

Portioning Using Ordinal Preferences: Fairness and Efficiency

Stéphane Airiau¹, Haris Aziz², Ioannis Caragiannis³,
Justin Kruger¹, Jérôme Lang¹ and Dominik Peters⁴

¹LAMSADE, CNRS, Université Paris-Dauphine, PSL University

²UNSW Sydney and Data61 CSIRO

³University of Patras

⁴University of Oxford

Abstract

A public divisible resource is to be divided among projects. We study rules that decide on a distribution of the budget when voters have ordinal preference rankings over projects. Examples of such portioning problems are participatory budgeting, time shares, and parliament elections. We introduce a family of rules for portioning, inspired by positional scoring rules. Rules in this family are given by a scoring vector (such as plurality or Borda) associating a positive value with each rank in a vote, and an aggregation function such as leximin or the Nash product. Our family contains well-studied rules, but most are new. We discuss computational and normative properties of our rules. We focus on fairness, and introduce the SD-core, a group fairness notion. Our Nash rules are in the SD-core, and the leximin rules satisfy individual fairness properties. Both are Pareto-efficient.

1 Introduction

The members of an organization need to divide its budget among several projects. They have different opinions about the relative value of different projects, and would like to *vote* over the budget. What kind of voting rule could they use?

Some cities let citizens vote over the use of the city budget, giving rise to *participatory budgeting*. In deployed applications (such as in Paris or Brazil), the projects are *indivisible*, and can be either fully funded or not at all, such as refurbishing a school or adding a bike lane. We focus on *divisible* projects on which an arbitrary fraction of the budget could be spent, such as ‘education’ or ‘transport’ or ‘parks’. The result of the vote can be visualized as a pie chart showing which percentage of the budget is spent on each type of expense. The ‘budget’ need not be monetary, and we refer to this general scheme as *portioning*. There are many applications:


- A conference board deciding how much time to assign to talks, poster sessions, invited talks, and coffee breaks.
- A parliamentary election deciding what percentage of parliament seats should go to each party.
- Coauthors deciding how much space to devote to various topics in a textbook or article with fixed total length.
- A company annually donates money to charity, and lets employees vote over which charities should receive a donation.

We ask voters to report their preferences over projects as *rankings*, the most common format considered in social choice. If a project is ranked more highly, the voter thinks it is more worthwhile and should receive a larger fraction of the budget.


The space of sensible aggregation rules is large, so let us illustrate some important design considerations by an example.

An Example


A family is planning a road trip by car. The family members have different musical tastes; they need to decide which type of music to play for how long. The genres under consideration are a, b, c, d, e . The three children all think $a \succ b \succ c \succ d \succ e$; mother thinks $e \succ b \succ c \succ d \succ a$; and father thinks $c \succ a \succ e \succ d \succ b$.



One simple way to split the time is to allocate each person the same share of time (20%) and let them decide what music to play, as a temporary dictator. During their time, each person plays their favorite music. To the social choice theorist, this rule sounds familiar: it is formally identical to *Random Dictatorship*, whose output is usually seen not as a division of a budget, but as a probability distribution. Indeed, any probabilistic social choice function can be repurposed to divide budgets; but these are often not attractive for portioning since many of them were designed as tie-breaking devices.




The output of random dictatorship can be a good choice, especially if our family strongly prefers their top choice to any other music. But it is also plausible that mother and the children agree that b is good common ground. Random Dictatorship, using plurality scores, ignores this. Instead, we could impute *Borda scores* on our family: for example, the children give 4 ‘utility’ points to a , 3 to b , 2 to c , 1 to d , and 0 to e . *Proportional Borda* then allocates time in proportion to the total Borda score of the genres. This leads to a significant time share for b . On the other hand, the family now also listens to d , which is dominated: everyone agrees that c is better than d ! So Proportional Borda is inefficient.




To restore efficiency, it makes sense to maximize a notion of social welfare. Suppose the utility enjoyed by a family member is the weighted average of the Borda scores of the music played on the trip, where the weights come from the fraction of time spent on each genre. *Utilitarian Borda* then

picks the distribution where the sum of utilities is greatest. In our example, we listen to a during 100% of the time. While this is Pareto-efficient, it is unfair to mother, who only gets to listen to her least-preferred style. In fact, many rules suffer from this phenomenon of completely overriding some voters’ preferences: For example, the ‘maximal lotteries’ rule also only plays a since it is the Condorcet winner.



To avoid frustration during the trip, we may take a more egalitarian approach, and aim to give every family member a significant share. *Borda-Egalitarian* picks the distribution maximizing the utility of the worst-off passenger. In our example, we can give every passenger an average Borda-utility of 2.4. To ensure Pareto-efficiency, we can refine this to *Borda-Leximin* (which does not change the outcome on this example).



We can also maximize *Nash social welfare*, the product of utilities. This is often seen as a compromise between maximizing utilitarian and egalitarian welfare notions. While egalitarian rules perform well when we wish to be fair to each *individual*, Nash rules tend to be fair to *groups*. In our example, the children form a large group, and Borda-Nash plays a almost half the time. If there were more children with the same preferences in the car, Borda-Nash would increase the time share of a . In contrast, Borda-Leximin avoids playing a to benefit the mother, and the output of egalitarian rules does not change with the number of children. Depending on the context, either of these behaviors might be more appropriate.

Our Contributions

We introduce a class of aggregation rules called *positional social decision schemes*. Rules in this class first convert each input ranking into scores for the alternatives, using a scheme such as plurality or Borda scores. Then, they select a distribution of the budget maximizing social welfare given those scores, where different notions of welfare can be used; classically, we consider utilitarian, egalitarian (leximin), and Nash welfare. Our class contains known rules such as random dictatorship, but most have not been studied.

We begin by noting basic properties of the rules in our class, giving closed forms and equivalent definitions in some cases. We also show that the rules in this class can be calculated or approximated in polynomial time. For rules based on Nash welfare, we show that their output can involve irrational percentages; we prove that those rules are guaranteed to be rational if the scoring vector used is plurality or veto, but that no other scoring vector guarantees rational output.

We then formalize intuitive notions of fairness in the budgeting context. The axioms we propose require that no individual is ignored by the procedure, in the sense of having none of the budget allocated to favored causes. We also give some group fairness notions. Our strongest axiom is the *SD-core* which, roughly, requires that a group of $\alpha\%$ of the voters can control what happens with $\alpha\%$ of the budget. We show that the rules in our class based on Nash welfare satisfy the SD-core, while the egalitarian rules satisfy the individual fairness notions.

We close by studying the performance of our rules on standard social choice properties, such as Pareto-efficiency, strategyproofness, and monotonicity.

Related Work

Bogomolnaia *et al.* [2005] introduced the portioning problem, motivated by time-sharing. They assume *dichotomous preferences*, and agents report a subset of the alternatives (an approval vote), rather than rankings. They study the compatibility of Pareto-efficiency and strategyproofness, with positive results (for example, spending the entire budget on the approval winner satisfies both requirements). However, after adding a fairness axiom, they get an impossibility result. A related impossibility is proved by Duddy [2015]. Aziz *et al.* [2019] introduce some new rules based on welfare maximization, and introduce new fairness axioms (including a core notion), and a weakened strategyproofness axiom. Brandl *et al.* [2019] study cases when the budget is owned by the voters.

Fain *et al.* [2016] study portioning in a *cardinal* model which allows agents to give a full utility function over alternatives (which may also feature decreasing returns). They study the core and connect it to the Lindahl equilibrium from the study of public goods, and prove that a core solution always exists. For a broad class of utility functions, they show that a core solution can be found in polynomial time by solving a suitable convex program. They also use differential privacy to design a mechanism for this setting which satisfies approximate versions of efficiency, truthfulness, and the core.

With rankings as input, this setting has been studied in the formally isomorphic guise of *probabilistic social choice* [see Brandt, 2018 for a recent survey]. In this literature, the outcome distribution is interpreted as a random device, which is used to eventually implement a single outcome. This makes notions of fairness and proportionality less relevant, and it is seen as desirable for a rule to randomize as little as possible. For example, the *maximal lotteries* rule [Kreweras, 1965; Brandl *et al.*, 2016], while attractive according to consistency axioms, spends the entire budget on the Condorcet winner if it exists. This is often undesirable in a budgeting context. On the other hand, results like Gibbard’s [1977] random dictatorship theorem are important in the portioning context. Some papers on probabilistic social choice also discuss fairness concerns [see, e.g., Aziz *et al.*, 2018b; Aziz and Stursberg, 2014].

In a related setting, projects are indivisible and come with a fixed cost; they can either be fully funded or not at all. This is known as *combinatorial public projects* [Papadimitriou *et al.*, 2008] in a literature focussing on computational aspects and truthfulness. This setting is also what is often meant by “participatory budgeting” [Cabannes, 2004; Goel *et al.*, 2016; Benade *et al.*, 2017]. Several recent papers have studied fairness in this setting [Aziz *et al.*, 2018a; Fain *et al.*, 2018; Conitzer *et al.*, 2017], inspired by *multi-winner elections*, for which fairness and proportionality are well-studied [Aziz *et al.*, 2017; Faliszewski *et al.*, 2017].

The literature on cake-cutting and item allocation is mostly *unrelated* to our work: in those settings, goods are allocated to specific agents for their exclusive use. In our setting, resources are spent on projects which can be enjoyed by all agents. On a technical level, the idea of scoring followed by aggregation has been explored in fair division [Brams and King, 2005; Darmann and Schauer, 2015; Baumeister *et al.*, 2016], and work on group fairness raises related issues [see, e.g., Manurangsi and Suksompong, 2019; Bade and Segal-Halevi, 2018].

2 Positional Social Decision Schemes

Let $X = \{x_1, \dots, x_m\}$ be a set of alternatives and $N = \{1, \dots, n\}$ be a set of voters. Let $\mathcal{L}(X)$ be the set of linear orders over X . For $\succ \in \mathcal{L}(X)$, the *rank* of alternative x_j is $r(\succ, x_j) = 1 + |\{x_i \in X : x_i \succ x_j\}|$. A *profile* $P = (\succ_1, \dots, \succ_n) \in \mathcal{L}(X)^n$ is a collection of linear orders, one for each voter. We write abc as shorthand for $a \succ b \succ c$. Let $\Delta(X) = \{p : X \rightarrow [0, 1] : \sum_{x \in X} p_x = 1\}$ be the set of (probability) distributions over X . We use notation like $\frac{1}{2}x_1 + \frac{1}{2}x_2$ to specify a distribution, and write x_j for the distribution with $p_{x_j} = 1$. We say that $z : X \rightarrow [0, 1]$ is a *partial distribution* if $\sum_{x \in X} z_x \leq 1$. A *social decision scheme* (SDS) is a function F assigning to each $P \in \mathcal{L}(X)^n$ a nonempty subset of $\Delta(X)$, usually a singleton.

A *scoring vector* for m alternatives is a vector $s = (s_1, \dots, s_m)$ of numbers with $s_1 \geq s_2 \geq \dots \geq s_m$ and $s_1 > s_m$. We usually assume $s_m = 0$. A scoring vector s is *strictly decreasing* if $s_j > s_{j+1}$ for all $j < m$. The Borda vector is **bor** = $(m-1, m-2, \dots, 0)$; the plurality vector is **plu** = $(1, 0, \dots, 0)$; the veto vector is **vet** = $(1, \dots, 1, 0)$.

For a fixed profile P , we write $s[i, j] = s_{r(\succ_i, x_j)}$ for the s -score that voter $i \in N$ assigns to alternative $x_j \in X$. These scores can be lifted to distributions in a natural way; the s -score of $p \in \Delta(X)$ for i is $s[i, p] = \sum_{j=1}^m p_j s[i, j]$. Finally, define the utility vector $s[p] = (s[1, p], \dots, s[n, p])$.

A *welfare ordering* is a weak order \geq_W ordering utility vectors $(\alpha_1, \dots, \alpha_n) \in \mathbb{R}_{\geq 0}^n$. The main examples are utilitarianism which orders vectors by their sum, egalitarianism which uses the minimum, the Nash product which uses multiplication, and leximin which sorts the components of the utility vector and then orders sorted vectors lexicographically.

By combining a scoring vector and a welfare ordering, we can define a positional social decision scheme.

Definition 1. For scoring vector s and a welfare ordering \geq_W , define the social decision scheme F_{s, \geq_W} so that for all P ,

$$F_{s, \geq_W}(P) = \{p \in \Delta(X) : s[p] \geq_W s[q] \text{ for all } q \in \Delta(X)\}.$$

For the specific \geq_W mentioned, we usually call these rules s -utilitarianism, s -egalitarianism, s -leximin, and s -Nash.

Example 1. Consider the profile $P = (ab, ab, ba)$ over two alternatives, with $s = (1, 0)$. Then s -utilitarianism selects a , s -egalitarianism selects $\frac{1}{2}a + \frac{1}{2}b$, and s -Nash selects $\frac{2}{3}a + \frac{1}{3}b$.

For normative analysis, it is useful to extend voters' rankings of the alternatives to (partial) preferences over distributions. We assume *linear preferences*: there is an unknown utility function $u_i : X \rightarrow \mathbb{R}$ consistent with \succ_i such that i prefers those distributions p with higher average utility $\sum_{x \in X} u_i(x)p_x$. A classical way of ranking distributions despite not knowing u_i uses *stochastic dominance* (SD).

If p and q are (possibly partial) distributions, we write

$$p \succ_i^{\text{SD}} q \iff \sum_{x_k \succ_i x_j} p_{x_k} \geq \sum_{x_k \succ_i x_j} q_{x_k} \text{ for all } x_j \in X.$$

This definition is justified by the following standard equivalence: We have $p \succ_i^{\text{SD}} q$ if and only if $\sum_{x \in X} u_i(x)p_x \geq \sum_{x \in X} u_i(x)q_x$ for all utility functions $u_i : X \rightarrow \mathbb{R}$ satisfying $\min_{x \in X} u_i(x) = 0$ and $u_i(x_k) > u_i(x_j)$ iff $x_k \succ_i x_j$. The

condition that the utility of the worst alternative is 0 is necessary to allow SD-comparisons of partial distributions: we assume that voters are indifferent between not spending part of the budget or spending it on their worst alternative. This is crucial for the definition of the SD-core in Section 4.

3 Computation and Basic Properties

In this section, we look at elementary properties of the family of rules we have defined. We will note that several of the rules are familiar from the probabilistic context. We also study the computational complexity of finding an optimal distribution.

3.1 Utilitarianism

From a utilitarian perspective, it never pays to spend part of the budget on alternatives whose total s -score is not maximal: shifting that spending to an s -maximal alternative increases utilitarian welfare. Thus, up to ties, s -utilitarianism never mixes and spends all resources on the s -winner. Formally, s -utilitarianism selects those distributions p for which $p_{x_j} > 0$ only if the score $\sum_{i \in N} s[i, j]$ is maximum.

Since the behavior of s -utilitarianism is familiar from work on scoring rules in voting, we will not study it in much detail.

3.2 Egalitarianism

Plurality-egalitarianism is easy to understand: it returns the uniform distribution over all alternatives that are ranked top by at least one voter. In the probabilistic context, this rule is known as *egalitarian simultaneous reservation* [Azziz and Stursberg, 2014]. For other scoring vectors, s -egalitarianism is less simple, and it need not return a uniform distribution (see the example of Section 1). However, one can easily evaluate s -egalitarianism using linear programming:

$$\begin{aligned} \text{maximize } t^* \text{ s.t. } & \sum_{j=1}^m s[i, j] \cdot p_j \geq t^* \text{ for } i \in N \\ & \sum_{j=1}^m p_j = 1, \text{ and } p_j \geq 0 \text{ for } x_j \in X \end{aligned}$$

Now, s -egalitarianism is not very decisive, and may select Pareto-inferior outcomes. When $P = (abcd, acbd, bdac)$, and $s = (1, 1, 0, 0)$, it selects all distributions of the form

$$p \cdot a + q \cdot b + (\frac{1}{2} - p) \cdot c + (\frac{1}{2} - q) \cdot d$$

where $0 \leq p, q \leq 1$ and $\frac{1}{2} \leq p + q \leq 1$. Note that d can get a positive fraction even though every voter prefers b to d (so that d is Pareto-dominated). A standard way of making egalitarianism more decisive and more efficient is by using leximin instead. In the above example, s -leximin uniquely selects $\frac{1}{2}a + \frac{1}{2}b$. It is easy to see that s -leximin will never give a positive fraction to a Pareto-dominated alternative.

It is still possible to evaluate s -leximin in polynomial time, by solving $O(n^2)$ linear programs successively. Our algorithm uses the convexity of $\Delta(X)$, which allows it to greedily fix the identity of the agent who is worst-off in the current iteration.

Theorem 1. For every s , one can compute a distribution selected by s -leximin in polynomial time.

Proof. The algorithm is specified as Algorithm 1 that requires running at most $n(n+1)/2$ linear programs.

Relabel the voters in N as i_1, \dots, i_n in the order that they were added to N' by the algorithm. We argue by induction on k that every distribution p selected by s -leximin satisfies

Algorithm 1 Computing an s -leximin distribution

Set $N' \leftarrow \emptyset$. For $i \in N$, we will set t_i once i is added to N' .

while $N' \neq N$ **do**

Using linear programming, find the maximum value t^* such that there exists a distribution (p_1, \dots, p_m) satisfying

$$\begin{aligned} \sum_{j=1}^m s[i, j] p_j &\geq t^* && \text{for } i \in N \setminus N' \\ \sum_{j=1}^m s[i, j] p_j &= t_i && \text{for } i \in N' \end{aligned}$$

for each $i' \in N \setminus N'$ **do**

Using linear programming, find the maximum ε such that there exists a distribution (p_1, \dots, p_m) satisfying

$$\begin{aligned} \sum_{j=1}^m s[i', j] p_j &\geq t^* + \varepsilon \\ \sum_{j=1}^m s[i, j] p_j &\geq t^* && \text{for } i \in N \setminus N' \\ \sum_{j=1}^m s[i, j] p_j &= t_i && \text{for } i \in N' \end{aligned}$$

If $\varepsilon = 0$, add i' to N' and set $t_{i'} \leftarrow t^*$.

return the solution (p_1, \dots, p_m) of the last LP solved

$\sum_{j=1}^m s[i_k, j] p_j = t_{i_k}$. Suppose $k \geq 1$, and we have shown this holds for i_1, \dots, i_{k-1} . Let p be a distribution selected by s -leximin. In the iteration before i_k is added to N' , because p satisfies the inductive hypothesis, from the upper LP, we know that the least s -score obtained by a voter in $N \setminus N'$ under p is t^* . But which voter? The lower LP tests, for each $i' \in N \setminus N'$, whether i' obtains s -score exactly t^* in all leximin distributions. Such a voter must exist: suppose not, and for each $i' \in N \setminus N'$, let $p_{(i')}$ be a leximin distribution where i' obtains s -score strictly higher than t^* . Write $p' = \sum_{i' \in N \setminus N'} \frac{1}{|N \setminus N'|} p_{(i')}$. Then p' satisfies the inductive hypothesis (because each $p_{(i')}$ does), but the least s -score obtained by a voter in $N \setminus N'$ under p' is strictly higher than t^* , contradicting the choice of t^* . Thus, there is a voter $i^* \in N \setminus N'$ who obtains s -score t^* in all leximin distributions. This voter is named i_k and t_{i_k} is set to t^* , establishing the inductive step. \square

3.3 Nash product

The defining optimization problem

$$\begin{aligned} &\text{maximize } \sum_{i \in N} \log \left(\sum_{j=1}^m s[i, j] \cdot p_j \right) \\ &\text{s.t. } \sum_{j=1}^m p_j = 1, \text{ and } p_j \geq 0 \text{ for } x_j \in X \end{aligned}$$

of s -Nash is a convex program which can be efficiently solved using standard solvers. Formally, one can approximate optimum Nash welfare within an additive factor of ε in time polynomial in n , m , and $1/\varepsilon$. Thus, all the usual decision problems associated with computing s -Nash are easy. However, writing down the precise output in decimal expansion is impossible, as there are instances where s -Nash uniquely returns a distribution with irrational fractions. For instance, for $P = (abc, acb, cab, cab)$ and $s = (2, 1, 0)$, s -Nash uniquely returns $\frac{1+\sqrt{33}}{8}a + \frac{7-\sqrt{33}}{8}c$.

To further understand s -Nash, let us analyze the first-order conditions of the convex program. Write down the Lagrangian

$$\mathcal{L} = \sum_{i \in N} \log \left(\sum_{j=1}^m s[i, j] \cdot p_j \right) - \lambda \cdot (1 - \sum_{j=1}^m p_j).$$

At an optimal solution p , we have

$$\frac{\partial \mathcal{L}}{\partial p_j} = \sum_{i \in N} \frac{s[i, j]}{s[i, p]} - \lambda \leq 0, \text{ with equality if } p_j > 0.$$

This implies $\lambda p_j = \sum_{i \in N} \frac{s[i, j]}{s[i, p]} p_j$. Summing over all j , thus

$$\lambda = \lambda(p_1 + \dots + p_m) = \sum_{j=1}^m \sum_{i \in N} \frac{s[i, j]}{s[i, p]} p_j = n,$$

since $s[i, p] = \sum_{j=1}^m s[i, j] \cdot p_j$ by definition. It follows that

$$n \geq \sum_{i \in N} \frac{s[i, j]}{s[i, p]}, \text{ with equality if } p_j > 0. \quad (1)$$

For example, using (1), we can characterize plurality-Nash [see also Moulin, 2003, Example 3.6]:

Theorem 2. *Plurality-Nash selects p with $p_j = \text{pl}(x_j)/n$ for all j , where $\text{pl}(x_j)$ is the number of voters placing x_j top.*

Proof. Let p be optimal for plurality-Nash. If some voter i puts x_j top then $p_j > 0$, or else $s[i, p] = 0$ and the Nash product equals 0. By (1), we get $n = \sum_{i \in N} \frac{s[i, j]}{s[i, p]} = \text{pl}(x_j)/p_j$, and so $p_j = \text{pl}(x_j)/n$. It follows that $p_j = 0$ whenever no voter places x_j top. \square

Thus, we see that plurality-Nash is the same rule as *random dictatorship*, familiar from the probabilistic context.

The *veto-Nash* rule seems sensible when alternatives are nuisances, where each agent wants to minimize the amount spent on the worst option. In some sense, veto-Nash for nuisances is as relevant as plurality-Nash for goods, in the portioning context. Mathematically, veto-Nash is also well-behaved. While we do not provide a closed formula, the following result shows that an exact optimum for veto-Nash can be found in polynomial time (and that it is rational). It gives a collection of at most m different explicit rational distributions, and guarantees that the veto-Nash optimum is among them.

Theorem 3. *Let P be a profile, and let $\text{vt}(x_j)$ be the number of voters placing x_j bottom. Relabel alternatives so that $\text{vt}(x_1) \leq \dots \leq \text{vt}(x_m)$. If $\text{vt}(x_j) = 0$ for some x_j , veto-Nash selects all distributions over such alternatives. Otherwise, there is $k \in [m]$ with $(k-1)\text{vt}(x_k) < \sum_{j=1}^k \text{vt}(x_j)$, such that veto-Nash selects the distribution p with*

$$p_j = 1 - \frac{(k-1)\text{vt}(x_j)}{\sum_{i=1}^k \text{vt}(x_i)} \text{ if } j \in [k], \text{ and } p_j = 0 \text{ otherwise.}$$

Proof. If $\text{vt}(x_j) = 0$ for some x_j , then the best-possible Nash product of 1 can be achieved, and is achieved precisely by distributions whose support consists of never-vetoeed alternatives.

Now suppose that $\text{vt}(x_j) > 0$ for all x_j . Let p be a distribution selected by veto-Nash, and take k maximal such that $p_k > 0$. Then we must also have $p_j > 0$ for all $j = 1, \dots, k-1$. (If not, and $p_j = 0$ for some j , consider the distribution q with $q_l = p_l$ for all l , except that $q_j = q_k = \frac{1}{2}p_k$. Then, since $\text{vt}(x_j) \leq \text{vt}(x_k)$, q has strictly higher Nash product than p , contradiction.) Thus, for $i = 1, \dots, k$, equation (1) applies and can be written as

$$n = \sum_{j \in [k] \setminus \{i\}} \frac{\text{vt}(x_j)}{1-p_j} + \sum_{j=k+1}^m \text{vt}(x_j). \quad (2)$$

Summing the equations (2) for $i = 1, \dots, k$, we get

$$nk = (k-1) \sum_{j \in [k]} \frac{\text{vt}(x_j)}{1-p_j} + k \sum_{j=k+1}^m \text{vt}(x_j).$$

Using $n = \sum_{j=1}^k \text{vt}(x_j) + \sum_{j=k+1}^m \text{vt}(x_j)$, rearrange this as

$$\sum_{j \in [k]} \frac{\text{vt}(x_j)}{1-p_j} = \frac{k}{k-1} \sum_{j=1}^k \text{vt}(x_j).$$

From the symmetry of the equations (2), the values $\frac{\text{vt}(x_i)}{1-p_i}$ must be equal for all $i \in [k]$. Since we know their sum, we get

$$\frac{\text{vt}(x_i)}{1-p_i} = \frac{1}{k-1} \sum_{j=1}^k \text{vt}(x_j) \quad \text{for all } i \in [k].$$

Rearranging, we arrive at the conclusion that

$$p_i = 1 - \frac{k-1}{\sum_{j=1}^k \text{vt}(x_j)} \text{vt}(x_i) \quad \text{for all } i \in [k].$$

These values sum to 1, and are non-negative provided that $(k-1)\text{vt}(x_k) < \sum_{j=1}^k \text{vt}(x_j)$. If this condition is not satisfied, the choice of k cannot lead to a veto-Nash optimum. \square

This gives an algorithm for computing veto-Nash exactly: if some alternatives are never vetoed, return any distribution over these. Otherwise iterate over all $k \in [m]$ satisfying the condition of the theorem and calculate the corresponding distribution, and return the one with highest Nash product.

Example 2. If 2, 3, 3 and 5 voters rank x_1, x_2, x_3 and x_4 last, respectively, then $k = 2$ and 3 satisfy the condition of Thm. 3. Thus, either $p = \frac{3}{5}x_1 + \frac{2}{5}x_2$ or $p' = \frac{1}{2}x_1 + \frac{1}{4}x_2 + \frac{1}{4}x_3$ is optimal. The former has higher Nash product, so p is optimal.

Theorems 2 and 3 show that both plurality-Nash and veto-Nash are rational. Are there any other score vectors s such that s -Nash is guaranteed to be rational? The answer is no: for every s other than plurality and veto, we can construct a profile where s -Nash uniquely returns an irrational distribution. This result suggests that a convex programming solver is the best way of computing s -Nash for s other than plurality and veto.

Theorem 4. Let $m \geq 3$, and let $s = (s_1, \dots, s_m) \in \mathbb{Q}^m$ be a score vector with $s_m = 0$ and normalized so that $s_1 = 1$. Unless $s = (1, 0, \dots, 0)$ or $s = (1, \dots, 1, 0)$, there exists a profile $P \in \mathcal{L}(X)^n$ for some $n \in \mathbb{N}$ such that s -Nash returns a unique distribution p with $p \notin \mathbb{Q}^m$.

Proof. We construct four infinite families of examples, for different shapes of score vectors s . Due to space constraints, we only consider the case $m = 3$ here, and only sketch the algebra required. The other families require a more involved construction, but work using similar calculus.

Suppose $m = 3$, and let $s = (1, \frac{r}{s}, 0)$, where $0 < \frac{r}{s} < 1$ and $\frac{r}{s}$ is in lowest terms. Let c be a large-enough integer. Consider the following profile: c voters with abc , one voter bac , one voter with bca . Note that b Pareto-dominates c , so that $p_c = 0$. Let $(x, 1-x, 0)$ be the distribution selected by s -Nash. One can show that $0 < x < 1$ if c is large enough. Now, the Nash product obtained by this distribution is $(x + \frac{r}{s}(1-x))^c \cdot ((1-x) + \frac{r}{s}x) \cdot (1-x)$. By optimality, x must make the derivative d/dx vanish. After a calculation, cancelling non-zero factors, this implies that

$$\begin{aligned} & ((c+2)(r-s)^2) \cdot x^2 \\ & + (-r-s)((c+3)r - 2(c+1)s) \cdot x \\ & + (r^2 - 2rs - crs + cs^2) = 0 \end{aligned}$$

This is a quadratic equation with integer coefficients. Solutions to the equation $ax^2 + bx + c = 0$ involve the term $\sqrt{b^2 - 4ac}$; thus, they are rational if and only if $b^2 - 4ac$ is a perfect square. In our case, the term under the square root simplifies to

$$(c+1)^2 r^2 + 4(rs + s^2).$$

The first summand is a large perfect square, and the second summand is a constant. Since the distance between consecutive perfect squares is large (in the sense that $(z+1)^2 - z^2 = 2z + 1 = \Theta(z)$), the discriminant cannot be a perfect square for large enough c . Hence, x is irrational. \square

4 Fairness, Proportionality, and the SD-core

Usually, s -utilitarianism spends 100% on a single alternative. Some agents might rank this alternative in a very low position, or even in last place. In some contexts, this is unfair and might rule out s -utilitarianism. In this section, we formalize several notions of fairness, and show that s -egalitarianism satisfies individual fairness, and that s -Nash satisfies group fairness.

A minimal fairness axiom is *positive share* [adapted from Bogomolnaia *et al.*, 2005] which requires that if voter i ranks x in last position, then $p_x < 1$. Hence, for every voter, a positive amount is spent on alternatives not ranked in last position. As suggested above, s -utilitarianism fails positive share for any s . However, provided that $s_m = 0$, positive share is satisfied by s -egalitarianism, s -leximin, and s -Nash. To see this, note that the uniform distribution has positive egalitarian and Nash welfare, whereas a distribution violating positive share has zero egalitarian and Nash welfare.

We can strengthen positive share to *individual fair share*, requiring that if voter i ranks x in last position, then $p_x \leq 1 - \frac{1}{n}$. Thus, for each voter, at least $\frac{1}{n}$ is spent on alternatives not ranked last. Note that the distribution identified by random dictatorship satisfies this condition and has egalitarian welfare at least $\frac{1}{n}$, normalizing $s_1 = 1$. Thus, the optimum s -egalitarian welfare is at least $\frac{1}{n}$, and hence s -egalitarianism and s -leximin satisfy individual fair share (recalling that $s_m = 0$). Below, we show that s -Nash also satisfies it.

Consider $X = \{a, b\}$, with 9 voters ab and 1 voter ba . Then s -egalitarianism returns $\frac{1}{2}a + \frac{1}{2}b$. While this is individually fair, the group of 9 voters is underrepresented. If we desire fairness to groups, we need a stronger axiom. One option is this: if k out of n voters rank x last, then $p_x \leq 1 - \frac{k}{n}$, so at least $\frac{k}{n}$ is spent on alternatives other than x . This condition is failed by s -egalitarianism and s -leximin, but s -Nash satisfies it. In our example, s -Nash picks $\frac{9}{10}a + \frac{1}{10}b$.

All the notions above focus on avoiding voters' last-ranked alternative. Despite working in an ordinal setting, using the SD-extension, we can define a group fairness notion that uses more than just the last-ranked alternative. An important underlying intuition is that agents are "entitled" to $1/n$ of the budget, and this share should be spent in accordance to their preferences. Similarly, a group $S \subseteq N$ of k agents could pool together and be entitled to k/n of the budget.

The intuitive notion of entitlement can be formalized using a core-style concept. A coalition $S \subseteq N$ of voters is supposed to be able to 'control' a fraction of $|S|/n$ of the entire budget. The notion of control is ambiguous since coalitions




Figure 1: The SD-core of the profile (abc, acb, bca) within the simplex of all distributions. The shaded area shows the distributions that are in the SD-core. The blue line shows the output of s -Nash for all $s = (1, q, 0)$ with $q \in [0, 1]$. Plurality-Nash selects $\frac{2}{3}a + \frac{1}{3}b$, Borda-Nash selects $0.58a + 0.42b$, and veto-Nash selects $\frac{1}{3}a + \frac{1}{3}b + \frac{1}{3}c$.

may overlap and each share of the budget is simultaneously controlled by several coalitions. However, the entitlement of S is certainly violated under p if S can come up with a way of using only its entitlement $|S|/n$ which all members prefer to the way that p uses the entire budget.

Definition 2. A coalition $S \subseteq N$ SD-blocks a distribution p if there exists a partial distribution z with $\sum_{x \in X} z_x = |S|/n$ such that $z \succ_i^{\text{SD}} p$ for all $i \in S$, and $z \succ_j^{\text{SD}} p$ for some $j \in S$. A distribution p is in the SD-core if no coalition SD-blocks p .

If a distribution p lies in the SD-core, then it also satisfies our other properties: Suppose not, and there is a coalition S of voters that rank x last, where $p_x > 1 - |S|/n$. Then S can SD-block p : Write $\varepsilon = p_x - (1 - |S|/n) > 0$ and define a deviation z with $z_y = p_y + \varepsilon/(m-1)$ for all $y \in A \setminus \{x\}$, and $z_x = 0$. Then $\sum_{a \in A} z_a = \varepsilon + \sum_{y \in A \setminus \{x\}} p_y = \varepsilon + (1 - p_x) = |S|/n$, so that z has the required total weight. It is easy to check that $z \succ_i^{\text{SD}} p$ for all $i \in S$. Thus, p is not in the SD-core.

For an example, take the profile with voters abc, acb, bca . Which distributions p are in the SD-core? First, singleton coalitions $\{i\}$ block p if $p_x > \frac{2}{3}$ for i 's bottom alternative x , using $z = \frac{1}{3}y$ where y is i 's top alternative. The coalition of abc and acb blocks all p with $p_a + p_b \leq \frac{2}{3}$ and $p_a + p_c \leq \frac{2}{3}$ (one inequality strict), using $z = \frac{2}{3}a$. All other distributions are in the SD-core. Figure 1 shows the simplex of all distributions, with the SD-core shaded (non-convex in this example).

Figure 1 shows the outputs of s -Nash for all s as a blue line. The blue line is entirely contained in the SD-core. In fact, s -Nash is always in the SD-core. We give a direct argument using equation (1). The result can also be obtained via the theory of Lindahl equilibrium [Fain *et al.*, 2016; Foley, 1970].

Theorem 5. For any s with $s_m = 0$, any distribution selected by s -Nash is in the SD-core.

Proof. Suppose p is selected by s -Nash. For a contradiction, assume that $S \subseteq N$ is a blocking coalition of agents, deviating using $(z_1, \dots, z_m) \in [0, 1]^m$ with $\sum_{j=1}^m z_j = |S|/n$, such that $z \succ_i^{\text{SD}} p$ for all $i \in S$, and $z \succ_j^{\text{SD}} p$ for some $j \in S$. Now, s defines utilities compatible with the voters' ordinal

preferences, and thus $s[i, z] \geq s[i, p]$ for all $i \in S$, and $s[j, z] > s[j, p]$ for some $j \in S$. Then

$$|S| = n \cdot \sum_{j=1}^m z_j \stackrel{(1)}{\geq} \sum_{i \in N} \frac{\sum_{j=1}^m s[i, j] z_j}{s[i, p]} = \sum_{i \in N} \frac{s[i, z]}{s[i, p]} > |S|.$$

The last inequality follows because the sum contains only non-negative terms, $|S|$ of which are at least 1, and one of which is strictly larger than 1. This is a contradiction. \square

Thus, the s -Nash rules are particularly fair to groups. The SD-core can also be seen as a *proportionality* requirement: the common resource should be divided so that the share of an alternative is proportional to its support. For example, this is of interest in politics, to divide parliament seats among parties.

5 Axiomatic Properties

We now briefly study other axiomatic properties of our rules. A more careful treatment appears in the full version; here we ignore ties when defining strategyproofness and monotonicity.

Pareto-efficiency. A distribution q SD-dominates p if $q \succ_i^{\text{SD}} p$ for all $i \in N$, and $q \succ_j^{\text{SD}} p$ for some $j \in N$. A distribution p is SD-efficient if no distribution dominates it. Note that SD-core implies SD-efficiency (with $S = N$), and so s -Nash rules are SD-efficient when $s_m = 0$. More generally, one can show that s -utilitarianism, s -leximin, and s -Nash are SD-efficient provided that s is strictly decreasing.

Strategyproofness. A social decision scheme is (*strongly*) SD-strategyproof if, when a voter misreports their ranking, the SDS selects a distribution that the voter believes is weakly SD-worse than the distribution resulting from a truthful report. Plurality-Nash (i.e., random dictatorship) is strategyproof in this sense. A well-known result of Gibbard [1977] shows that this is the only SDS that is strategyproof and also anonymous and Pareto-efficient. Hence, all other SD-efficient rules we have considered are manipulable.

Monotonicity. An SDS F is *monotone* if, when we change a profile P into P' by moving up an alternative x in a voters' ranking (by swapping), then the share of x weakly increases, i.e., $F(P')_x \geq F(P)_x$. This is clearly satisfied by s -utilitarianism, and also by plurality-Nash. However, other s -Nash rules (and also s -leximin) may fail it. If $s = (2, 1, 0)$ and $P = (abc, abc, abc, acb, bac, cba)$, then s -Nash selects an irrational distribution which rounds to $0.642a + 0.333b + 0.024c$. If the bac voter moves c up one place (to get bca), then s -Nash selects $0.5a + 0.5b$. Thus, c 's share has strictly decreased. Monotonicity is a kind of fairness to alternatives (x gets more if it performs better), while our rules aim for fairness to voters.

6 Conclusions

We have introduced a class of aggregation rules which can be used to make budget decisions. We have found that our rules are attractive on efficiency and fairness grounds. We have introduced concepts such as the SD-core, a group fairness and proportionality notion, satisfied by Nash-based rules. Egalitarian rules satisfy individual fairness, which may be particularly desirable in small electorates.

References

- [Aziz and Stursberg, 2014] H. Aziz and P. Stursberg. A generalization of probabilistic serial to randomized social choice. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 559–565. AAAI Press, 2014.
- [Aziz et al., 2017] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.
- [Aziz et al., 2018a] H. Aziz, B. E. Lee, and N. Talmon. Proportionally representative participatory budgeting: Axioms and algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 23–31. IFAAMAS, 2018.
- [Aziz et al., 2018b] H. Aziz, P. Luo, and C. Rizkallah. Rank maximal equal contribution: A probabilistic social choice function. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 910–916, 2018.
- [Aziz et al., 2019] H. Aziz, A. Bogomolnaia, and H. Moulin. Fair mixing: the case of dichotomous preferences. In *Proceedings of the 20th ACM Conference on Economics and Computation (ACM-EC)*, 2019.
- [Bade and Segal-Halevi, 2018] S. Bade and E. Segal-Halevi. Fair and efficient division among families. *arXiv:1811.06684*, 2018.
- [Baumeister et al., 2016] D. Baumeister, S. Bouveret, J. Lang, N.-T. Nguyen, T. T. Nguyen, J. Rothe, and A. Saffidine. Positional scoring-based allocation of indivisible goods. *Autonomous Agents and Multi-Agent Systems*, pages 1–28, 2016.
- [Benade et al., 2017] G. Benade, S. Nath, A. D. Procaccia, and N. Shah. Preference elicitation for participatory budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 376–382. AAAI Press, 2017.
- [Bogomolnaia et al., 2005] A. Bogomolnaia, H. Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.
- [Brams and King, 2005] S. J. Brams and D. L. King. Efficient fair division: Help the worst off or avoid envy? *Rationality and Society*, 17(4):387–421, 2005.
- [Brandl et al., 2016] F. Brandl, F. Brandt, and H. G. Seedig. Consistent probabilistic social choice. *Econometrica*, 84(5):1839–1880, 2016.
- [Brandl et al., 2019] F. Brandl, F. Brandt, D. Peters, C. Stricker, and W. Suksompong. Donor coordination: Collective distribution of individual contributions. 2019. Working paper.
- [Brandt, 2018] F. Brandt. Collective choice lotteries: Dealing with randomization in economic design. In J.-F. Laslier, H. Moulin, R. Sanver, and W. S. Zwicker, editors, *The Future of Economic Design*. Springer-Verlag, 2018. Forthcoming.
- [Cabannes, 2004] Y. Cabannes. Participatory budgeting: a significant contribution to participatory democracy. *Environment and Urbanization*, 16(1):27–46, 2004.
- [Conitzer et al., 2017] V. Conitzer, R. Freeman, and N. Shah. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (ACM-EC)*, pages 629–646, 2017.
- [Darmann and Schauer, 2015] A. Darmann and J. Schauer. Maximizing Nash product social welfare in allocating indivisible goods. *European Journal of Operational Research*, 247(2):548–559, 2015.
- [Duddy, 2015] C. Duddy. Fair sharing under dichotomous preferences. *Mathematical Social Sciences*, 73:1–5, 2015.
- [Fain et al., 2016] B. Fain, A. Goel, and K. Munagala. The core of the participatory budgeting problem. In *Proceedings of the 12th International Workshop on Internet and Network Economics (WINE)*, Lecture Notes in Computer Science (LNCS), pages 384–399, 2016.
- [Fain et al., 2018] B. Fain, K. Munagala, and N. Shah. Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (ACM-EC)*, pages 575–592, 2018.
- [Faliszewski et al., 2017] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner voting: A new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 2. 2017.
- [Foley, 1970] D. K. Foley. Lindahl’s solution and the core of an economy with public goods. *Econometrica*, 38(1):66–72, 1970.
- [Gibbard, 1977] A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, 1977.
- [Goel et al., 2016] A. Goel, A. K. Krishnaswamy, S. Sakshuwong, and T. Aitamurto. Knapsack voting: Voting mechanisms for participatory budgeting. Unpublished manuscript, 2016.
- [Kreweras, 1965] G. Kreweras. Aggregation of preference orderings. In *Mathematics and Social Sciences I: Proceedings of the seminars of Menthon-Saint-Bernard, France (1–27 July 1960) and of Gössing, Austria (3–27 July 1962)*, pages 73–79, 1965.
- [Manurangsi and Suksompong, 2019] P. Manurangsi and W. Suksompong. Computing a small agreeable set of indivisible items. *Artificial Intelligence*, 268:96–114, 2019.
- [Moulin, 2003] H. Moulin. *Fair division and collective welfare*. MIT Press, 2003.
- [Papadimitriou et al., 2008] C. Papadimitriou, M. Schapira, and Y. Singer. On the hardness of being truthful. In *Proceedings of the 49th Symposium on Foundations of Computer Science (FOCS)*, pages 250–259. IEEE, 2008.