows that $x_1 = 1 - \sum_{j=2}^m x_j \leqslant 1 - \frac{k^2}{n}$. Thus,

$$\mathrm{sw}(I, \boldsymbol{x}) = \sum_{i \in N} u_i(I, \boldsymbol{x}) \leqslant k \cdot \left(1 - \frac{k^2}{n}\right) + k^2 \cdot \frac{1}{n} = \frac{2}{1 + \frac{1}{k}}.$$

Now, $\mathrm{sw}^*(I) = k$ since o_1 is approved by k voters. Noting that $\sqrt{m} = \sqrt{k^2 + r} \leqslant \sqrt{(k+1)^2} = k + 1$, we have

$$\widehat{\mathrm{sw}}(I, \boldsymbol{x}) = \frac{\mathrm{sw}(I, \boldsymbol{x})}{k} \leqslant \frac{2}{k+1} \leqslant \frac{2}{\sqrt{m}}.$$

This completes the proof. □

Note that in the hard instance constructed in Theorem 1, every voter's approval set is a singleton. Intuitively, if voters are inflexible and single-minded, then the conflict between individual fairness and utilitarian welfare is strongest.

A priori, one would expect that imposing stronger fairness notions than IFS (such as GFS, AFS, or the core) would come at a greater cost to utilitarian welfare than imposing only IFS. As we will see, this need not be the case: there are rules (such as the Nash rule) which achieve the bound of Theorem 1 while satisfying those stronger properties.

Several of the fairness axioms we mentioned have the form of imposing lower bounds on the utility of individuals or of groups of agents. AFS, in particular, requires that cohesive groups of agents have a high level of utility on average. Thus, imposing AFS does not only impose a cost on welfare, but might also give a guarantee: the welfare of an AFS distribution cannot be too low. Remarkably, the guarantee provided by AFS asymptotically matches the lower bound for IFS.

Theorem 2. *Let I be an instance, and let $\boldsymbol{x}$ be a distribution satisfying AFS. Then* $\widehat{\mathrm{sw}}(I, \boldsymbol{x}) \geqslant \frac{2}{\sqrt{m}} - \frac{2}{m}$.

Proof. Consider the project in O with the highest approval score. Without loss of generality, assume it is o_1; let $a_1 = |N(o_1)|$. Applying AFS to the coalition $S = N(o_1)$, we see that the total utility $\sum_{i \in S} u_i(\boldsymbol{x})$ of members of $N(o_1)$ is at least $(a_1)^2/n$. Now, remove the voters from $N(o_1)$ from our further reasoning, and consider the project approved by the most voters from $N \setminus N(o_1)$; without loss of generality assume it is o_2. Let a_2 denote the number of voters from $N \setminus N(o_1)$ who approve o_2. By the same reasoning as before, we get that the total utility of the voters from $N(o_2) \setminus N(o_1)$ equals at least $(a_2)^2/n$. By applying the same reasoning recursively we get that that the total utility of the voters equals at least:

$$\mathrm{sw}(I, \boldsymbol{x}) \geqslant \sum_{i=1}^{m} \frac{(a_i)^2}{n}.$$

Further, since $\sum_{i=1}^{m} a_i = n$, we get that:

$$\mathrm{sw}(I, \boldsymbol{x}) \geqslant \frac{\sum_{i=1}^{m} (a_i)^2}{\sum_{i=1}^{m} a_i}, \text{ and hence } \widehat{\mathrm{sw}}(I, \boldsymbol{x}) \geqslant \frac{\sum_{i=1}^{m} (a_i)^2}{a_1 \sum_{i=1}^{m} a_i}.$$

Let us set $b_j = a_j/a_1$ for each $j = 1, \ldots, m$, and $s = \sum_{j=2}^{m} b_j$. From the Cauchy-Schwarz inequality we know that

$$\sum_{j=2}^{m} b_j^2 \geqslant \frac{s^2}{m-1}.$$

Thus, we have that:

$$\widehat{\mathrm{sw}}(I, \boldsymbol{x}) \geqslant \frac{\sum_{i=1}^{m} a_i^2}{a_1 \sum_{i=1}^{m} a_i} = \frac{1 + \sum_{j=2}^{m} b_j^2}{1 + \sum_{j=2}^{m} b_j} \geqslant \frac{m - 1 + s^2}{(m-1)(s+1)}.$$

When we consider the function $g(s) = \frac{m-1+s^2}{(m-1)(s+1)}$ for $s \in [0, m-1]$, we see that it has value 1 on both ends. Let us compare the derivative of the function to 0 to find the extremum (omiting the denominator to simplify the proof):

$$\begin{aligned} & 2s(m-1)(s+1) - (m-1)(m-1+s^2) = 0 \\ \iff & s^2 + 2s - (m-1) = 0 \end{aligned}$$

This gives $s = \sqrt{m} - 1$ (the second root is negative). Now,

$$\begin{aligned} g(\sqrt{m} - 1) &= \frac{m - 1 + m + 1 - 2\sqrt{m}}{\sqrt{m}(m-1)} = \frac{2 + \frac{2 - 2\sqrt{m}}{m-1}}{\sqrt{m}} \\ &\geqslant \frac{2 - \frac{2\sqrt{m}}{m}}{\sqrt{m}} = \frac{2}{\sqrt{m}} - \frac{2}{m}. \end{aligned}$$

Hence, $\widehat{\mathrm{sw}}(I, \boldsymbol{x}) \geqslant \inf_s g(s) \geqslant \frac{2}{\sqrt{m}} - \frac{2}{m}$, as required. ◻

The core similarly guarantees a minimum welfare level, though our bound is weaker by a factor of 2. The proof is similar to the proof of Theorem 2, using an idea also used in the context of committee elections and extended justified representation (Sánchez-Fernández et al., 2017, Thm. 6).

Theorem 3. *Let I be an instance, and let $\boldsymbol{x}$ be a distribution satisfying the core. Then* $\widehat{\mathrm{sw}}(I, \boldsymbol{x}) \geqslant \frac{1}{\sqrt{m}} - \frac{1}{m}$.

4 Guarantees for Voting Rules

In this section, we analyze the voting rules introduced in Section 2, to obtain guarantees on their normalized social welfare. Formally, for each number m of alternatives, we calculate the rules' *efficiency guarantee* (Lackner and Skowron 2019)

$$\kappa_{\mathrm{eff}}(f, m) = \min_{I \in \mathcal{I}_m} \widehat{\mathrm{sw}}(I, f(I)).$$

Efficiency of NASH

Based on Section 3, the efficiency guarantee of NASH is easy to determine. Since NASH satisfies IFS, by Theorem 1, we have $\kappa_{\mathrm{eff}}(\mathrm{NASH}, m) \leqslant \frac{2}{\sqrt{m}}$. On the other hand, since NASH satisfies AFS, by Theorem 2, we have $\kappa_{\mathrm{eff}}(\mathrm{NASH}, m) \geqslant \frac{2}{\sqrt{m}} - \frac{2}{m}$. Thus, we have proved the following:

Theorem 4. *The efficiency guarantee for* NASH *is between* $\frac{2}{\sqrt{m}} - \frac{2}{m}$ *and* $\frac{2}{\sqrt{m}}$.

Efficiency of CUT

Just as for NASH, we obtain an upper bound on the efficiency guarantee of CUT from Theorem 1, since CUT satisfies IFS. However, CUT fails AFS and the core, so we cannot use the results in Section 3 to obtain a lower bound. However, we can use them indirectly, by comparing the CUT rule to NASH:

Lemma 1. *For every instance I, we have* $\mathrm{sw}(I, \mathrm{CUT}(I)) \geqslant \mathrm{sw}(I, \mathrm{NASH}(I))$.

Proof. As suggested by Brandl et al. (2019), the CUT rule maximizes utilitarian welfare among implementable distributions: For every implementable $\boldsymbol{x}$, we have $\mathrm{sw}(I, \mathrm{CUT}(I)) \geqslant \mathrm{sw}(I, \boldsymbol{x})$. This is because, writing $\boldsymbol{x} = \frac{1}{n} \sum_{i \in N} \boldsymbol{x}_i$ like in the definition of implementability, maximization of social welfare requires that each $\boldsymbol{x}_i$ maximizes welfare, subject to the requirement that the support of $\boldsymbol{x}_i$ is contained in A_i. This is precisely what the CUT rule does. The result of the lemma follows if we choose $\boldsymbol{x} = \mathrm{NASH}(I)$, noting that $\mathrm{NASH}(I)$ is implementable (Brandl et al. 2019; Guerdjikova and Nehring 2014). ◻

Thus, the lower bound for NASH also applies to CUT.

Theorem 5. *The efficiency guarantee for* CUT *is between* $\frac{2}{\sqrt{m}} - \frac{2}{m}$ *and* $\frac{2}{\sqrt{m}}$.

Efficiency of EGAL

A notable feature of the egalitarian rule EGAL is that it does not attempt to be fair to groups. The egalitarian objective even treats different voters with the same approval set as if they were a single voter; thus, if we copy a voter many times, the output of EGAL does not change. This is in contrast to the behavior of (for example) NASH and CUT, where a voter gets more influence when copied.

Our proofs of the welfare guarantees of NASH and CUT were based on guarantees for groups. A similar argument cannot work for EGAL. Supporting this intuition, we find that the efficiency guarantee of EGAL is substantially worse.

Theorem 6. *The efficiency guarantee of* EGAL *is* $\frac{1}{m}$.

Proof. Let I be an instance. Observe that for $\boldsymbol{x} = (\frac{1}{m}, \dots, \frac{1}{m})$, we have $u_i(I, \boldsymbol{x}) \geqslant \frac{1}{m}$ for each voter $i \in N$. So the egalitarian welfare of $\boldsymbol{x}$ is at least $\frac{1}{m}$. Let $\boldsymbol{x}_{\mathrm{egal}}$ be a distribution returned by EGAL for I. Its egalitarian welfare must also be at least $\frac{1}{m}$, so $u_i(I, \boldsymbol{x}_{\mathrm{egal}}) \geqslant \frac{1}{m}$ for each $i \in N$. Hence,

$$\mathrm{sw}(I, \boldsymbol{x}_{\mathrm{egal}}) = \textstyle\sum_{i \in N} u_i(I, \boldsymbol{x}) \geqslant \frac{n}{m}.$$

Since for every distribution $\boldsymbol{x}$ we have $\mathrm{sw}(I, x) \leqslant n$, we get:

$$\widehat{\mathrm{sw}}(I, \boldsymbol{x}_{\mathrm{egal}}) \geqslant \tfrac{n}{mn} = \tfrac{1}{m}.$$

It follows that $\kappa_{\mathrm{eff}}(\mathrm{EGAL}, m) \geqslant \frac{1}{m}$.

For the lower bound, let us fix $k > 0$, and consider an instance with m projects and $k + m - 1$ voters, constructed as follows. The first k voters approve $\{o_1\}$. For each $i = 1, \dots, m-1$, voter $k+i$ approves $\{o_{i+1}\}$. Let $\boldsymbol{x}$ be a distribution. Then there exists some project o_i with $x_i \leqslant \frac{1}{m}$, and thus there exists a voter $i \in N$ with $u_i(\boldsymbol{x}) \leqslant \frac{1}{m}$. Thus, the egalitarian welfare of $\boldsymbol{x}$ is at most $\frac{1}{m}$, and this value is achieved only when $\boldsymbol{x} = \boldsymbol{x}_{\mathrm{egal}} = (\frac{1}{m}, \frac{1}{m}, \dots, \frac{1}{m})$. We have $\mathrm{sw}(I, \boldsymbol{x}_{\mathrm{egal}}) = \sum_{i \in N} u_i(I, \boldsymbol{x}) = \frac{k+m-1}{m}$. On the other hand, $\mathrm{sw}^*(I) = k$ (which is achieved for $\boldsymbol{x}^* = (1, 0, 0, \dots, 0)$). This leads us to:

$$\kappa_{\mathrm{eff}}(\mathrm{EGAL}, m) \leqslant \frac{k + m - 1}{mk} = \frac{1}{m} + \frac{1}{k} - \frac{1}{mk} \xrightarrow{k \to \infty} \frac{1}{m}.$$

Combining both directions, we obtain $\kappa_{\mathrm{eff}}(\mathrm{EGAL}) = \frac{1}{m}$. □

Efficiency of PV

Inspecting Table 1, we see that point voting (PV) does not satisfy any of our fairness axioms. A simple example is the profile $I = (abc, d, e)$ with three voters. PV returns the uniform distribution over $\{a, b, c, d, e\}$. This violates individual fair share (IFS) for the voter approving $\{d\}$, since d gets $\frac{1}{5} < \frac{1}{3}$. One interpretation of the rule PV is that it aims to be fair to *projects*, not voters: a project a receives more funding than b if and only if more voters approve a than approve b.

Interestingly, we find that despite the different philosophy behind PV, its worst-case efficiency guarantee matches the one we found for NASH and CUT.

Theorem 7. *The efficiency guarantee for* PV *is between* $\frac{2}{\sqrt{m}} - \frac{2}{m}$ *and* $\frac{2}{\sqrt{m}}$.

Proof. We first prove the lower bound. Let I be an instance with $O = \{o_1, \dots, o_m\}$. For $j = 1, \dots, m$, we write $n_j = |N(o_j)|$ for the approval score of o_j. Label the projects so that $n_1 \geqslant n_2 \geqslant \dots \geqslant n_m$. Then, the maximum social welfare is $\mathrm{sw}^*(I) = n_1$. Now, PV chooses the distribution $\boldsymbol{x} = \frac{1}{\sum_{j=1}^m n_j}(n_1, n_2, \dots, n_m)$. Then we have

$$\mathrm{sw}(I, \boldsymbol{x}) = \sum_{i \in N} \sum_{o_j \in A_i} \frac{n_j}{\sum_{k=1}^m n_k}$$

$$= \sum_{o_j \in O} \sum_{i \in N(o_j)} \frac{n_j}{\sum_{k=1}^m n_k} = \sum_{o_j \in O} \frac{(n_j)^2}{\sum_{k=1}^m n_k}$$

Hence, using the estimation from the second part of the proof of Theorem 2,

$$\widehat{\mathrm{sw}}(I, \boldsymbol{x}) = \sum_{o_j \in O} \frac{(n_j)^2}{n_1 \sum_{k=1}^m n_k} \geqslant \frac{2}{\sqrt{m}} - \frac{2}{m}.$$

For the upper bound, we observe that while PV fails IFS in general, it satisfies IFS on profiles in which every approval set is a singleton. Since the hard instance in the proof of Theorem 1 only contains singleton approval sets, the upper bound of Theorem 1 applies to PV. □

Efficiency of ES

The rule ES gives each voter a $1/n$ share of the resource to spend on approved projects. The same is true for CUT, where each voter uses their share in a utilitarian way, taking other voters' preferences into account. In contrast, under ES, there is no coordination, and voters ignore others' preferences.

While we do not give a tight estimate of the utilitarian efficiency of ES, we separate it from that of CUT: In the worst case, the normalized social welfare under ES is worse than under CUT. This underscores the value of coordination, and it shows that implementability alone is not enough to obtain the positive results we obtained for NASH and CUT.

Theorem 8. *The efficiency guarantee of* ES *is* $O\big(\frac{1}{m^{2/3}}\big)$.

Proof. Let us fix two integers, c and ℓ such that $c + \ell \leqslant m - 1$, and let us consider the following instance. The first $(m-1)\binom{m-1}{c}$ voters are divided into $m-1$ groups, and within each group, the voters are indexed with c-element subsets of $\{2, \dots m\}$. For each such a subset $S = \{i_1, \dots, i_c\}$, the voter with index S approves projects $o_{i_1}, \dots, o_{i_c}$ and o_1. Thus, each of the first $(m-1)\binom{m-1}{c}$ voters approves exactly $c+1$ projects, and each project other than o_1 is approved by $\binom{m-1}{c}c$ such voters. The remaining voters are divided into $m-1$ groups, each consisting of $\ell\binom{m-1}{c}$ voters; the voters from the i-th group approve project o_{i+1}. Thus, altogether o_1 is approved by $(m-1)\binom{m-1}{c}$ voters, and each other candidate is approved by $\binom{m-1}{c}c + \ell\binom{m-1}{c}$ voters. Clearly, we have

$$\begin{aligned} n &= (m-1)\tbinom{m-1}{c} + (m-1)\ell\tbinom{m-1}{c} \\ &= (m-1)(\ell+1)\tbinom{m-1}{c}. \end{aligned}$$

ES chooses the distribution $\boldsymbol{x} = (p_1, p_2, p_2, \dots, p_2)$, where

$$p_1 = \frac{(m-1)\binom{m-1}{c}}{(c+1)n} = \frac{1}{(c+1)(\ell+1)},$$

$$p_2 = \frac{\ell\binom{m-1}{c} + \frac{c}{c+1}\binom{m-1}{c}}{n} = \frac{\ell + 1 - \frac{1}{c+1}}{(m-1)(\ell+1)}.$$

Then the total utility achieved by this distribution equals:

$$(m-1)\binom{m-1}{c}p_1 + (m-1)p_2\left(\binom{m-1}{c}c + \ell\binom{m-1}{c}\right) = \binom{m-1}{c}\left(\frac{m-1}{(c+1)(\ell+1)} + \frac{(\ell+c)(\ell+1-\frac{1}{c+1})}{(\ell+1)}\right).$$

The highest possible utilitarian score is $(m-1)\binom{m-1}{c}$, so the guarantee is at most:

$$\frac{1}{(c+1)(\ell+1)} + \frac{(\ell+c)(\ell+1-\frac{1}{c+1})}{(m-1)(\ell+1)} < \frac{1}{c\ell} + \frac{\ell+c}{m-1} \leqslant \frac{1}{c\ell} + (1+\varepsilon)\frac{\ell+c}{m},$$

for $\varepsilon = \frac{1}{m-1} > 0$, which is close to 0 when $m \to \infty$. When we set $c = \ell = m^{(1/3)}$, we get the upper bound

$$\frac{1}{m^{(2/3)}} + (1+\varepsilon)\frac{2}{m^{(2/3)}} = \frac{3+2\varepsilon}{m^{(2/3)}}$$

for our guarantee. This completes the proof. □

5 A Trade-Off Between Group Fairness and Utilitarian Efficiency

In this section, we consider a family of rules. Let $f : [0,1] \to \mathbb{R}$ be a non-decreasing function. The f-UTIL rule selects the distributions $\boldsymbol{x}$ that maximize $\sum_{i\in N} f(u_i(\boldsymbol{x}))$. Thus, the rule transforms voter utilities using f and then maximizes the sum. If f is the identity, this is UTIL; if $f = \log$ then it is NASH. Bogomolnaia, Moulin, and Stong (2002) considered this family and showed that UTIL is its only member satisfying strategyproofness, and that NASH is its only member satisfying GFS (or the stronger AFS). Here, for well-behaved f, we will show that f-UTIL satisfies a weakening of AFS, and give a lower bound for its utilitarian efficiency.

Our first result gives a guarantee for the average welfare of groups of voters who approve at least one project in common.

Theorem 9. *Let f be a continuous non-decreasing function with a convex derivative. Let $\boldsymbol{x}$ be a lottery returned by f-UTIL. For every $S \subseteq N$ with $\bigcap_{i\in S} A_i \neq \emptyset$, we have:*

$$\frac{1}{|S|}\sum_{i\in S} u_i(\boldsymbol{x}) \geqslant (f')^{-1}\Big(\frac{nM}{|S|}\Big),$$

where $M = \max_{x\in[0,1]}(f'(x)x)$.

If we plug in $f = \log$ into this theorem, we confirm that NASH satisfies AFS. For another example, if we plug in $f(x) = \sqrt{x}$, we find that a cohesive group S will obtain an average welfare of at least $(|S|/n)^2$. Note that this is strictly weaker than the guarantee implied by AFS (except when $|S| = n$). Note also that larger groups S obtain a relatively stronger guarantee than smaller groups. More generally, for $f(x) = x^\alpha$ with $0 < \alpha < 1$, we get that a cohesive group S obtains average welfare of at least $(|S|/n)^{1/(1-\alpha)}$. As $\alpha \to 0$, this guarantee converges to full AFS.

Proof of Theorem 9. Let $\boldsymbol{x}$ and S be as in the theorem. For a contradiction, assume that $\frac{1}{|S|}\sum_{i\in S} u_i(\boldsymbol{x}) < (f')^{-1}\left(\frac{nM}{|S|}\right)$. The Lagrangian for the optimization of $\sum_{i\in N} f(u_i(\boldsymbol{x}))$ is

$$\mathcal{L}(\boldsymbol{x},\lambda) = \sum_{i\in N} f(u_i(\boldsymbol{x})) - \lambda\Big(1 - \sum_{j\in[m]} x_j\Big).$$

From the KKT conditions, $\frac{\partial\mathcal{L}}{\partial x_j} \leqslant 0$ for each $j \in [m]$, and so $\sum_{i\in N(o_j)} f'(u_i(\boldsymbol{x})) - \lambda \leqslant 0$, which holds with equality if $x_j > 0$. Thus, for each x_j we get $\sum_{i\in N(o_j)} f'(u_i(\boldsymbol{x}))x_j = \lambda x_j$. Summing up over all $j \in [m]$, we get:

$$\lambda = \lambda\sum_{j\in[m]} x_j = \sum_{i\in N}\sum_{o_j\in A_i} f'(u_i(\boldsymbol{x}))x_j = \sum_{i\in N} f'(u_i(\boldsymbol{x}))u_i(\boldsymbol{x}) \leqslant nM.$$

Thus for each $o_j \in O$, we have $nM \geqslant \sum_{i\in N(o_j)} f'(u_i(\boldsymbol{x}))$. Since there exists a project approved by all the members of S, it must be the case that $\sum_{i\in S} f'(u_i(\boldsymbol{x})) \leqslant nM$. Using Jensen's inequality, we find that

$$\sum_{i\in S} f'(u_i(\boldsymbol{x})) \geqslant |S| f'\Big(\frac{\sum_{i\in S} u_i(\boldsymbol{x})}{|S|}\Big) > |S| f'\Big((f')^{-1}\Big(\frac{nM}{|S|}\Big)\Big) = nM.$$

This gives a contradiction and completes the proof. □

Theorem 9 gives us examples of rules that provide well-defined group fairness guarantees, weaker than AFS. In particular, these guarantees do not in general imply IFS, so that the bound on utilitarian welfare given by Theorem 1 does not apply. Hence, we might hope that the weakening of AFS allows for asymptotically higher utilitarian welfare. The next theorem gives a lower bound on the utilitarian efficiency of f-UTIL for well-behaved f.

Theorem 10. *Let f be a continuous non-decreasing function with a convex derivative. Suppose $\kappa^* \in [0,1]$ is such that*

$$2\kappa^* \cdot f'(\tfrac{1}{2m}) = f'(\kappa^*). \tag{1}$$

Then the efficiency guarantee of f-UTIL is at least κ^*.

If we plug in $f = \log$, then $\kappa^* = \frac{1}{2\sqrt{m}}$ solves (1), which matches the guarantee of Theorem 4 up to a factor of 4. If we plug in $f(x) = \sqrt{x}$, we find that κ^* must satisfy $\kappa^*\sqrt{2m} = 1/(2\sqrt{\kappa^*})$, which gives $\kappa^* = \Omega(m^{-1/3})$. Thus, the sqrt-UTIL rule provides a substantially stronger utilitarian guarantee than the NASH rule. More generally, for $f(x) = x^\alpha$ with $0 < \alpha < 1$, we obtain a utilitarian efficiency of $\Omega(m^{1/(2-\alpha)-1})$.

Proof of Theorem 10. Let o_w be the most popular project, and let $\boldsymbol{x}$ be a lottery returned by f-UTIL.

If $\sum_{i\in N(o_w)} u_i(\boldsymbol{x}) \geqslant |N(o_w)|\kappa^*$, then we already obtain our guarantee. Thus, from now on we will assume that

$$\sum_{i\in N(o_w)} u_i(\boldsymbol{x}) < |N(o_w)|\kappa^*.$$

Consider a project o_j such that $|N(o_j)| < 2|N(o_w)|\kappa^*$. We will first show that $x_j \leqslant \frac{1}{2m}$. For the sake of contradiction, let us assume that $x_j > \frac{1}{2m}$. Then, we have that:

$$\sum_{i \in N(o_w)} f'(u_i(\boldsymbol{x})) \leqslant \sum_{i \in N(o_j)} f'(u_i(\boldsymbol{x})) < \sum_{i \in N(o_j)} f'\left(\tfrac{1}{2m}\right) = f'\left(\tfrac{1}{2m}\right) \cdot |N(o_j)|.$$

We continue the analysis using the Jensen's inequality:

$$\begin{aligned} f'\left(\tfrac{1}{2m}\right)|N(o_j)| &> \sum_{i \in N(o_w)} f'(u_i(\boldsymbol{x})) \\ &\geqslant |N(o_w)| f'\left(\frac{\sum_{i \in N(o_w)} u_i(\boldsymbol{x})}{|N(o_w)|}\right) \\ &\geqslant |N(o_w)| f'(\kappa^*). \end{aligned}$$

This would however imply that:

$$\frac{|N(o_j)|}{|N(o_w)|} > \frac{f'(\kappa^*)}{f'(\frac{1}{2m})} = 2\kappa^*,$$

a contradiction. Consequently, the projects o_j with $|N(o_j)| \leqslant 2|N(o_w)|\kappa^*$ get a total probability mass of at most $\frac{1}{2}$. As a result, a mass of at least $\frac{1}{2}$ is distributed among projects which are approved by at least $2|N(o_w)|\kappa^*$ voters. Thus, the total utility of the voters is at least equal to $\frac{1}{2} \cdot 2|N(o_w)|\kappa^*$, which yields the guarantee from the theorem statement. □

6 Average Guarantees: Experiments

In this section we extend our approach beyond the worst-case analysis. In a series of computer simulations we assess the *average* efficiency and egalitarian fairness of randomized rules assuming that voters' preferences come from certain distributions. We checked the following three distributions:

Euclidean Model. The voters and projects are associated with points in two-dimensional Euclidean space. The points are drawn uniformly at random from a square. Then, for each voter $i \in N$, we select $k_i \in \{1, \dots, m\}$ uniformly at random. The voter approves the k_i closest projects. (We also considered the same models where k_i is the same for all voters, or where voters approve all projects within a random radius. These models show a similar behavior.)

Impartial Culture. For each $i \in N$, choose $k_i \in \{1, \dots, m\}$ uniformly at random, and independently for each project o, voter i approves o with probability $\frac{k}{m}$. If, in the end, no project is approved, the voter approves one random project.

Mallow's (1957) Model. We first sample three reference rankings of the projects. Then, for each voter $i \in N$ we randomly choose one of these rankings, say π, and sample a ranking π_i from the Mallow's model with parameter π. Finally, we assume i approves the first k projects from π_i, where k is drawn uniformly at random from $\{1, \dots, m\}$.

For each distribution we consider 224 configurations for the values of n and m: we take $m = 2, 3, 4, 5, 10, 30, 50, 100$, and $n = 10, 20, \dots, 90, 100, 150, 200, \dots, 950, 1000$. For each configuration we draw 500 instances, and for each instance I and each rule f we calculate the normalized welfare $\widehat{\mathrm{sw}}(I, f(I))$.

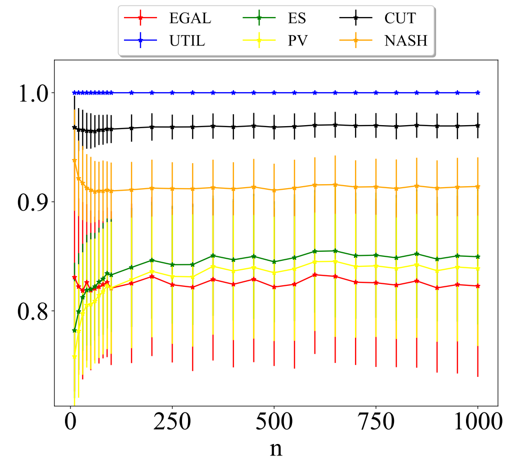

(a) $m = 10$

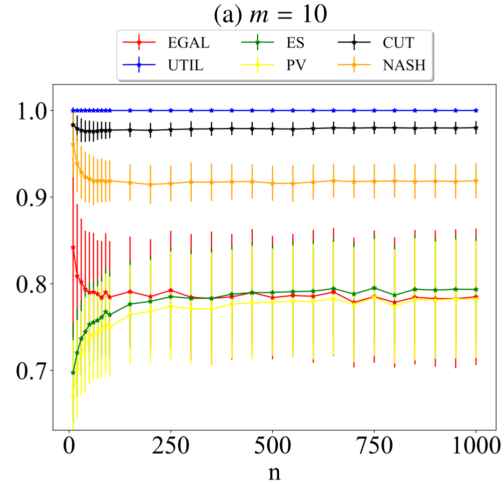

(b) $m = 50$

Figure 2: The average normalized social welfare of our voting rules, with preferences drawn from the Mallow's model.

We present our results in Figure 2 and in Table 2. The figure shows plots for the Mallow's Model for $m = 10, 50$. The table shows numerical values for different distributions, for the average and the worst case over the sampled instances.

The most striking feature of the plots in Figure 2 is the very high utilitarian welfare achieved by CUT, both in absolute terms and compared to the other rules. As shown in Table 2, CUT achieves 96%+ on average for all three probabilistic models. This suggests that, in practice, the utilitarian loss of CUT may be tiny. While NASH and PV have the same asymptotic worst-case utilitarian efficiency as CUT, we see that in our experiments, PV does much worse than NASH, and NASH does worse than CUT. However, the NASH rule still has high welfare, achieving 86%+ on average, suggesting that the extra utilitarian loss from the stronger fairness guarantees of NASH is moderate. While in the worst case, EGAL and ES have a worse utilitarian efficiency than PV, these three rules appear to have a similar performance on average.

We found that the differences between probabilistic models are small in terms of the price of fairness. Thus, we expect that our conclusions are fairly robust.

Rule		Euclidean Model				Impartial Culture				Mallow's Model			
		avg		worst		avg		worst		avg		worst	
	$m =$	10	50	10	50	10	50	10	50	10	50	10	50
CUT		0.92	0.93	0.73	0.7	0.99	0.99	0.96	0.98	0.97	0.98	0.94	0.95
NASH		0.86	0.85	0.67	0.66	0.97	0.96	0.91	0.93	0.91	0.92	0.86	0.86
PV		0.76	0.74	0.59	0.55	0.97	0.95	0.91	0.91	0.84	0.78	0.67	0.61
ES		0.76	0.73	0.59	0.55	0.97	0.95	0.91	0.91	0.85	0.79	0.71	0.64
EGAL		0.72	0.72	0.49	0.57	0.97	0.95	0.9	0.92	0.82	0.78	0.57	0.54

Table 2: The average efficiency guarantees for selected rules and distributions of voters' preferences, with $n = 1000$. The column "worst" gives the lowest value that was obtained over 500 samples.

7 Conclusion

We have studied the effect of imposing fairness constraints on the utilitarian social welfare in a simple model of budget division problems. Our surprisingly universal answer is that sensible rules in this context are guaranteed to provide a roughly $\frac{2}{\sqrt{m}}$ fraction of the optimum social welfare. This bound is obtained for the rule which maximizes Nash welfare, and for the simple conditional utilitarian rule. In each case, we can show that the guarantee is asymptotically tight.

While there is no asymptotic worst-case separation between these rules on the criterion we have studied, differences do emerge in an empirical analysis. Using a variety of probabilistic models, we find that the conditional utilitarian rule significantly outperforms all other fair rules proposed in the literature. In situations where high utilitarian welfare is particularly desired, this can provide a compelling reason to use the conditional utilitarian rule over, say, the Nash rule.

Acknowledgements. We thank Ayumi Igarashi for useful discussions. Marcin Michorzewski and Piotr Skowron were supported by the Foundation for Polish Science within the Homing programme (Project title: "Normative Comparison of Multiwinner Election Rules).

References

Airiau, S.; Aziz, H.; Caragiannis, I.; Kruger, J.; Lang, J.; and Peters, D. 2019. Portioning using ordinal preferences: Fairness and efficiency. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 11–17.

Aziz, H.; Bogomolnaia, A.; and Moulin, H. 2019. Fair mixing: the case of dichotomous preferences. In *Proceedings of the 20th ACM Conference on Economics and Computation (ACM-EC)*, 753–781.

Aziz, H.; Lee, B. E.; and Talmon, N. 2018. Proportionally representative participatory budgeting: Axioms and algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 23–31.

Benade, G.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Preference elicitation for participatory budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 376–382.

Bertsimas, D.; Farias, V. F.; and Trichakis, N. 2011. The price of fairness. *Operations Research* 59(1):17–31.

Bogomolnaia, A.; Moulin, H.; and Stong, R. 2002. Collective choice under dichotomous preferences. Extended working paper version.

Bogomolnaia, A.; Moulin, H.; and Stong, R. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory* 122(2):165–184.

Brandl, F.; Brandt, F.; Peters, D.; Stricker, C.; and Suksompong, W. 2019. Donor coordination: Collective distribution of individual contributions.

Brandt, F. 2017. Rolling the dice: Recent results in probabilistic social choice. In Endriss, U., ed., *Trends in Computational Social Choice*. AI Access. chapter 1, 3–26.

Cabannes, Y. 2004. Participatory budgeting: a significant contribution to participatory democracy. *Environment and Urbanization* 16(1):27–46.

Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; and Kyropoulou, M. 2012. The efficiency of fair division. *Theory Computing Systems* 50(4):589–610.

Duddy, C. 2015. Fair sharing under dichotomous preferences. *Mathematical Social Sciences* 73:1–5.

Fain, B.; Goel, A.; and Munagala, K. 2016. The core of the participatory budgeting problem. In *Proceedings of the 12th International Workshop on Internet and Network Economics (WINE)*, 384–399.

Fain, B.; Munagala, K.; and Shah, N. 2018. Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (ACM-EC)*, 575–592.

Goel, A.; Krishnaswamy, A. K.; Sakshuwong, S.; and Aitamurto, T. 2019. Knapsack voting for participatory budgeting. *ACM Transactions on Economics and Computation* 7(2):8:1–8:27.

Guerdjikova, A., and Nehring, K. 2014. Weighing experts, weighing sources: The diversity value.

Lackner, M., and Skowron, P. 2019. A quantitative analysis of multi-winner rules. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 407–413.

Mallows, C. L. 1957. Non-null ranking models. *Biometrika* 44(1/2):114–130.

Sánchez-Fernández, L.; Elkind, E.; Lackner, M.; Fernández, N.; Fisteus, J. A.; Basanta Val, P.; and Skowron, P. 2017. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*.

Talmon, N., and Faliszewski, P. 2019. A framework for approval-based budgeting methods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, volume 33, 2181–2188.