Balanced Donor Coordination

Felix Brandt[1], Matthias Greger[1], Erel Segal-Halevi[2], Warut Suksompong[3]

Abstract

Charity is typically done either by individual donors, who donate money to the charities that they support, or by centralized organizations such as governments or municipalities, which collect the individual contributions and distribute them among a set of charities. On the one hand, individual charity respects the will of the donors but may be inefficient due to a lack of coordination. On the other hand, centralized charity is potentially more efficient but may ignore the will of individual donors. We present a mechanism that combines the advantages of both methods by distributing the contribution of each donor in an efficient way such that no subset of donors has an incentive to redistribute their donations. Assuming Leontief utilities (i.e., each donor is interested in maximizing an individually weighted minimum of all contributions across the charities), our mechanism is group-strategyproof, preference-monotonic, contribution-monotonic, maximizes Nash welfare, and can be computed using convex programming.

The full version of this paper is available at `https://arxiv.org/abs/2305.10286`

1 Introduction

Private charity, given by individual donors to underprivileged people in their vicinity, has existed long before institutionalized charity via municipal or governmental organizations. Its main advantage is transparency—the donors know exactly where their money goes to, which may increase their willingness to donate. A major disadvantage of private charity is the lack of coordination: donors may donate to certain people or charities without knowing that these recipients have already received ample money from other donors. Centralized charity via governments or municipalities is potentially more efficient but, if not done carefully, may disrespect the will of the donors.

As an example, consider the following scenario involving two donors and four charities. The first donor is willing to contribute \$900 and supports charities A, B, and C, whereas the second donor is willing to contribute \$100 and supports charities C and D.

A central organization may collect the contributions of the donors and divide them equally among the four charities, so that each charity receives \$250. While this outcome is the most balanced possible for the charities, it goes against the will of the first donor, since \$150 of her contribution is used to support charity D.

By contrast, without any coordination, each donor may split her individual contribution equally between the charities that she approves. As a result, charities A and B receive \$300 each, charity C receives \$350, and charity D receives \$50. However, if the second donor knew that charity C would already receive \$300 from the first donor, she would probably prefer to donate more to charity D, for which she is the only contributor.

Our suggested mechanism would give \$300 to each of charities A, B, and C, and \$100 to charity D. This distribution can be implemented in such a way that the contribution of each donor only goes to charities that the donor approves, and subject to that, the donations are divided as equally as possible. The distribution can also be understood as recommendations

[1]Technical University of Munich, Germany
[2]Ariel University, Israel
[3]National University of Singapore, Singapore

to the individual donors: the first donor should distribute her contribution uniformly over charities A, B, and C whereas the second donor should transfer all contribution to charity D.

Evaluating and comparing donor coordination mechanisms requires some assumptions on the utility functions of the donors. In this paper, we assume that each donor's utility is given by the smallest amount of money allocated to one of the donor's approved charities. For example, for the distribution $(300, 300, 300, 100)$, the first agent's utility is 300 and the second agent's utility is 100. More generally, our model allows donors to attribute different values than merely 1 and 0 (which indicate approval and disapproval, respectively) to different charities. If a donor i values a project x at $v_{i,x}$, then i's utility from a distribution δ equals $\min_x \delta(x)/v_{i,x}$, where the minimum is taken over all projects x for which $v_{i,x} > 0$. Such utilities are known as *Leontief utilities* and are often studied in resource allocation problems [e.g., 4, 10, 14, 15, 17, 20, 23]. Whenever $v_{i,x} \in \{0, 1\}$ for all agents i and projects x, we refer to this as (Leontief) utility functions with *binary weights*.

Given the contribution and utility function of each donor, our goal is to distribute the money among the charities in a way that respects the individual donors' preferences. The idea of "respecting the donors' preferences" is captured by the notion of an *equilibrium distribution.* We say that a distribution is *in equilibrium* if it can be implemented by telling each donor how to distribute her contribution among the charities, such that the prescribed distribution maximizes the donor's utility given that the distributions of the other donors remain fixed. One can check that, in the above example, the unique equilibrium distribution is $(300, 300, 300, 100)$.

A priori, it is not clear that an equilibrium distribution (in pure strategies) always exists. Our first main result is that each profile admits a *unique* equilibrium distribution. Moreover, we prove that the unique equilibrium distribution coincides with the unique distribution that maximizes the product of individual utilities weighted by their contributions (*Nash welfare*), which implies that it is Pareto efficient, and can be computed via convex programming.

In our example, the equilibrium distribution $(300, 300, 300, 100)$ also maximizes the minimum utility of all agents (*egalitarian welfare*) subject to each donor only contributing to her approved charities. We show that this is true in general when weights are binary, and extends to an infinite class of welfare measures "in between" Nash welfare and egalitarian welfare. Moreover, for the case of binary weights, we show that the equilibrium distribution coincides with the distribution that allocates individual contributions to approved projects such that the minimum contribution to projects is maximized lexicographically. This allows for simpler computation via linear programming. Further, we propose a simple dynamics for binary weights, based on best responses, that converges to the equilibrium distribution.

Based on existence and uniqueness, we can define the *equilibrium distribution rule (EDR)*—the mechanism that returns the unique equilibrium distribution of a given profile. Our second main result is that *EDR* exhibits remarkable axiomatic properties:

- *Group-strategyproofness*: agents and coalitions thereof are never better off by misrepresenting their preferences, and are strictly better off by contributing more money,
- *Preference-monotonicity*: the amount donated to a project can only increase when agents increase their valuation for the project, and
- *Contribution-monotonicity*: the amount donated to a project can only increase when agents increase their contributions.

Our results can be applied not only to private charity, but also to donation programs—prominent examples include *AmazonSmile*[4] and *cinque per mille* [5] by the Italian Revenue

[4]The program generated over $15 million in donations before being discontinued in 2023.

[5]`https://www.agenziaentrate.gov.it/portale/web/guest/contributo-del-5-per-mille-2022`

Agency. In these programs, participants can redirect a portion of their payments (purchase price and income tax, respectively) to charitable organizations of their choice.[6] Note that, in contrast to private charity, participants of donation programs do not have the option of taking their money out of the system, which means that the important issue lies in finding a desirable distribution of the contributions rather than in incentivizing the participants to take part in the donation exercise in the first place.

Another example of a potential application is the transmission of signals in a network. Consider a directed graph and a set of agents where each agent intends to transmit a signal along an individual path in the graph. Agents are able to invest in the “transmission quality” of each edge. Their utilities are given by the quality of the signal at the last vertex on their path, which equals the minimal transmission quality of an edge along that path.

The remainder of this paper is structured as follows. After discussing related work in Section 2, we formally introduce our model in Section 3. Section 4 lays the foundation for the proposed distribution rule by showing existence and uniqueness of equilibrium distributions as well as characterizing Pareto efficient distributions in our setting. Subsequently, we define the *equilibrium distribution rule* as the rule that always returns the equilibrium distribution and examine it axiomatically in Section 5. The special case of Leontief utilities with binary weights is covered in Section 6. Binary weights allow for alternative characterizations of *EDR* that enable its computation via linear programming. Furthermore, we justify *EDR* from a welfare point of view and present a simple best response dynamics that converges to the equilibrium distribution. Finally, we conclude in Section 7 and point out further directions for future research.

All proofs appear in the full version of our paper.

2 Related work

The work most closely related to ours is that of Brandl et al. [8, 9] who initiated the axiomatic study of donor coordination mechanisms. In their model, the utility of each donor is defined as the weighted *sum* of contributions to projects, where the weights correspond to the agent’s inherent utilities for a unit of contribution to each project. Under this assumption, the only efficient distribution in the introductory example is to allocate the entire donation of \$1000 to charity C, since this distribution gives the highest possible utility, 1000, to all donors. However, this distribution leaves charities A, B, and D with no money at all, which may not be what the donors intended. With sum-based utilities, as studied by Brandl et al., charities are perfect *substitutes*: when a donor assigns the same utility to several charities, she is completely indifferent to how money is distributed among these charities. By contrast, in our model of *minimum-based* utilities, charities are perfect *complements*: donors want their money to be evenly distributed among charities they like equally much. Preferences over charities can be expressed by setting weights for Leontief utility functions. It can be argued that this assumption better reflects the spirit of charity by not leaving anyone behind. The modified definition of utility functions critically affects the nature of elementary concepts such as efficiency or strategyproofness and fundamentally changes the landscape of attractive mechanisms.

The main result by Brandl et al. [9] shows that, in their model of sum-based utilities, the Nash product rule incentivizes agents to contribute their entire budget, even when attractive outside options are available. On the other hand, the Nash product rule fails to be strategyproof [3] and violates simple monotonicity conditions [8]. In fact, a sweeping

[6] For AmazonSmile and cinque per mille, each participant can choose only one charitable organization. However, as Brandl et al. [9] argued, permitting them to indicate support for multiple organizations can increase efficiency of the process.

impossibility by Brandl et al. [8] shows that, even in the simple case of binary valuations, no distribution rule that spends money to at least one approved project of each agent can simultaneously satisfy efficiency and strategyproofness. This confirms a conjecture by Bogomolnaia et al. [7] and demonstrates the severe limitations of donor coordination with sum-based utilities. Interestingly, as we show in this paper, Leontief utilities allow for much more positive results.

Originating from the *Nash bargaining solution* [16], the Nash product rule can be interpreted as a tradeoff between maximizing utilitarian and egalitarian welfare, a recurring idea when it comes to finding efficient *and* fair solutions. When allocating divisible private goods to agents with additive valuations, the Nash product rule returns the set of all *competitive equilibria from equal incomes* [12]; thus, it results in an efficient and *envy-free* allocation [13]. For the case of indivisible private goods, Caragiannis et al. [11] showed that maximizing Nash welfare returns an allocation that is not only efficient but also satisfies *envy-freeness up to one good*, and Suksompong [22] and Yuen and Suksompong [24] obtained characterizations of the Nash product rule using the latter axiom. The Nash product rule is also a sensible mechanism in our context and, as shown in Section 4, its outcome is completely characterized by another concept due to Nash [16]: when defining a game in which the players' strategies are redistributions of their individual contributions, there is a unique Nash equilibrium which coincides with the distribution maximizing Nash welfare.

A natural special case of our model is that of Leontief utilities with *binary weights*, where agents only approve or disapprove projects and the utility of each agent is given by the minimal amount transferred to any of her approved projects. Under the assumption that agents only contribute to projects they approve and that all individual contributions are equal, this can be interpreted as a (many-to-many) matching problem on a bipartite graph where agents (and their contributions) need to be assigned to projects with unlimited capacity. Bogomolnaia and Moulin [6] proposed a solution to such matching problems that maximizes egalitarian welfare of the projects (rather than the agents). The reasons for the intriguing connection between these two types of egalitarianism are addressed in Section 6. These authors also showed that their solution constitutes a competitive equilibrium from equal incomes (from the project holders' point of view).

A problem remotely related to the setting we study in this paper is that of *private provision of public goods* [e.g., 5, 21]. In this stream of research, each agent decides on how much money she wants to contribute to funding a public good. Typically, this leads to under-provision of the public good in equilibrium, resulting in inefficient outcomes. In our model, we assume that agents have already set aside a budget to support public projects, either voluntarily or compulsorily (as part of their taxes or payments to a company). The inefficiency that we are worried about is an inefficient allocation among different public goods. As a result, our problem has the flavor of both social choice and fair division.

Finally, in *participatory budgeting* [e.g., 2], it is typically assumed that project funding is discrete, that is, each project has a fixed cost (e.g., constructing a new bridge), and it can be either fully funded or not at all. Moreover, most participatory budgeting papers assume that money is owned by a central authority rather than by the agents. Aziz and Ganguly [1] studied a donor coordination version where the budget belongs to the agents, but still considered discrete project funding. This is in contrast to our setting, in which each charity can receive any amount of money, and there are no pre-specified costs.

3 The model

Let N be a set of n *agents*. Each agent i contributes an amount $C_i \geq 0$. For every subset of agents $N' \subseteq N$, we denote $C_{N'} := \sum_{i \in N'} C_i$. The sum of all contributions, C_N, is called the *endowment*.

Further, consider a set A of m potential recipients of the contributions, which we call *projects*. A *distribution* is a function δ assigning a nonnegative real number to each project, such that $\sum_{x\in A}\delta(x) = C_N$. The support $\{x\colon \delta(x) > 0\}$ of δ is denoted by $\mathrm{supp}(\delta)$, and the space of all possible distributions is denoted by $\Delta(C_N)$. For a subset of projects $A' \subseteq A$, we define $\delta(A') := \sum_{x\in A'}\delta(x)$ as the total amount allocated to projects in A'.

For every $i \in N$ and $x \in A$, there is a real number $v_{i,x} \geq 0$ that represents the value of project x to agent i. We assume that each agent i has at least one project x for which $v_{i,x} > 0$. For every agent $i \in N$, we define $A_i := \{x\colon v_{i,x} > 0\}$ as the set of projects to which i attributes a positive value.

The utility that agent i derives from distribution δ is denoted by $u_i(\delta)$ and is given by the Leontief utility function[7]

$$u_i(\delta) = \min_{x\in A_i}\frac{\delta(x)}{v_{i,x}}.$$

Note that, for every project $x \in A$ and every agent $i \in N$, $\delta(x) \geq v_{i,x}\cdot u_i(\delta)$.

If all $v_{i,x}$ are in $\{0, 1\}$, we refer to Leontief utilities with *binary* weights. A *profile* P consists of $\{C_i\}_{i\in N}$ and $\{v_{i,x}\}_{i\in N, x\in A}$. Throughout this paper, agents with contribution zero do not have any influence on the outcome and can thus be treated as agents who choose not to participate in the mechanism. A *distribution rule* f maps every profile to a distribution $\Delta(C_N)$ of the total endowment C_N.

4 Equilibrium distributions

Donor coordination differs from other participatory budgeting problems in that the budget is initially owned by the agents. This makes it all the more important that agents are able to observe the distribution of their individual contribution. We formalize this intuition using the notion of a decomposition.

Definition 4.1 (Decomposition)**.** A *decomposition* of a distribution δ is a vector of distributions $(\delta_i)_{i\in N}$ with

$$\sum_{i\in N}\delta_i(x) = \delta(x) \qquad \text{for all } x \in A; \tag{1}$$

$$\sum_{x\in A}\delta_i(x) = C_i \qquad \text{for all } i \in N. \tag{2}$$

Note that each distribution admits at least one decomposition. We aim for a decomposition in which no agent can increase her utility by changing δ_i, given C_i and the δ_j for $j \neq i$. In other words, we look for a pure strategy Nash equilibrium of the game in which the strategy space of each agent i is the set of δ_i satisfying (2).

Definition 4.2 (Equilibrium distribution)**.** A distribution δ is *in equilibrium* if it admits a decomposition $(\delta_i)_{i\in N}$ such that, for every agent i and for every alternative distribution δ'_i satisfying $\sum_{x\in A}\delta'_i(x) = C_i$,

$$u_i(\delta) \geq u_i(\delta - \delta_i + \delta'_i).$$

A priori, it is not clear whether an equilibrium distribution always exists. The present section is devoted to proving the following theorem.

[7] The case of *sum-based*, rather than *min-based*, utility functions $u_i(\delta) = \sum_{x\in A} v_{i,x}\cdot\delta(x)$ is discussed in Section 2.

Theorem 4.3. *Every profile admits a unique equilibrium distribution. This distribution is Pareto efficient and can be computed via convex programming.*

As a consequence, we can define the *equilibrium distribution rule* as the distribution rule that selects for each profile its unique equilibrium distribution. In Section 5, we prove that this rule satisfies desirable strategic and monotonicity properties.

4.1 Critical projects

We start by characterizing equilibrium distributions based on *critical projects.*

Given a distribution δ, we define the set of agent i's *critical projects*

$$T_{\delta,i} := \arg\min_{x \in A_i} \frac{\delta(x)}{v_{i,x}}.$$

Each project $x \in T_{\delta,i}$ is critical for agent i in the sense that the utility of i would decrease if the amount allocated to x were to decrease. Every agent has at least one critical project. For every agent i and project x such that either $v_{i,x} > 0$ or $\delta(x) > 0$, the following implications hold:

$$\begin{aligned} x \in T_{\delta,i} \quad &\Leftrightarrow \quad \delta(x) = v_{i,x} \cdot u_i(\delta); \\ x \notin T_{\delta,i} \quad &\Leftrightarrow \quad \delta(x) > v_{i,x} \cdot u_i(\delta). \end{aligned} \tag{3}$$

We establish that a distribution is in equilibrium if and only if each agent contributes only to her critical projects.

Lemma 4.4. *[Characterization of equilibrium distributions] A distribution δ is in equilibrium if and only if it has a decomposition $(\delta_i)_{i \in N}$ such that $\delta_i(x) = 0$ for every project $x \notin T_{\delta,i}$. Equivalently, it has a decomposition satisfying the following, instead of* (2)*:*

$$\sum_{x \in T_{\delta,i}} \delta_i(x) = C_i \qquad \textit{for all } i \in N. \tag{4}$$

Corollary 4.5. *In an equilibrium distribution, every project that receives a positive amount of contribution is critical for at least one agent.*

Remark 4.6. Lemma 4.4 implies that an equilibrium distribution satisfies an even stronger equilibrium property: no *group of agents* can deviate without making at least one of its members worse off. This is because any deviation decreases the contribution to a critical project of at least one group member.

4.2 Efficiency

One of the main objectives of a centralized distribution rule is economic efficiency.

Definition 4.7 (Efficiency)**.** Given a profile P, a distribution $\delta \in \Delta(C_N)$ is *(Pareto) efficient* if there does not exist another distribution $\delta' \in \Delta(C_N)$ that *(Pareto) dominates* δ, i.e., $u_i(\delta') \geq u_i(\delta)$ for all $i \in N$ and $u_i(\delta') > u_i(\delta)$ for at least one $i \in N$. A distribution rule is efficient if it returns an efficient distribution for every profile P.

The following lemma characterizes efficient distributions of an arbitrary profile.

Lemma 4.8. *[Characterization of efficient distributions] A distribution δ is efficient if and only if every project $x \in \operatorname{supp}(\delta)$ is critical for some agent.*

Combining Corollary 4.5 with Lemma 4.8 gives the following implication.

Corollary 4.9. *Every equilibrium distribution is efficient.*

The following lemma shows that every efficient utility vector is generated by at most one distribution.

Lemma 4.10. *Let δ and δ' be efficient distributions inducing the same utility vector, that is, $u_i(\delta) = u_i(\delta')$ for all $i \in N$. Then, $\delta = \delta'$.*

Consequently, an efficient distribution rule can be defined as mapping a profile to a utility vector.

4.3 Existence, uniqueness, and computation

One common way to obtain an efficient distribution is to maximize a welfare function. Formally, for any strictly-increasing function g on $\mathbb{R}_{\geq 0}$, we say that a distribution δ is *g-welfare-maximizing* if it maximizes the weighted sum $\sum_{i \in N} C_i \cdot g(u_i(\delta))$. Clearly, any such distribution is efficient. Whenever g is strictly concave, there is a unique g-welfare-maximizing distribution; the proof is straightforward and is given in the full version of our paper.

We focus on the special case in which g is the log function. The *Nash welfare* of a distribution δ is defined as the sum of logarithms of the agents' utilities, weighted by their contributions:

$$Nash(\delta) := \sum_{i \in N} C_i \cdot \log u_i(\delta).$$

The *Nash rule* selects a distribution δ that maximizes $Nash(\cdot)$ or, equivalently, the weighted product of the agents' utilities $\prod_{i \in N} u_i^{C_i}$ (with the convention $0 \log 0 = 0$ and $0^0 = 1$). The Nash rule has strong fairness guarantees in various settings (see Section 2). As we will see, this is also the case in our model.

The following two lemmas show that a distribution is in equilibrium if and only if it maximizes Nash welfare.

Lemma 4.11. *Every distribution that maximizes Nash welfare is in equilibrium.*

Lemma 4.12. *Every equilibrium distribution maximizes Nash welfare.*

Since the log function is strictly concave, there is a unique distribution that maximizes Nash welfare. Therefore, Lemmas 4.11 and 4.12 imply that there is a unique equilibrium distribution, and it is efficient, as claimed in Theorem 4.3.

Since the equilibrium distribution maximizes a weighted sum of logarithms, it can be approximated arbitrarily well by considering the corresponding convex optimization problem. With sum-based utilities, Brandl et al. [9] show that it is impossible to compute the Nash-optimal distribution exactly even for binary valuations, since it may involve irrational numbers. Interestingly, for Leontief utilities the Nash-optimal distribution is rational whenever the agents' valuations and contributions are rational; see the full version of our paper for a proof.

In the case of binary weights, the equilibrium distribution can be computed exactly using a polynomial number of linear programs; see Section 6. We do not know whether the same is true for non-binary weights.

5 The Equilibrium Distribution Rule

Based on Theorem 4.3, we define the *equilibrium distribution rule (EDR)* as the distribution rule that, for each profile, returns the unique equilibrium distribution for this profile. In this section, we investigate the axiomatic properties of *EDR*.

5.1 Strategyproofness

A distribution rule is *group-strategyproof* if no coalition of agents can gain utility by misreporting their valuations or contributing less. This incentivizes truthful reports and allows for a correct estimation of agents' utilities under different distributions. Furthermore, a group-strategyproof rule ensures that every agent donates the maximal possible contribution, thereby guaranteeing maximal gains from coordination.

Definition 5.1 (Group-strategyproofness)**.** (a) Given a distribution rule f, a profile P, and a group $G \subseteq N$, a profile P' is called a *manipulation of P by G* if $C'_G \leq C_G$ (the contribution of G may decrease), and the valuations of agents in G may change, while the contributions and valuations of all agents in $N \setminus G$ remain the same. Such a manipulation is called *successful* if $u_j(f(P')) \geq u_j(f(P))$ for all $j \in G$ and $u_i(f(P')) > u_i(f(P))$ for at least one $i \in G$, where $\mathbf{u}$ refers to the utilities in P.

(b) A distribution rule f is *group-strategyproof* if in any profile, no group of agents has a successful manipulation.

Theorem 5.2. *EDR is group-strategyproof.*

In fact, the proof of Theorem 5.2 shows that if the total contribution C_G decreases, then the utility of at least one agent in G has to *strictly* decrease under *EDR*. In particular, an agent receives *strictly* more utility when she increases her contribution.

Theorem 5.3. *Under EDR, agents are strictly better off by increasing their contribution.*

Remark 5.4. In the context of sum-based utilities, Brandl et al. [8] have proposed an even stronger participation axiom called *contribution incentive-compatibility*. This axiom allows agents who do not participate in the mechanism to receive additional utility by spending her saved contribution independently. Unfortunately, in our setting, this property is incompatible with efficiency and also with strategyproofness, even for binary weights. For more details, we refer to the full version of our paper.

5.2 Preference-monotonicity

An important property for project holders is *preference-monotonicity*, which requires that for every agent i and project $x \in A$, $\delta(x)$ weakly increases when $v_{i,x}$ increases. In other words, a project can only receive more donations when becoming more popular, which, for example, incentivizes advertising projects.

Definition 5.5 (Preference-monotonicity)**.** A distribution rule f satisfies *preference-monotonicity* if for every two profiles P and P' which are identical except that $v'_{i,x} > v_{i,x}$ for one agent i and one project x, we have $f(P')(x) \geq f(P)(x)$.

Theorem 5.6. *EDR satisfies preference-monotonicity.*

5.3 Contribution-monotonicity

For some applications, it would be desirable if increased contributions do not result in the redistribution of funds that have already been allocated. For example, if agents arrive over time or increase their contributions over time, ideally the mechanism only needs to take care of the additional contributions. This would allow a deployment of the mechanism as an ongoing process in which donations arrive over time and projects can make immediate use of the donations they receive. We formalize this property in the following definition.

Definition 5.7 (Contribution-monotonicity)**.** A distribution rule f satisfies *contribution-monotonicity* if for every two profiles P and P' where P' can be obtained from P by increasing the contribution of one agent (possibly from 0), $f(P')(x) \geq f(P)(x)$ for all projects $x \in A$.

Theorem 5.8. *EDR satisfies contribution-monotonicity.*

6 Leontief utilities with binary weights

In this section, we consider the special case of having binary weights, i.e., $v_{i,x} \in \{0, 1\}$ for all agents $i \in N$ and projects $x \in A$. Equivalently, each agent i has a non-empty set of *approved projects* $A_i \subseteq A$ and her utility from a distribution δ is

$$u_i(\delta) = \min_{x \in A_i} \delta(x).$$

For each project $x \in A$, we denote by $N_x \subseteq N$ the set of agents who approve project x. Note that, for every project $x \in A$ and every agent $i \in N_x$,

$$\delta(x) \geq u_i(\delta). \tag{5}$$

Binary weights allow for further insights into the structure of the equilibrium distribution, which in turn yield new interpretations and additional properties of *EDR*.

For sum-based utilities with binary weights, a distribution is in equilibrium if and only if each agent contributes only to projects she approves. Brandl et al. [8] refer to this axiom as *decomposability*.

Definition 6.1 (Decomposable distribution)**.** Given a profile with binary weights ($v_{i,x} \in \{0, 1\}$), a distribution δ is *decomposable* if it has a decomposition $(\delta_i)_{i\in N}$ such that $\delta_i(x) = 0$ for every project $x \notin A_i$. Equivalently, it has a decomposition satisfying the following, instead of (2):

$$\sum_{x \in A_i} \delta_i(x) = C_i \qquad \text{for all } i \in N.$$

The equivalence of decomposable distributions and equilibrium distributions no longer holds with Leontief utilities: there are decomposable distributions that are not in equilibrium even when there is only one agent.

Example 6.2. There is a single agent with $C_1 = 2$, $A = \{a, b\}$, with valuations $v_{1,a} = 1$ and $v_{1,b} = 1$. The distribution $\delta = (2, 0)$ is decomposable, but it is not in equilibrium, since the single agent is better off by the equilibrium distribution $\delta' = (1, 1)$.

Nevertheless, decomposability can be used to establish two alternative interpretations of *EDR* for binary weights.

6.1 Egalitarianism for projects

Motivated by the example from the introduction, we aim at a rule which distributes money on the projects as equally as possible while still respecting the preferences of the donors. One rule that comes to mind selects a distribution that, among all decomposable distributions, maximizes the smallest amount allocated to a project. Subject to this, it maximizes the second-smallest allocation to a project, and so on. We define it formally below.

Definition 6.3. Given two vectors $\mathbf{x}, \mathbf{y}$ of the same size, we say that $\mathbf{x}$ *is leximin-higher than* $\mathbf{y}$ (denoted $\mathbf{x} \succ_{lex} \mathbf{y}$) if the smallest value in $\mathbf{x}$ is larger than the smallest value in $\mathbf{y}$; or the smallest values are equal, and the second-smallest value in $\mathbf{x}$ is larger than the second-smallest value in $\mathbf{y}$; and so on. $\mathbf{x} \succeq_{lex} \mathbf{y}$ means that either $\mathbf{x} \succ_{lex} \mathbf{y}$ or the multiset of values in $\mathbf{x}$ is the same as that in $\mathbf{y}$.

Definition 6.4. The *project egalitarian rule* selects a distribution δ^* that, among all decomposable distributions, maximizes the distribution vector by the leximin order, that is: $\delta^* \succeq_{lex} \delta$ for every decomposable distribution δ.

The leximin order on the closed and convex set of decomposable distributions is connected, every two vectors are comparable, and there exists a unique maximal element (otherwise, any convex combination of two different maximal elements would be leximin-higher than the maximal elements). Therefore, the project egalitarian rule selects a unique distribution and is well-defined. The following theorem states that the returned distribution is the equilibrium distribution, resulting in an alternative characterization of *EDR* for binary weights.

Theorem 6.5. *With binary weights, the project egalitarian rule and EDR are equivalent.*

Remarkably, this new interpretation of *EDR* ignores the Leontief utilities of the agents and does not directly take into account the different contributions. Instead, they enter indirectly through the constraints induced by decomposability.

Theorem 6.5 implies that *EDR* can be computed by solving the following program, with variables δ_x for all $x \in A$ and $\delta_{i,x}$ for all $i \in N, x \in A$:

$$\begin{aligned}
&\operatorname{lex\,max\,min}\{\delta_x\}_{x\in A} && \text{subject to} \\
&\delta_x = \sum_{i\in N} \delta_{i,x} && \forall x \in A \\
&\sum_{x\in A_i} \delta_{i,x} = C_i && \forall i \in N \\
&\delta_{i,x} \geq 0, \qquad \delta_x \geq 0 && \forall i \in N, \qquad \forall x \in A_i
\end{aligned}$$

where "lex max min" refers to finding a solution vector that is maximal in the leximin order subject to the constraints, and the second constraint represents decomposability. It is well-known that such leximin optimization with k objectives and linear constraints can be solved by a sequence of poly(k) linear programs.[8]

Corollary 6.6. *With binary weights, the equilibrium distribution can be computed by solving at most $|A|$ linear programs.*

[8] Ogryczak et al. [19] showed that every leximin optimization problem with k objectives has an equivalent *lexicographic* optimization problem, denoted (32) in their paper, with $k^2 + k$ additional variables and k^2 additional constraints. In a lexicographic optimization problem, the objectives have a fixed priority order. The goal is to maximize the most important objective, then subject to this, maximize the second most important objective, and so on. A lexicographic optimization problem can be solved by a simple sequential algorithm using at most k linear programs (Algorithm 1 in their work). For the special case of a convex optimization problem, Ogryczak et al. [19] presented Algorithm 4, which solves the problem using at most k linear programs without additional variables and constraints.

6.2 Egalitarianism for agents

While *EDR* is egalitarian from the point of view of the projects, one could also consider a rule that is egalitarian from the point of view of the agents. The *conditional egalitarian rule* (CEG) aims to balance the agents' utilities without disregarding their approvals. It selects a decomposable distribution that, among all decomposable distributions, maximizes the utility vector by the leximin order, that is: $\mathbf{u}(\delta^{CEG}) \succeq_{lex} \mathbf{u}(\delta)$ for every decomposable distribution δ.

Theorem 6.7. *With binary weights, the CEG rule and EDR are equivalent.*

Theorem 6.7 implies that the equilibrium distribution can be computed by solving the following program, with variables u_i for all $i \in N$ and $\delta_{i,x}$ for all $i \in N, x \in A_i$.

$$
\begin{aligned}
\text{lex}\max\min\{u_i\}_{i\in N} & \qquad\qquad \text{subject to} \\
u_i \leq \delta_{i,x} & \qquad\qquad \forall i \in N, \quad \forall x \in A_i \\
\sum_{x \in A_i} \delta_{i,x} = C_i & \qquad\qquad \forall i \in N \\
\delta_{i,x} \geq 0, \quad u_i \geq 0 & \qquad\qquad \forall i \in N, \quad \forall x \in A_i
\end{aligned}
$$

Using the algorithms in the works by Ogryczak et al. [19] and Ogryczak and Śliwiński [18], this program can be solved using at most $|N|$ linear programs.

Thus, we have three algorithms for computing the equilibrium distribution in the case of binary weights: one requires at most $|A|$ linear programs; one requires at most $|N|$ linear programs; and one requires a single convex (non-linear) program. It would be interesting to check which of these algorithms is most efficient in practice.

6.3 Welfare functions maximized by *EDR*

Based on the observation that *EDR* coincides with both the Nash product rule and the conditional egalitarian rule for binary weights, a natural question to ask is which other welfare notions are maximized by *EDR* subject to decomposability.

For this, we take a closer look at g-welfare (see Section 4.3), but this time subject to decomposability. Clearly, every g-welfare-maximizing distribution is efficient. We prove that efficiency remains even if we maximize among the decomposable distributions.

Lemma 6.8. *Let g be any strictly-increasing function, and let δ be a distribution that maximizes the g-welfare among all decomposable distributions. Then δ is unique and efficient.*

The Nash product rule is often considered a compromise between maximizing utilitarian welfare ($\sum_{i\in N} C_i \cdot u_i$) and egalitarian welfare (maximize the utility of the agent with smallest utility; notice that the conditional egalitarian rule is a refinement). This can be seen when considering the family of g-welfare functions $\sum_{i\in N} C_i \cdot \text{sgn}(p) \cdot u^p$ for $p \neq 0$ where the limit $p \to 0$ corresponds to $\sum_{i\in N} C_i \cdot \log(u_i)$ and $p \to -\infty$ approaches egalitarian welfare.

It is interesting to check whether the equivalence between conditional egalitarian welfare and Nash welfare extends to a larger class of g-welfare functions. This is indeed the case, as the following theorem shows.

Theorem 6.9. *Assume $g : \mathbb{R}_{\geq 0} \to \mathbb{R} \cup \{-\infty\}$*

1. *is strictly increasing on $\mathbb{R}_{\geq 0}$ and differentiable on $\mathbb{R}_{>0}$, and*

2. *$xg'(x)$ is non-increasing on $\mathbb{R}_{>0}$.*

Then the equilibrium distribution maximizes g-welfare among all decomposable distributions.

Property (1) ensures that the social welfare is indeed increasing when an individual's utility increases and small changes in individual utilities only cause small changes in the total social welfare. Property (2) implies that increasing utilities are discounted "at least logarithmically" when being translated to welfare.

In particular, Theorem 6.9 holds for all g-welfare functions $\sum_{i \in N} C_i \cdot \mathrm{sgn}(p) \cdot u^p$ with $p < 0$. However, it ceases to hold when $p > 0$, as the following proposition shows.

Proposition 6.10. *For each $p > 0$, maximizing the g-welfare with respect to $g(u) = u^p$ subject to decomposability does not always return the equilibrium distribution.*

Theorem 6.9 stresses the fact that *EDR* can be motivated not only from a game-theoretic and axiomatic point of view, but also from a welfarist perspective.

6.4 Convergence to equilibrium

In this section, we propose a simple best-response-based spending dynamics for binary weights that converges to the equilibrium distribution δ^*. This enables a decentralized implementation in which agents do not have to reveal their preferences explicitly.

Denote by δ^t the distribution at time step t (along with its associated decomposition), e.g., δ^0 equals the null vector as no agent $i \in N$ has yet distributed her contribution C_i. At each time step t, allow one agent i_t to contribute or redistribute her entire contribution in such a way that her utility is maximized for the new distribution δ^{t+1}, i.e.,

$$\delta_{i_t}^{best} := \arg \max_{\delta_{i_t} \in \Delta(C_{i_t})} u_{i_t} \left(\delta_{i_t} + \sum_{j \neq i_t} \delta_j^t \right);$$
$$\delta^{t+1} = \delta_{i_t}^{best} + \sum_{j \neq i_t} \delta_j^t.$$

Lemma 6.11. *For every time step t and agent i_t, there exists a unique best response $\delta_{i_t}^{best}$.*

Theorem 6.12. *Let $\mathcal{S} = (i_0, i_1, i_2, \dots)$ be an infinite sequence of agents updating their individual distributions by best responses. If each agent $i \in N$ appears infinitely often in $\mathcal{S}$, the dynamics converges to the equilibrium distribution, that is, $\lim_{t \to \infty} \delta^t = \delta^*$.*

Just like the question of whether the equilibrium distribution can be computed by a linear program for general Leontief utilities, it is open whether the best response dynamics converges to the equilibrium distribution in the case of general Leontief utilities.

7 Conclusion and further directions

All in all, *EDR* turns out to be an exceptionally attractive rule for donor coordination with Leontief utilities. It satisfies virtually all desirable properties one could hope for and can be computed via convex programming. In the case of binary weights, *EDR* maximizes a wide range of possible welfare functions and can be computed via linear programming or a simple spending dynamics. These results stand in sharp contrast to the previously studied case of sum-based utilities, where a far-reaching impossibility has shown the incompatibility of efficiency, strategyproofness, and a very weak form of fairness [8].

An interesting question for future work is to find more general types of utility functions for which maximizing Nash welfare results in equilibrium distributions.

Acknowledgments

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/11-1 and BR 2312/12-1, by the Israel Science Foundation under grant number 712/20, by the Singapore Ministry of Education under grant number MOE-T2EP20221-0001, and by an NUS Start-up Grant. We are grateful to Florian Brandl for helpful discussions including the proposal of the best response dynamics together with a proof idea, and the participants of the 3rd Ariel Conference on the Political Economy of Public Policy (September 2022), the joint Microeconomics Seminar of ETH Zurich and the University of Zurich (March 2023), and the Bar-Ilan Computer Science Seminar (April 2023) for their insightful comments.

References

[1] Haris Aziz and Aditya Ganguly. Participatory funding coordination: Model, axioms and rules. In *Proceedings of the 7th International Conference on Algorithmic Decision Theory (ADT)*, pages 409–423, 2021.

[2] Haris Aziz and Nisarg Shah. Participatory budgeting: Models and approaches. In Tamás Rudas and Gábor Péli, editors, *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*, pages 215–236. Springer International Publishing, 2021.

[3] Haris Aziz, Anna Bogomolnaia, and Hervé Moulin. Fair mixing: the case of dichotomous preferences. *ACM Transactions on Economics and Computation*, 8(4):18:1–18:27, 2020.

[4] Xiaohui Bei, Zihao Li, and Junjie Luo. Fair and efficient multi-resource allocation for cloud computing. In *Proceedings of the 18th International Conference on Web and Internet Economics (WINE)*, pages 169–186, 2022.

[5] Theodore Bergstrom, Lawrence Blume, and Hal Varian. On the private provision of public goods. *Journal of Public Economics*, 29(1):25–49, 1986.

[6] Anna Bogomolnaia and Hervé Moulin. Random matching under dichotomous preferences. *Econometrica*, 72(1):257–279, 2004.

[7] Anna Bogomolnaia, Hervé Moulin, and Richard Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.

[8] Florian Brandl, Felix Brandt, Dominik Peters, and Christian Stricker. Distribution rules under dichotomous preferences: Two out of three ain't bad. In *Proceedings of the 22nd ACM Conference on Economics and Computation (ACM-EC)*, pages 158–179, 2021.

[9] Florian Brandl, Felix Brandt, Matthias Greger, Dominik Peters, Christian Stricker, and Warut Suksompong. Funding public projects: A case for the Nash product rule. *Journal of Mathematical Economics*, 99:102585, 2022.

[10] Simina Brânzei, Yuezhou Lv, and Ruta Mehta. To give or not to give: Fair division for single minded valuations. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 123–129, 2016.

[11] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation*, 7(3):12:1–12:32, 2019.

[12] Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The pari-mutuel method. *Annals of Mathematical Statistics*, 30(1):165–168, 1959.

[13] Duncan Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7: 45–98, 1967.

[14] Ashish Goel, Reyna Hulett, and Benjamin Plaut. Markets beyond Nash welfare for Leontief utilities. In *Proceedings of the 15th International Conference on Web and Internet Economics (WINE)*, page 340, 2019.

[15] Jin Li and Jingyi Xue. Egalitarian division under Leontief preferences. *Economic Theory*, 54(3):597–622, 2013.

[16] John F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences (PNAS)*, 36:48–49, 1950.

[17] Antonio Nicoló. Efficiency and truthfulness with Leontief preferences. A note on two-agent, two-good economies. *Review of Economic Design*, 8(4):373–382, 2004.

[18] Włodzimierz Ogryczak and Tomasz Śliwiński. On direct methods for lexicographic min-max optimization. In *Proceedings of the 6th International Conference on Computational Science and Its Applications (ICCSA), Part III*, pages 802–811, 2006.

[19] Włodzimierz Ogryczak, Michał Pióro, and Artur Tomaszewski. Telecommunications network design and max-min optimization problem. *Journal of Telecommunications and Information Technology*, (3):43–56, 2005.

[20] David C. Parkes, Ariel D. Procaccia, and Nisarg Shah. Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. *ACM Transactions on Economics and Computation*, 3(1):3:1–3:22, 2015.

[21] Paul A. Samuelson. The pure theory of public expenditure. *The Review of Economics and Statistics*, 36(4):387–389, 1954.

[22] Warut Suksompong. A characterization of maximum Nash welfare for indivisible goods. *Economics Letters*, 222:110956, 2023.

[23] Steven Yin, Shatian Wang, Lingyi Zhang, and Christian Kroer. Dominant resource fairness with meta-types. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 486–492, 2021.

[24] Sheung Man Yuen and Warut Suksompong. Extending the characterization of maximum Nash welfare. *Economics Letters*, 224:111030, 2023.

Felix Brandt
Technical University of Munich
Munich, Germany
Email: `brandt@tum.de`

Matthias Greger
Technical University of Munich
Munich, Germany
Email: `matthias.greger@tum.de`

Erel Segal-Halevi
Ariel University
Ariel, Israel
Email: erelsgl@gmail.com

Warut Suksompong
National University of Singapore
Singapore, Singapore
Email: warut@comp.nus.edu.sg