Computing Lindahl Equilibrium with and without Funding Caps

Christian Kroer[1] , **Dominik Peters**[2]
[1]Columbia University
[2]CNRS, LAMSADE, Université Paris Dauphine - PSL
ck2945@columbia.edu, dominik.peters@lamsade.dauphine.fr

Abstract

Lindahl equilibrium is a solution concept for allocating a fixed budget across several divisible public goods. It always lies in the core, meaning that the equilibrium allocation satisfies desirable stability and proportional fairness properties. In the *uncapped setting*, each of the public goods can absorb any amount of funding. In this case, it is known that Lindahl equilibrium is equivalent to maximizing Nash social welfare, and this allocation can be computed by a public-goods variant of the proportional response dynamics. We introduce a new convex programming formulation for computing this solution and show that it is related to Nash welfare maximization through duality and reformulation. We then show that the proportional response dynamics is equivalent to running mirror descent on our new formulation, thereby providing a new and very immediate proof of the convergence guarantee for the dynamics. In the *capped setting*, each public good has an upper bound on the amount of funding it can receive, which is a type of constraint that appears in fractional committee selection and participatory budgeting. We prove that our new convex program continues to work when the cap constraints are added, and its optimal solutions are Lindahl equilibria. Thus, we establish that Lindahl equilibrium can be efficiently computed in the capped setting.

1 Introduction

We consider a setting where a fixed budget $B > 0$ needs to be spent on m divisible public goods. Thus, an outcome is a vector $x = (x_1, \dots, x_m) \in \mathbb{R}^m_{\geqslant 0}$ summing to at most B. Some of the public goods may additionally have *caps*, i.e., upper bounds on the amount of funding they can receive. How to distribute the spending across the goods is decided based on the preferences of n agents. We will consider agents with separable linear utility functions over the goods. Agents may have heterogeneous weights (which can be interpreted as endowments). We will study the solution concept of *Lindahl equilibrium*, which is based on a virtual market with personalized prices (Foley, 1970). This equilibrium notion is known to lead to allocations that are fair to voters, in the sense of the core from cooperative game theory (Fain et al., 2016).

The classic economics literature on public goods, starting with Samuelson (1954), focusses on how to arrive at the socially efficient amount of spending in the face of free-riding incentives. In contrast, we consider a fixed budget and are mostly concerned with how to divide it between different public goods. This approach, sometimes called *portioning* or *fair mixing*, has received increasing attention in computer science over recent years (see, e.g., Fain et al., 2016; Aziz et al., 2020; Brandl et al., 2021; Airiau et al., 2023), due to its many concrete applications. These include *participatory budgeting*, a method used by many cities to let residents vote over how the government will spend a fixed part of its budget (Rey and Maly, 2023; Aziz and Shah, 2021), and *donation platforms*, where donors can influence the distribution of a fixed matching fund (Brandl et al., 2022; Brandt et al., 2024). The model also captures *committee elections* (i.e., multiwinner voting) in its fractional version (Aziz et al., 2023a; Suzuki and Vollen, 2024), as well as the *cake sharing* problem (Bei et al., 2024). Voting methods for the public goods model can also be used to settle small-stakes issues such as a lecturer letting students vote over the distribution of class time across topics, or a team to vote over the frequencies with which they will go to different lunch venues. Companies and non-profit organizations can use the principles derived in this model to decide how to fairly and efficiently divide resources among units or grantees.

In many of these applications, it is desirable to select an outcome that is *representative* of the voters, in that every agent has an equal influence on the overall spending (or an influence that is proportional to the weight assigned to the agent). This can be formalized as a group fairness guarantee. In particular, we can require that the spending distribution lie in the *core*, which means that no subset of voters can construct an alternative way of spending their endowments in a way that they all prefer. We know that a core outcome always exists thanks to Foley (1970), who gave a definition of what he called *Lindahl equilibrium* (because he was inspired by ideas of Lindahl (1919)), proved its existence, and showed that it always lies in the core. This result was introduced to the computer science literature by Fain et al. (2016).

In the setting where each public good has no cap on funding (we call this the *uncapped setting*), Fain et al. (2016)

showed that Lindahl equilibrium is equivalent to the rule that maximizes Nash social welfare (i.e., the product of agent utilities). The Nash rule has its root in the Nash (1950) bargaining solution, and its objective function has attractive mathematical properties such as scale-freeness (Moulin, 2004). The Nash rule as applied to the public goods model had already been discussed earlier and independently from Lindahl equilibrium due to its attractive group fairness properties (Bogomolnaia et al., 2005; Guerdjikova and Nehring, 2014). The connection between Lindahl equilibrium and the Nash rule is convenient since the latter can be efficiently computed via a convex program reminiscent of the classic Eisenberg–Gale program (Eisenberg and Gale, 1959; Eisenberg, 1961) for computing a market equilibrium for private goods. In addition, Brandl et al. (2022) showed that the Nash rule can be computed by running a simple proportional response dynamics which converges to the Nash outcome. They pointed out that the same convex program had been considered in several unrelated contexts such as in the portfolio selection literature, where this dynamics had also been discovered and shown to converge (Cover, 1984). While the dynamics converges rapidly in practice, a formal bound on the speed of convergence had not been established by 2022, with Li et al. (2018, page 11) noting that the "algorithm [of Cover] possesses a guarantee of convergence but [no] convergence rate."

In the *capped setting*, the Nash rule loses its fairness properties and is not equivalent to Lindahl equilibrium. In contrast, Lindahl equilibrium retains its fairness properties, and its existence is known via fixed-point theorems (Foley, 1970). However, this existence result only applies to strictly monotonic utility functions and thus does not allow agents to have valuations equal to 0 for some goods, and it does not allow for caps except through approximating them through appropriate 'saturating' utility functions (Fain et al., 2016; Munagala et al., 2022b). Most importantly, the existence result is not algorithmic, and how to compute a Lindahl equilibrium was an open question. Fain et al. (2016) asked: "Is computing the Lindahl equilibrium for public goods computationally hard or is there a polynomial time algorithm even [when the public goods are capped]?" Since then there has been no progress on this question. Indeed, Jiang et al. (2020) again noted that "we do not know how to compute the Lindahl equilibrium efficiently". It was even open how to compute any allocation that lies in the core, not necessarily a Lindahl equilibrium allocation.

1.1 Contributions

In the uncapped setting, we prove that the proportional response dynamics converges to a Lindahl equilibrium at a rate of $\log(nm)/t$. We show this by developing a new convex program, distinct from the standard Eisenberg–Gale-style program for Nash welfare maximization, and show that applying mirror descent to this program is equivalent to the proportional response dynamics, thereby allowing us to obtain the convergence rate from known results about mirror descent.[1]

Our new convex program is related to the Eisenberg–Gale-style program through double duality: we show that it can be obtained by taking the dual, introducing new redundant variables, making a change of variable, and performing another dual derivation on this reformulated dual.

The duality and mirror descent relationship that we discover for public goods mirrors existing relationships known in the literature on *private goods* allocation using Fisher market equilibrium. For the private-good setting, equilibrium is also equivalent to maximizing Nash welfare. An alternative convex program for this equilibrium was developed by Shmyrev (1983, 2009). A proportional response dynamics exists for the private goods case as well (Wu and Zhang, 2007; Zhang, 2011), and Birnbaum et al. (2011) showed that it is equivalent to mirror descent on the Shmyrev program. Our new program for the uncapped public goods setting is "Shmyrev-like".

In the capped setting, we answer the open problem raised by Fain et al. (2016) positively: Lindahl equilibrium can be computed efficiently in the capped setting. Indeed, the caps can be naturally added as constraints to our new convex program, and the resulting program correctly computes a Lindahl equilibrium respecting the caps. (In contrast, as is well-known, adding caps as constraints to the Nash welfare program does not lead to a Lindahl equilibrium.) We present numerical experiments on real-life data from participatory budgeting, showing that solving our program is feasible even for large instances.

1.2 Related Work

Lindahl equilibrium Lindahl equilibrium was introduced by Foley (1970), who named this equilibrium concept after Lindahl (1919) who put forward related ideas of personalized taxation. However, note that there are other distinct ways of formalizing Lindahl's ideas (see van den Nouweland, 2015), including ratio and cost share equilibrium (Kaneko, 1977; Mas-Colell and Silvestre, 1989). We use the Foley definition.

Uncapped setting Our interest in Lindahl equilibrium is motivated mainly by their proportional fairness properties (notably the core). Such fairness properties have been studied in many related models, notably the "fair mixing" or "portioning" models (Bogomolnaia et al., 2005; Fain et al., 2016; Aziz et al., 2020; Brandl et al., 2021; Airiau et al., 2023; Gul and Pesendorfer, 2020) that correspond to what we call the uncapped setting. In this setting, Lindahl equilibrium coincides with the maximum Nash welfare solution which has been axiomatically characterized (Guerdjikova and Nehring, 2014) and noted for its strong participation incentives (Brandl et al., 2022) as well as its lowest-possible price of fairness (Michorzewski et al., 2020). The Nash solution is also well-known to provide fair outcomes in other models, such as for private goods (Caragiannis et al., 2019).

Capped setting The capped setting has also been studied in various special cases under various names, such as cake sharing (Bei et al., 2024), fractional committee elections (Pierczyński and Skowron, 2022; Suzuki and Vollen, 2024), or

[1] Zhao (2023) has recently obtained the same convergence rate bound of $\log(nm)/t$, using a direct first-order analysis of the multiplicative gradient (MG) method. Zhao notes that "the extraordinary numerical performance of the MG method is rather surprising and somewhat mysterious [because it] is extremely simple". Our results demystify the performance of the dynamics, by showing that it is equivalent to mirror descent, but on a different convex program.

divisible participatory budgeting (Fain et al., 2016; Aziz and Shah, 2021). These works have mostly not considered Lindahl equilibrium, since there was no known way of computing one.

Discrete models In discrete models, the public goods can either be fully funded or not at all. This model captures the way many cities run their participatory budgets, and has thus been well-studied including via core-like fairness notions such as EJR (Rey and Maly, 2023; Peters et al., 2021a), that were developed in the large literature on approval-based committee elections (Lackner and Skowron, 2023; Aziz et al., 2017; Peters, 2025). There also exist proposals for definitions of Lindahl equilibrium for discrete models (Peters et al., 2021b; Munagala et al., 2022a).

Computation In the uncapped setting, the maximum Nash welfare solution (and thus Lindahl equilibrium) can be efficiently computed via an Eisenberg–Gale-style convex program. This program has a simple structure (maximizing a natural objective function over the standard simplex), and Zhao (2023) has cataloged its appearance in many unrelated areas, including portfolio selection for maximizing log investment returns (Cover, 1984), information theory (Csiszár, 1974) and statistics (Vardi and Lee, 1993), and in medical imaging for positron emission tomography (Vardi et al., 1985). Cover (1984) proposed a dynamics converging to the optimal solution of this program. Convergence proofs were also given by Csiszár (1984) and Brandl et al. (2022). Later, Zhao (2023) obtained a convergence rate of $\log(nm)/t$ for this dynamics. This is the same rate that we establish, though his approach does not connect the dynamics to mirror descent. In the capped setting, very little was known about computation, except for a heuristic algorithm proposed by Fain et al. (2016).

2 Setup

Most proofs are omitted due to space constraints. They can be found in the arXiv version at https://arxiv.org/abs/2503.16414.

Let M be a set of m *public goods*, which we sometimes refer to as *projects*. We have an overall *budget* $B > 0$ that we can spend on the public goods. Let $N = \{1, \dots, n\}$ be a set of n *agents*. Each agent $i \in N$ has an individual budget $B_i > 0$ representing i's weight or endowment. These sum to the overall budget, $\sum_{i \in N} B_i = B$. In many applications, the entitlements are equal: $B_i = B/n$. Each agent i has a valuation $v_{ij} \geqslant 0$ for each public good $j \in M$. We write $v_i = (v_{ij})_{j \in M}$ for the vector of i's valuations. The *utility* of an agent $i \in N$ for an outcome $x \in \mathbb{R}^m_{\geqslant 0}$ is $u_i(x) = \langle v_i, x \rangle = \sum_{j \in M} v_{ij} x_j$. Thus, we use separable linear utilities. We write $M_i = \{j \in M : v_{ij} > 0\}$ for the projects that agent $i \in N$ likes, and we write $N_j = \{i \in N : v_{ij} > 0\}$ for the agents that support project $j \in M$.

In the *uncapped public goods* setting, an *allocation* is a vector $x = (x_j)_{j \in M}$ with $x_j \geqslant 0$ for all $j \in M$ and $\sum_{j \in M} x_j \leqslant B$. Here, x_j denotes the total spending on project b. In this definition, the public goods have no upper bound on how much of them we can spend on them, so in principle the entire budget B could be spent on a single good.

In the *capped public goods* setting, we add the additional constraint that each good $j \in M$ has a maximum amount $\text{cap}_j > 0$ that can be spent on it. Thus, in this setting, an *allocation* is a vector $x = (x_j)_{j \in M}$ with $0 \leqslant x_j \leqslant \text{cap}_j$ for all $j \in M$ and $\sum_{j \in M} x_j \leqslant B$. We assume that $\sum_{j \in M} \text{cap}_j \geqslant B$ (if not then we simply fully fund all the goods).

2.1 Lindahl Equilibrium

Our goal is to find a *Lindahl equilibrium* which is known to yield a fair and efficient allocation of public goods, in the sense that it yields an allocation that is Pareto efficient and lies in the (weak) core (Foley, 1970; Fain et al., 2016). Let $p = (p_{ij})_{i \in N, j \in M}$ be a collection of non-negative *personalized prices*, with $p_{ij} \geqslant 0$ denoting the price that agent i needs to pay per unit of project j, and $p_i = (p_{ij})_{j \in M}$ denoting the vector of prices facing i.

Definition 2.1 (Lindahl Equilibrium)**.** Let x be an allocation and let p be a collection of non-negative personalized prices. Then (x, p) is a *Lindahl equilibrium* if

- x is *affordable*: we have $\langle p_i, x \rangle \leqslant B_i$ for every $i \in N$,
- x is *utility-maximizing*: for every $i \in N$ and every $y \in \mathbb{R}^m_{\geqslant 0}$ such that $0 \leqslant y_j \leqslant \text{cap}_j$ for all $j \in M$ and such that $\langle p_i, y \rangle \leqslant B_i$, we have $u_i(x) \geqslant u_i(y)$,
- x is *profit-maximizing*: for every $j \in M$, we have $\sum_{i \in N} p_{ij} \leqslant 1$, and if $x_j > 0$ then $\sum_{i \in N} p_{ij} = 1$.

An allocation x is a *Lindahl equilibrium allocation* if there exist prices p such that (x, p) is a Lindahl equilibrium.

The distinctive property of a Lindahl equilibrium is that prices are personalized, but every agent demands the exact same bundle x of public goods. The interpretation of the profit maximization condition is less clear. Its most important effect is that it imposes some amount of efficiency: an equilibrium can only spend a positive amount of budget on projects that have the maximum total price.

Example 2.2 (Personal projects)**.** *Consider the uncapped setting, and suppose that each agent likes exactly one project that nobody else likes, so we have $N = M$, with $v_{ii} = 1$ for each $i \in N$ and $v_{ij} = 0$ for all $i \neq j$. In a Lindahl equilibrium, for each $i \in N$, utility maximization requires $x_i > 0$ and that the entire endowment B_i is spent on the personal project, so $p_{ii} = B_i/x_i$ and $p_{ij} = 0$ when $i \neq j$. By profit maximization, since $x_i > 0$, we get that $p_{ii} = 1$. Thus, $B_i/x_i = 1$ and so $x_i = B_i$. Therefore, there is a unique Lindahl equilibrium allocation x with $x_i = B_i$ for each $i \in N$.*

Every Lindahl equilibrium (x, p) can be *decomposed*: For each $i \in N$ and $j \in M$, write

$$b_{ij} = p_{ij} x_j$$

for the *contribution* of i towards j. This is a decomposition of x (similar to a notion considered by Brandl et al. (2022, Definition 2)) because the values $(b_{ij})_{ij}$ satisfy:

- For each $j \in M$, we have $x_j = \sum_{i \in N} b_{ij}$.
- For each $i \in N$, we have $\sum_{j \in M} b_{ij} \leqslant B_i$.

With this interpretation, we can see that p_{ij} equals the fraction of spending on project j that is contributed by agent i.

Foley (1970) proved the existence of Lindahl equilibrium using a fixed-point theorem, in a model that is more general than ours. However, his result only applies to strictly monotonic preferences, and thus only establishes existence when $v_{ij} > 0$ for all $i \in N$ and $j \in M$. We will allow $v_{ij} = 0$. In the presence of zeros, it makes sense to consider Lindahl equilibria (x, p) that are what we call zero-respecting.

Definition 2.3 (Zero-respecting)**.** A Lindahl equilibrium (x, p) is *zero-respecting* if for all $i \in N$ and $j \in M$, whenever $v_{ij} = 0$ and $x_j > 0$ then $p_{ij} = 0$.

This is a natural condition in view of the decomposition we considered above, because in a zero-respecting Lindahl equilibrium, an agent contributes only to projects with positive utility: if $v_{ij} = 0$ then $b_{ij} = 0$.

The following example shows that not every Lindahl equilibrium is zero-respecting, and that zero-respecting Lindahl equilibria may violate Pareto efficiency.

Example 2.4 (Lindahl equilibrium may underspend)**.** *Consider the following instance:*

	B_i	Project 1	Project 2
Agent 1	0.5	1	0
Agent 2	0.5	0	1
cap_j		0.25	∞

On this instance, the unique zero-respecting Lindahl equilibrium allocation is $x = (0.25, 0.5)$. To see this, note that each agent will demand the project that the agent likes, no matter the prices. Thus $x_1, x_2 > 0$. By profit maximization and the zero-respecting condition, we have $p_{11} = p_{22} = 1$ and $p_{12} = p_{21} = 0$. Then by the affordability and utility maximization conditions of Lindahl equilibrium, we get $x = (0.25, 0.5)$. Note that the total spending in this instance is 0.75, strictly less than the available budget of $B = 0.5 + 0.5$. In particular, x is Pareto-dominated by the allocation $y = (0.25, 0.75)$.

If we remove the zero-respecting condition, there exist other Lindahl equilibria. In particular, $x' = (0.25, 0.75)$ forms an equilibrium with the prices $p_1 = (1, \frac{1}{3})$ and $p_2 = (0, \frac{2}{3})$.

Note that the formal model of Foley (1970) does not directly support caps, but these can be simulated ε-approximately through concave utility functions (Munagala et al., 2022b, Footnote 2). We will handle caps without approximations.

2.2 Pareto-Optimality and the Core

Next we discuss how the Lindahl equilibrium relates to Pareto optimality and the set of allocations that are in the core. In the uncapped setting with strictly increasing valuations, the relationship between these concepts is straightforward, and was already studied by Foley (1970). However, as we shall see, there is more nuance in the capped setting and in the presence of valuations equal to 0. We begin by introducing a sufficient condition that excludes examples like Theorem 2.4 where intuitively the caps of projects that receive non-zero valuations are too low. We will see that under this sufficient condition, every zero-respecting Lindahl equilibrium spends the entire budget, is Pareto efficient, and lies in the core.

For every $i \in N$, write $F_i = \{f \in N \mid M_i \cap M_f \neq \emptyset\}$ for the set of "friends" of i who agree that at least one common project has a positive valuation.

Definition 2.5. An instance is *cap-sufficient* if we have $\sum_{j \in M_i} \text{cap}_j \geqslant \sum_{f \in F_i} B_f$ for all $i \in N$.

There are many interesting settings in which instances are always cap-sufficient, including:

- The uncapped setting where $\text{cap}_j = +\infty$ for all $j \in M$.
- All valuations are positive: $v_{ij} > 0$ for all i and j.
- Each agent has positive utility for goods whose total cap reaches the budget: $\sum_{j \in M_i} \text{cap}_j \geqslant B$.

We will show that Lindahl equilibrium has particularly desirable properties on cap-sufficient instances. A key consequence of cap-sufficiency is that every voter spends their entire budget.

Proposition 2.6. *On a cap-sufficient instance, if (x, p) is a zero-respecting Lindahl equilibrium, then for every $i \in N$, we have $\langle p_i, x \rangle = B_i$. It follows that $\sum_{j \in M} x_j = B$.*

A major reason to be interested in Lindahl equilibrium is that it always lies in the weak core, which is a fairness or stability property formalizing proportional representation.

Definition 2.7 (Core)**.** An allocation x is in the *core* if there is no blocking coalition $S \subseteq N$ and no objection $z = (z_j)_{j \in M} \in \mathbb{R}^m_{\geqslant 0}$ with $0 \leqslant z_j \leqslant \text{cap}_j$ for all $j \in M$, such that $\sum_{j \in M} z_j \leqslant \sum_{i \in S} B_i$ (it can be afforded by the blocking coalition) and for all $i \in S$, we have $\langle v_i, z \rangle \geqslant \langle v_i, x \rangle$ and the inequality is strict for at least one $i \in S$. It is in the *weak core* if there are no such S and z such that $\langle v_i, z \rangle > \langle v_i, x \rangle$ for all $i \in S$.

Foley (1970, Section 6) proved that Lindahl equilibrium allocations are in the weak core, though his model implicitly assumed cap-sufficiency. We can show the following.

Proposition 2.8. *Let (x, p) be a Lindahl equilibrium. Then x lies in the weak core. If the instance is cap-sufficient and (x, p) is zero-respecting, then x lies in the core.*

As a special case, taking $S = N$ in Theorem 2.8, we see that Lindahl equilibrium allocations are (weakly) Pareto efficient, establishing a version of the First Welfare Theorem.

Corollary 2.9. *Let (x, p) be a Lindahl equilibrium. Then x is weakly Pareto optimal. If the instance is cap-sufficient and (x, p) is zero-respecting, then x is Pareto optimal.*

3 Convex Optimization Background

This section gives background on convex optimization and the mirror descent algorithm.

Basic definitions. Let $f : \mathbb{R}^n \to (-\infty, \infty]$ be a function. We use the convention $0 \log 0 = 0$. We write $B \cdot \Delta^m = \{x \in \mathbb{R}^m_{\geqslant 0} : \sum_j x_j = B\}$ for the scaled simplex.

Mirror descent. The mirror descent (MD) algorithm is a first-order method for convex minimization which generalizes projected gradient descent. The goal is to minimize a convex function f over a convex set X via first-order updates. MD relies on a Bregman divergence $D_h(x \| y)$, which is a convex function that measures the difference between x and y. The function D_h is constructed from some 1-strongly convex *reference function* h as $D_h(x \| y) = h(x) - h(y) - \langle \nabla h(y), x - y \rangle$. For example, taking the negative entropy reference function $h(x) = \sum_i x_i \log x_i$, the Bregman divergence becomes the

KL divergence, $D_h(x\|y) = \sum_i x_i \log(x_i/y_i)$. The update rule for MD is

$$x^{t+1} = \arg\min_{x \in X} \langle \nabla f(x^t), x \rangle + \frac{1}{\eta} D_h(x\|x^t), \qquad (1)$$

where $\eta > 0$ is a stepsize parameter. There are a variety of convergence results for MD. We will be interested in the case where a special relationship holds between the objective f and the reference function h, knowns as *relative smoothness*. The function f is said to be 1-smooth relative to the reference function h when it holds for all $x, y \in \operatorname{relint} X$ that

$$f(a) \leqslant f(b) + \langle \nabla f(b), a - b \rangle + D_h(a\|b).$$

The following theorem from Birnbaum et al. (2011) shows that when the reference function h is chosen such that relative smoothness holds, the sequence of iterates generated by mirror descent converges at a rate of $O(1/t)$:

Theorem 3.1 (Birnbaum et al., 2011, Theorem 3)**.** *Suppose that f is 1-smooth relative to the reference function h, and we run mirror descent using h as the distance-generating function. Let x^* be an optimal solution. Then the sequence of iterates generated by mirror descent satisfies:*

$$f(x^t) - f(x^*) \leqslant \frac{D_h(x^*\|x^0)}{t}$$

4 Uncapped Public Goods

We begin by analyzing the uncapped setting, and begin by characterizing the Lindahl equilibrium prices, which will be helpful for understanding the convex programs we discuss. Note that if (x, p) is a Lindahl equilibrium, then each agent will only demand public goods that maximize the "bang-per-buck" ratio v_{ij}/p_{ij}. (When $v_{ij} = 0$ and $x_j > 0$, the agent also demands project j but only if $p_{ij} = 0$. Thus, in the uncapped setting, every Lindahl equilibrium is zero-respecting.) Thus, the quantity v_{ij}/p_{ij} must be equal for all projects j with $x_j > 0$ and $v_{ij} > 0$. Therefore, $p_{ij} \propto v_{ij}$, say with factor of proportionality α. Because each agent spends their entire budget (Theorem 2.6 (i), which applies since in the uncapped setting Lindahl equilibrium is always zero-respecting), we have $B_i = \sum_{j \in M_i} p_{ij} x_j = \sum_{j \in M_i} \alpha v_{ij} x_j = \alpha \langle v_i, x \rangle$. Thus we deduce that in the uncapped setting,

$$p_{ij} = B_i \cdot \frac{v_{ij}}{\langle v_i, x \rangle} \quad \text{for all } j \in M \text{ with } x_j > 0. \qquad (2)$$

Now consider a project $j \in M$ with $x_j = 0$. Because i does not demand it, its bang-per-buck must be weakly below the bang-per-buck of funded projects. By (2),

$$p_{ij} \geqslant B_i \cdot \frac{v_{ij}}{\langle v_i, x \rangle} \quad \text{for all } j \in M \text{ with } x_j = 0. \qquad (3)$$

As we explained in Section 2.1, any Lindahl equilibrium can be decomposed into individual contributions $b_{ij} = p_{ij} x_j$. From (2), it follows that in Lindahl equilibrium,

$$b_{ij} = B_i \cdot \frac{v_{ij} x_j}{\langle v_i, x \rangle} \quad \text{for all } j \in M, \qquad (4)$$

or more simply that $b_{ij} \propto v_{ij} x_j$, so contributions are proportional to the utility i obtains in x from j. One can view (4) as a kind of fixed-point property implied by Lindahl equilibrium (Guerdjikova and Nehring, 2014), and it suggests the proportional response dynamics that we will study later.

4.1 Nash Welfare and the Eisenberg–Gale Program

In the uncapped setting (i.e. $\text{cap}_j = +\infty$ for all $j \in M$), Lindahl equilibrium allocations can be nicely characterized as those maximizing the Nash social welfare $\prod_i u_i(x)$ (Fain et al., 2016). Such an allocation can be computed by solving the following convex program:

$$\begin{aligned} \max_{x \geqslant 0} \; & \sum_{i \in N} B_i \log \langle v_i, x \rangle \\ \text{s.t.} \; & \sum_{j \in M} x_j \leqslant B \end{aligned} \qquad (5)$$

This program is the public-goods analogue of the Eisenberg–Gale convex program for computing a Fisher market equilibrium with private goods (Eisenberg and Gale, 1959; Eisenberg, 1961). Based on this description of the prices, we can now one can analyze the KKT conditions of Program 5 to show that it exactly computes Lindahl equilibrium.

Theorem 4.1. [Fain et al., 2016, Corollary 2.3] *In the uncapped setting, an allocation x is a Lindahl equilibrium allocation if and only if it is an optimal solution to Program 5.*

Fain et al. (2016, Theorem 2.2) also present Eisenberg–Gale-style programs for computing Lindahl equilibria for certain non-linear utility functions called "scalar separable non-satiating" including CES and Cobb-Douglas utilities.

Interestingly, for Fisher market equilibrium, the Eisenberg–Gale program always admits a rational solution (Devanur et al., 2008; Vazirani, 2012). However, this is not the case in the public goods setting (Airiau et al., 2023, Theorem 5).

4.2 A New Convex Program

We will present a new convex program which also captures the Lindahl equilibrium concept in the uncapped setting. As we will see, this convex program will yield several useful results. First, we will use it to show that the proportional response dynamics for uncapped public goods can indeed be interpreted as mirror descent with the entropy distance, just as in the Fisher market setting. Secondly, extending this convex program will allow us to give the first computational results for the capped public goods setting. Our new convex program is in the spirit of the Shmyrev convex program for Fisher markets for private goods (Shmyrev, 1983, 2009), though there are important differences. The convex program is as follows:

$$\begin{aligned} \max_{b \geqslant 0, x \geqslant 0} \; & \sum_{i \in N, j \in M_i} b_{ij} \log v_{ij} - \sum_{i \in N, j \in M_i} b_{ij} \log (b_{ij}/x_j) \\ \text{s.t.} \; & \sum_{j \in M_i} b_{ij} = B_i, \qquad \forall i \in N \\ & \sum_{i \in N_j} b_{ij} = x_j, \qquad \forall j \in M \end{aligned} \qquad (6)$$

The program has two sets of variables, though one is implied by the other. The x_j variable has the same interpretation as in Program 5: it is the amount of budget allocated to project j. The b_{ij} variables can be interpreted as the share of agent i's budget B_i that they allocate towards project j. Note that each x_j variable is directly implied by the choice of the b_{ij}

variables across agents i. We could replace each occurrence of it in Program 6 by $\sum_{i \in N_j} b_{ij}$.

While our program has some similarity to the Shmyrev program for private goods (Shmyrev, 2009; Birnbaum et al., 2011), it has the following important differences. First, the Shmyrev program contains variables corresponding to prices, which do not appear in our program. Second, the original primal variables x_j appear directly in our program, whereas in Shmyrev's program these are a non-linear function of the corresponding b variables. Third, we have a somewhat unusual term that looks like a partially-normalized entropy in our objective, whereas Shmyrev's program only requires using a typical negative entropy term over prices.

4.3 Connecting the Programs via Duality

Program 5 and Program 6 can be related to each other through "double duality". This is analogous to a result for private goods, where the Shmyrev and the Eisenberg–Gale program also share a dual after reformulation (Cole et al., 2017).

Theorem 4.2. *Program 6 is the dual of the dual of the Eisenberg–Gale convex program for public goods, after reformulation.*

4.4 Proportional Response Dynamics

It is known that the Lindahl equilibrium for the uncapped public goods setting can be computed by a simple dynamics (Brandl et al., 2022) which we call the *proportional response* dynamics in analogy to a similar dynamics for private-good Fisher markets (Wu and Zhang, 2007; Zhang, 2011). At each iteration t, the proportional response dynamics have some current budget allocation $x^t = (x_1^t, \dots, x_j^t)$ summing to B. Let $u_i^t = \langle v_i, x^t \rangle$ be the current utility of agent i under this allocation. Then the next budget allocation in the dynamics is

$$x_j^{t+1} = \sum_{i \in N} B_i \frac{x_j^t v_{ij}}{u_i^t}.$$

This dynamics can be interpreted as each agent i independently deciding how they wish to allocate their share of the budget B_i in the next round. Specifically, agent i allocates spending proportional to how much utility each project provided them at round t. This spending allocation matches the property in (4) we derived earlier from the definition of Lindahl equilibrium. We will show that the proportional response dynamics is the mirror descent algorithm applied to our Program 6.

In order to derive this relationship, we first reformulate Program 6 to an equivalent version: we eliminate the redundant x_j variables, convert the problem into a minimization problem, and define the shorthand function $x_j(b) = \sum_{i \in N} b_{ij}$. Then we get the following convex program:

$$\begin{aligned} \min_{b \geqslant 0} \quad & f(b) := -\sum_{i \in N, j \in M_i} b_{ij} \left(\log v_{ij} - \log \left(b_{ij} / x_j(b) \right) \right) \\ \text{s.t.} \quad & \sum_{j \in M} b_{ij} = B_i, \ \forall i \in N \end{aligned} \tag{7}$$

Theorem 4.3. *Assume that the PR dynamics and mirror descent algorithm on Program 7 are both initialized at a point $b^0 \in \mathbb{R}^{n \times m}$ such that $b_i^0 \in B_i \cdot \Delta^m$ and $x_j(b^0) = \sum_{i \in N} b_{ij}^0 > 0$ for all $j \in M$. Then the proportional response dynamics are equivalent to applying the mirror descent algorithm with the entropy reference function to Program 7.*

Thus, the proportional response dynamics is equivalent to mirror descent with unit stepsize. Next we wish to apply the convergence-rate result from Theorem 3.1. Thus, we need to show that the objective in Program 7 is 1-smooth relative to the entropy reference function.

Lemma 4.4. *The function f is 1-smooth relative to the reference function $h(b) = \sum_{i \in N, j \in M} b_{ij} \log b_{ij}$, i.e., for all $a, b \in \mathbb{R}_{>0}^{n \times m}$ such that $a_i \in B_i \cdot \Delta^m, b_i \in B_i \cdot \Delta^m$ we have*

$$f(a) \leqslant f(b) + \langle \nabla f(b), a - b \rangle + D_h(a \| b)$$

Now we can combine Theorem 4.4 with Theorem 3.1 to get a $D_h(b^* \| b^0)/t$ rate of convergence for the proportional response dynamics. If we start the dynamics at the uniform allocation $b_{ij}^0 = B_i/m$, we can upper bound the Bregman divergence $D_h(b^* \| b^0)$ as follows:

$$\begin{aligned} D_h(b^* \| b^0) &= h(b^*) - h(b^0) \leqslant -h(b^0) \\ &= -\sum_{i \in N, j \in M_i} \frac{B_i}{m} \log(B_i/m) = \sum_{i \in N} B_i \log(m/B_i). \end{aligned}$$

Combining this with Theorem 3.1 and Theorem 4.4, we get a $\frac{\sum_{i \in N} B_i \log(m/B_i)}{t}$ rate of convergence for the proportional response dynamics. Suppose for simplicity that $B = 1$ and $B_i = 1/n$, then we get that proportional response dynamics converges at a rate of $\frac{\log(nm)}{t}$.

The same convergence rate was recently independently obtained by Zhao (2023), after it had been an open question for almost fifty years. Zhao (2023) derived this rate directly. Our result gives a deeper explanation of the performance of the proportional response dynamics: it is equivalent to mirror descent with the entropy reference function applied to Program 6.

5 Capped Public Goods

Next we study the capped public goods setting, where we have a constraint $x_j \leqslant \text{cap}_j$ for each good $j \in M$. One may naïvely attempt to add this constraint to Program 5 maximizing Nash welfare, but this does not lead to a Lindahl equilibrium and not even to a core solution (Suzuki and Vollen, 2024, Prop. 4.1), though it does produce an allocation that satisfies the 2-approximate core (Munagala et al., 2022b, Cor. 3.5). Garg et al. (2021, Comment A.1) also noted that Nash does not extend to capped settings, writing that Lindahl equilibrium "does not transform into a Fisher market".

5.1 Adapting the Convex Program

We will show in this section that Program 6 can be used to compute a Lindahl equilibrium in the capped public goods setting through a simple modification: we simply add a constraint $x_j \leqslant \text{cap}_j$ for all $j \in M$. Surprisingly, we will show that this works, even though the exact same constraint does not work for the original EG program (Program 5) for maximizing Nash welfare. Thus, we obtain the first efficient algorithm for capped public goods, thereby resolving an open problem first posed by Fain et al. (2016).

Our modified program for capped public goods is as follows, where as before we write $x_j(b) = \sum_{i \in N} b_{ij}$ as a shorthand:

$$\begin{aligned} \max_{b \geqslant 0} \quad & \sum_{i \in N, j \in M_i} b_{ij} \ (\log v_{ij} - \log(b_{ij}/x_j(b))) \\ \text{s.t.} \quad & \sum_{j \in M_i} b_{ij} \leqslant B_i \text{ for all } i \in N \\ & x_j(b) \leqslant \text{cap}_j \text{ for all } j \in M \end{aligned} \tag{8}$$

5.2 The Program Computes a Lindahl Equilibrium

We can prove that Program 8 computes a zero-respecting Lindahl equilibrium. This in particular proves the existence of such an equilibrium, which does not quite follow from the existence result of Foley (1970), since his model does not allow for caps and does not allow for valuations equal to 0 (since it assumes strictly monotonic valuations).

Our proof proceeds by analyzing KKT condition applied to Program 8. We will require that all valuations have been rescaled such that $v_{ij} > 1$ for all $v_{ij} \neq 0$, to ensure that the coefficients $\log v_{ij}$ in the objective function are positive. Rescaling is without loss of generality, since the Lindahl equilibrium is invariant to scaling valuations by a positive constant. A similar normalization is used by Brandl et al. (2022).

Theorem 5.1. *Assume that valuations are rescaled such that $v_{ij} > 1$ for all $v_{ij} \neq 0$. Let x^* be an optimal solution to Program 8. Then there exist zero-respecting prices p such that (x^*, p) forms a Lindahl equilibrium for the capped public goods setting.*

5.3 Discussion of the Convex Program

Comparison to Fisher markets. It is interesting to contrast our program with the Fisher market setting with private goods. There, the Eisenberg–Gale program also does not allow the introduction of saturating constraints on the primal variables (which correspond to a maximum amount of a good that an agent may receive). Yet it is not possible to add such constraints to the Shmyrev program for Fisher markets either, because that program does not contain the original primal variables encoding the allocation (in contrast to our public-goods program). Instead, the allocation is obtained through a nonlinear function of the optimization variables in the Shmyrev program. Thus, Program 6 allows for a type of saturating consumption constraint that has previously never been possible for either private or public goods.

Not all Lindahl equilibria are optimal solutions. In the uncapped setting, *every* Lindahl equilibrium forms an optimum of both Program 6 and the Eisenberg–Gale program. As the following example shows, this is not the case for the capped setting, where Program 8 captures only a strict subset of Lindahl equilibria. The example also shows that Lindahl equilibria are not unique in utilities.

Example 5.2 (Lindahl equilibrium is not unique in utilities)**.** *Consider the following instance:*

	B_i	Project 1	Project 2	Project 3
Agent 1	1	1	1	0
Agent 2	1	1	0	1
cap_j		1	∞	∞

This instance is cap-sufficient, as each agent has a positive valuation for an uncapped project. Let us determine the set of zero-respecting Lindahl equilibria (x, p). By Theorem 2.9, x is Pareto-optimal, and therefore $x_1 = 1$ and $x_2 + x_3 = 1$. For each $\gamma \in [0, 1]$, one can check that $x = (1, 1 - \gamma, \gamma)$ forms a Lindahl equilibrium together with the prices $p_1 = (\gamma, 1, 0)$ and $p_2 = (1 - \gamma, 0, 1)$. Thus, in the capped setting, Lindahl equilibria are not unique in utilities: in the equilibrium allocation $(1, 1, 0)$, agent 1 obtains utility 2, but in $(1, 0, 1)$, agent 1 obtains utility 1.

Note that the allocation $x^* = (1, \frac{1}{2}, \frac{1}{2})$ is the unique allocation that is intuitively fair and respects the symmetry of the instance, but this allocation is not the only Lindahl equilibrium. However, Program 8 uniquely selects x^*, since on this instance its objective function simplifies to $-b_{11} \log b_{11} - b_{21} \log b_{21}$, maximized by $b_{11} = b_{21} = 0.5$.

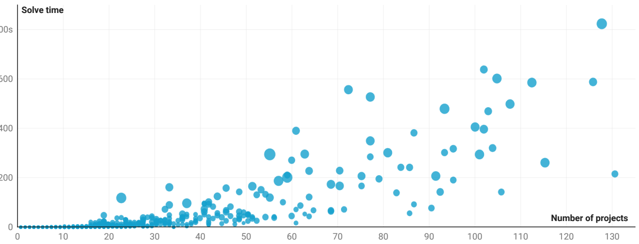

Figure 1: Results of our experiments on Pabulib instances, showing the solve time of the MOSEK solver as a function of the number of projects in the instance. The largest instances are from Warsaw and Amsterdam.

5.4 Computation and Experiments

Let us briefly discuss how to solve Program 8. Numerically, the program can be solved using any conic convex optimization solver supporting exponential cones, such as MOSEK, COPT, Clarabel, ECOS, or SCS. We built a simple online tool for solving moderate-size instances with the SCS solver (O'Donoghue et al., 2016), available at dominik-peters.de/demos/lindahl/. From a complexity-theoretic perspective, an ε-optimal solution to Program 8 can be computed in polynomial time using the ellipsoid method (see, e.g., Vishnoi, 2021, Theorem 13.1).

To evaluate the performance of computing Lindahl equilibrium via Program 8, we implemented it using the MOSEK solver and applied it to the participatory budgeting datasets in the Pabulib repository (Faliszewski et al., 2023). We find that the program can be solved quite quickly, with solve times shown in Section 5.4. The longest solve time we encountered was 822s (or 1489s including the time to write down the encoding) for an instance from Warsaw with 14 897 voters (with 11 426 distinct approval sets) and 134 projects.

6 Conclusion

We have developed a new class of convex programs that can be used to efficiently compute Lindahl equilibria both in the uncapped and the capped setting. These new programs open up many opportunities for future research.

In the uncapped setting, our new program might lead to new proofs of known results for the well-studied maximum Nash welfare rule. This might include the result about participation incentives of Brandl et al. (2022) or the axiomatic characterization of Guerdjikova and Nehring (2014). Perhaps our program could also shed light on the other uses of the Eisenberg–Gale program across statistics, information theory, and medical imaging, as discussed in Section 1.2. For the capped setting, our computability result has implications for the discrete public goods model, because it allows the efficient implementation of the 9.27-approximation to the core obtained by Munagala et al. (2022b), rather than having to rely on their 67.37-approximation. Munagala et al. (2022b) used Lindahl equilibrium as a black box to obtain their approximation result; reasoning about the structure of our convex program might lead to even better bounds.

Since our focus has been on computational questions, we have not considered strategic aspects. Lindahl equilibrium is well-known to have high informational requirements, and in particular we need to know the truthful valuations of the agents to compute it. Interpreted as a decision rule (Gul and Pesendorfer, 2020), Lindahl equilibrium is not strategyproof and can be manipulated both in a free-riding sense (Brandl et al., 2021, Section 5.3), and in some paradoxical ways (Aziz et al., 2020, Theorem 3(ii)), even in the uncapped setting. Manipulability is unavoidable if one desires a Pareto-efficient and core-stable solution, both in the uncapped setting (Brandl et al., 2021, Theorem 2 and Theorem 3) and in the capped setting (Bei et al., 2024, Theorem 6.2). These impossibilities apply even for approval (0/1) preferences. For more general linear utilities, strategyproofness is only attainable by dictatorial-type rules (Hylland, 1980), even in the uncapped setting.

We leave several interesting technical questions open. Is the optimum of our program unique in utilities? This is known to be true for the uncapped setting, by strict convexity (in utilities) of the Eisenberg–Gale program. Can we develop first-order methods for the capped settings, or derive a natural dynamics converging to an equilibrium? Applying mirror descent to our program does not appear to lead to a nice closed-form update like in the uncapped setting. Finally, can the cap constraint be generalized? For example, one could apply cap constraints on the total spending of *sets* of public goods. This would allow us to model multi-issue and multi-round decision making settings (see, e.g., Banerjee et al., 2023, Section 5). It would also allow us to embed private goods in the model (as in Conitzer et al., 2017), and potentially connect the notions of Fisher market equilibrium and Lindahl equilibrium.

References

Stéphane Airiau, Haris Aziz, Ioannis Caragiannis, Justin Kruger, Jérôme Lang, and Dominik Peters. 2023. Portioning using ordinal preferences: Fairness and efficiency. *Artificial Intelligence* 314 (2023), 103809. doi:10.1016/j.artint.2022.103809 [→ p. 1, 2, 5]

Haris Aziz, Anna Bogomolnaia, and Hervé Moulin. 2020. Fair mixing: The case of dichotomous preferences. *ACM Transactions on Economics and Computation (TEAC)* 8, 4, Article 18 (2020), 27 pages. doi:10.1145/3417738 [→ p. 1, 2, 8]

Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. 2017. Justified Representation in approval-based committee voting. *Social Choice and Welfare* 48, 2 (2017), 461–485. doi:10.1007/s00355-016-1019-3 [→ p. 3]

Haris Aziz, Xinhang Lu, Mashbat Suzuki, Jeremy Vollen, and Toby Walsh. 2023a. Best-of-both-worlds fairness in committee voting. arXiv:2303.03642 [cs.GT] https://arxiv.org/abs/2303.03642 [→ p. 1]

Haris Aziz, Evi Micha, and Nisarg Shah. 2023b. Group Fairness in Peer Review. In *Advances in Neural Information Processing Systems*, Vol. 36. 64885–64895. https://proceedings.neurips.cc/paper_files/paper/2023/file/ccba10dd4e80e7276054222bb95d467c-Paper-Conference.pdf [→ p.]

Haris Aziz and Nisarg Shah. 2021. *Participatory budgeting: Models and approaches*. Springer, 215–236. doi:10.1007/978-3-030-54936-7_10 [→ p. 1, 3]

Siddhartha Banerjee, Vasilis Gkatzelis, Safwan Hossain, Billy Jin, Evi Micha, and Nisarg Shah. 2023. Proportionally fair online allocation of public goods with predictions. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*. 20–28. doi:10.24963/ijcai.2023/3 [→ p. 8]

Xiaohui Bei, Xinhang Lu, and Warut Suksompong. 2024. Truthful cake sharing. *Social Choice and Welfare* (2024), 1–35. doi:10.1007/s00355-023-01503-0 [→ p. 1, 2, 8]

Benjamin Birnbaum, Nikhil R. Devanur, and Lin Xiao. 2011. Distributed algorithms via gradient descent for Fisher markets. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*. 127–136. doi:10.1145/1993574.1993594 [→ p. 2, 5, 6]

Anna Bogomolnaia, Hervé Moulin, and Richard Stong. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory* 122, 2 (2005), 165–184. doi:10.1016/j.jet.2004.05.005 [→ p. 2]

Florian Brandl, Felix Brandt, Matthias Greger, Dominik Peters, Christian Stricker, and Warut Suksompong. 2022. Funding public projects: A case for the Nash product rule. *Journal of Mathematical Economics* 99 (2022), 102585. doi:10.1016/j.jmateco.2021.102585 [→ p. 1, 2, 3, 6, 7, 8]

Florian Brandl, Felix Brandt, Dominik Peters, and Christian Stricker. 2021. Distribution rules under dichotomous preferences: two out of three ain't bad. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*. 158–179. doi:10.1145/3465456.3467653 [→ p. 1, 2, 8]

Felix Brandt, Matthias Greger, Erel Segal-Halevi, and Warut Suksompong. 2024. Coordinating charitable donations. arXiv:2305.10286 [econ.TH] [→ p. 1]

Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)* 7, 3 (2019), 1–32. doi:10.1145/3355902 [→ p. 2]

Ioannis Caragiannis, Evi Micha, and Nisarg Shah. 2024. Proportional Fairness in Non-Centroid Clustering. In *Proceedings of the 38th Annual Conference on Neural Information Processing Systems (NeurIPS)*. https://openreview.net/forum?id=Actjv6Wect [→ p.]

Bhaskar Ray Chaudhury, Linyi Li, Mintong Kang, Bo Li, and Ruta Mehta. 2022. Fairness in Federated Learning via Core-Stability. In *Advances in Neural Information Processing Systems*, Vol. 35. 5738–5750. PDF at proceedings.neurips.cc. [→ p.]

Xingyu Chen, Brandon Fain, Liang Lyu, and Kamesh Munagala. 2019. Proportionally fair clustering. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*. PMLR, 1032–1041. https://proceedings.mlr.press/v97/chen19d.html [→ p.]

Xi Chen and Shang-Hua Teng. 2009. Spending is not easier than trading: on the computational equivalence of Fisher and Arrow-Debreu equilibria. In *Proceedings of the 20th International Symposium on Algorithms and Computation (ISAAC)*. 647–656. [→ p.]

Yu Cheng, Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2020. Group fairness in committee selection. *ACM Transactions on Economics and Computation (TEAC)* 8, 4, Article 23 (2020), 18 pages. doi:10.1145/3417750 [→ p.]

Richard Cole, Nikhil R. Devanur, Vasilis Gkatzelis, Kamal Jain, Tung Mai, Vijay V. Vazirani, and Sadra Yazdanbod. 2017. Convex program duality, Fisher markets, and Nash social welfare. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*. 459–460. doi:10.1145/3033274.3085109 [→ p. 6]

Vincent Conitzer, Rupert Freeman, and Nisarg Shah. 2017. Fair public decision making. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*. 629–646. doi:10.1145/3033274.3085125 [→ p. 8]

Thomas Cover. 1984. An algorithm for maximizing expected log investment return. *IEEE Transactions on Information Theory* 30, 2 (1984), 369–373. doi:10.1109/TIT.1984.1056869 [→ p. 2, 3]

Imre Csiszár. 1974. On the computation of rate-distortion functions (corresp.). *IEEE Transactions on Information Theory* 20, 1 (1974), 122–124. doi:10.1109/TIT.1974.1055146 [→ p. 3]

Imre Csiszár. 1984. Information geometry and alternating minimization procedures. *Statistics and Decisions* Supplemental Issue No. 1 (1984), 205–237. https://dominik-peters.de/archive/csiszar1984.pdf [→ p. 3]

Argyrios Deligkas, John Fearnley, Alexandros Hollender, and Themistoklis Melissourgos. 2024. Constant inapproximability for Fisher markets. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*. 13–39. doi:10.1145/3670865.3673533 [→ p.]

Nikhil R. Devanur, Christos H. Papadimitriou, Amin Saberi, and Vijay V. Vazirani. 2008. Market equilibrium via a primal–dual algorithm for a convex program. *Journal of the ACM (JACM)* 55, 5 (2008), 1–18. doi:10.1145/1411509.1411512 [→ p. 5]

Soroush Ebadian and Evi Micha. 2025. Boosting Sortition via Proportional Representation. In *Proceedings of 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. Forthcoming. [→ p.]

Edmund Eisenberg. 1961. Aggregation of utility functions. *Management Science* 7, 4 (1961), 337–350. doi:10.1287/mnsc.7.4.337 [→ p. 2, 5]

Edmund Eisenberg and David Gale. 1959. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics* 30, 1 (1959), 165–168. doi:10.1214/aoms/1177706369 [→ p. 2, 5]

Brandon Fain, Ashish Goel, and Kamesh Munagala. 2016. The core of the participatory budgeting problem. In *Proceedings of the 12th International Conference on Web and Internet Economics (WINE)*. 384–399. doi:10.1007/978-3-662-54110-4_27 [→ p. 1, 2, 3, 5, 7]

Brandon Fain, Kamesh Munagala, and Nisarg Shah. 2018. Fair allocation of indivisible public goods. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*. 575–592. doi:10.1145/3219166.3219174 [→ p.]

Piotr Faliszewski, Jarosław Flis, Dominik Peters, Grzegorz Pierczyński, Piotr Skowron, Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. 2023. Participatory budgeting: Data, tools, and analysis. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*. 2667–2674. doi:10.24963/ijcai.2023/297 [→ p. 7]

Duncan K. Foley. 1970. Lindahl's solution and the core of an economy with public goods. *Econometrica* (1970), 66–72. doi:10.2307/1909241 [→ p. 1, 2, 3, 4, 7]

Nikhil Garg, Ashish Goel, and Benjamin Plaut. 2021. Markets for public decision-making. *Social Choice and Welfare* 56, 4 (2021), 755–801. doi:10.1007/s00355-020-01298-4 [→ p. 6]

Hilary Greaves and Owen Cotton-Barratt. 2023. A bargaining-theoretic approach to moral uncertainty. *Journal of Moral Philosophy* 1, aop (2023), 1–43. doi:10.1163/17455243-20233810 [→ p.]

Ani Guerdjikova and Klaus Nehring. 2014. Weighing experts, weighing sources: The diversity value. (2014). https://dominik-peters.de/archive/guerdjikova2014.pdf Working Paper. [→ p. 2, 5, 8]

Faruk Gul and Wolfgang Pesendorfer. 2020. Lindahl equilibrium as a collective choice rule. arXiv:2008.09932 [econ.TH] [→ p. 2, 8]

Jean-Baptiste Hiriart-Urruty. 1982. ε-subdifferential calculus. In *Convex Analysis and Optimization*, Jean-Pierre Aubin and Richard B. Vinter (Eds.). Research Notes in Mathematics, Vol. 57. Pitman, 43–92. https://dominik-peters.de/archive/hiriarturruty1982.pdf [→ p.]

Aanund Hylland. 1980. Strategyproofness of voting procedures with lotteries as outcomes and infinite sets of strategies. (1980). https://www.sv.uio.no/econ/personer/vit/aanundh/upubliserte-artikler-og-notater/Strategy%20Proofness%5B1%5D.pdf unpublished. [→ p. 8]

Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2020. Approximately stable committee selection. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. 463–472. doi:10.1145/3357713.3384238 [→ p. 2]

Mamoru Kaneko. 1977. The ratio equilibrium and a voting game in a public goods economy. *Journal of Economic Theory* 16, 2 (1977), 123–136. doi:10.1016/0022-0531(77)90001-1 [→ p. 2]

Leon Kellerhals and Jannik Peters. 2024. Proportional Fairness in Clustering: A Social Choice Perspective. In *Proceedings of the 38th Annual Conference on Neural Information Processing Systems (NeurIPS)*. https://openreview.net/forum?id=KsLX5pFpOs [→ p.]

Martin Lackner and Piotr Skowron. 2023. *Multi-Winner Voting with Approval Preferences*. Springer. doi:10.1007/978-3-031-09016-5 [→ p. 3]

Yen-Huan Li, Carlos A. Riofrio, and Volkan Cevher. 2018. A general convergence result for mirror descent with Armijo line search. arXiv:1805.12232 [math.OC] [→ p. 2]

Erik Lindahl. 1919. Just taxation—a positive solution. Translated from German (*Die Gerechtigkeit der Besteuerung*, Lund 1919, Part I, Chap. 4, pp. 85–98: Positive Lösung) by Elizabeth Henderson. In *Classics in the Theory of Public Finance*, Richard A. Musgrave and Alan T. Peacock (Eds.). Palgrave Macmillan (1958), 168–176. doi:10.1007/978-1-349-23426-4_11 [→ p. 1, 2]

Andreu Mas-Colell and Joaquim Silvestre. 1989. Cost share equilibria: A Lindahlian approach. *Journal of Economic Theory* 47, 2 (1989), 239–256. doi:10.1016/0022-0531(89)90019-7 [→ p. 2]

Marcin Michorzewski, Dominik Peters, and Piotr Skowron. 2020. Price of fairness in budget division and probabilistic social choice. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*. 2184–2191. doi:10.1609/aaai.v34i02.5594 [→ p. 2]

James C. Moore. 2006. *General Equilibrium and Welfare Economics: An Introduction*. Springer. doi:10.1007/978-3-540-32223-8 [→ p.]

Hervé Moulin. 2004. *Fair Division and Collective Welfare*. MIT Press. doi:10.7551/mitpress/2954.001.0001 [→ p. 2]

Kamesh Munagala, Yiheng Shen, and Kangning Wang. 2022a. Auditing for core stability in participatory budgeting. In *Proceedings of the 18th International Conference on Web and Internet Economics (WINE)*. 292–310. doi:10.1007/978-3-031-22832-2_17 [→ p. 3]

Kamesh Munagala, Yiheng Shen, Kangning Wang, and Zhiyi Wang. 2022b. Approximate core for committee selection via multilinear extension and market clearing. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2229–2252. doi:10.1137/1.9781611977073.89 [→ p. 2, 4, 6, 8]

John F. Nash. 1950. The bargaining problem. *Econometrica* 18, 2 (1950), 155–162. doi:10.2307/1907266 [→ p. 2]

Brendan O'Donoghue, Eric Chu, Neal Parikh, and Stephen Boyd. 2016. Conic Optimization via Operator Splitting and Homogeneous Self-Dual Embedding. *Journal of Optimization Theory and Applications* 169, 3 (June 2016), 1042–1068. http://stanford.edu/~boyd/papers/scs.html [→ p. 7]

Dominik Peters. 2019. Economic design for effective altruism. *The Future of Economic Design: The Continuing Development of a Field as Envisioned by Its Researchers* (2019), 381–388. doi:10.1007/978-3-030-18050-8_53 [→ p.]

Dominik Peters. 2025. The Core of Approval-Based Committee Elections with Few Candidates. arXiv:2501.18304 [cs.GT] [→ p. 3]

Dominik Peters, Grzegorz Pierczyński, Nisarg Shah, and Piotr Skowron. 2021b. Market-based explanations of collective decisions. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5656–5663. doi:10.1609/aaai.v35i6.16710 [→ p. 3]

Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. 2021a. Proportional Participatory Budgeting with Additive Utilities. In *Advances in Neural Information Processing Systems*, Vol. 34. 12726–12737. https://proceedings.neurips.cc/paper_files/paper/2021/file/69f8ea31de0c00502b2ae571fbab1f95-Paper.pdf [→ p. 3]

Grzegorz Pierczyński and Piotr Skowron. 2022. Core-stable committees under restricted domains. In *Proceedings of the 18th International Conference on Web and Internet Economics (WINE)*. 311–329. doi:10.1007/978-3-031-22832-2_18 [→ p. 2]

Simon Rey and Jan Maly. 2023. The (computational) social choice take on indivisible participatory budgeting. arXiv:2303.00621 [cs.GT] https://arxiv.org/abs/2303.00621 [→ p. 1, 3]

Ralph Tyrell Rockafellar. 1970. *Convex Analysis*. Princeton Mathematical Series, Vol. 28. Princeton University Press. doi:10.1515/9781400873173 [→ p.]

Andrzej Ruszczynski. 2011. *Nonlinear Optimization*. Princeton University Press. doi:10.2307/j.ctvcm4hcj [→ p.]

Paul A. Samuelson. 1954. The pure theory of public expenditure. *The Review of Economics and Statistics* 36, 4 (1954), 387–389. doi:10.2307/1925895 [→ p. 1]

Vadim I Shmyrev. 1983. On an approach to the determination of equilibrium in elementary exchange models. In *Doklady Akademii Nauk SSSR*, Vol. 268:5. 1062–1066. https://www.mathnet.ru/eng/dan10141 (In Russian). [→ p. 2, 5]

Vadim I. Shmyrev. 2009. An algorithm for finding equilibrium in the linear exchange model with fixed budgets. *Journal of Applied and Industrial Mathematics* 3, 4 (2009), 505. doi:10.1134/S1990478909040097 [→ p. 2, 5, 6]

Jean-Jacques Strodiot, V. Hien Nguyen, and Norbert Heukemes. 1983. ε-optimal solutions in nondifferentiable convex programming and some related questions. *Mathematical Programming* 25 (1983), 307–328. doi:10.1007/BF02594782 [→ p.]

Mashbat Suzuki and Jeremy Vollen. 2024. Maximum flow is fair: A network flow approach to committee voting. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*. 964–983. doi:10.1145/3670865.3673603 [→ p. 1, 2, 6]

Anne van den Nouweland. 2015. Lindahl and equilibrium. In *Individual and Collective Choice and Social Welfare: Essays in Honor of Nick Baigent*, Constanze Binder, Giulio Codognato, Miriam Teschl, and Yongsheng Xu (Eds.). Springer, 335–362. doi:10.1007/978-3-662-46439-7_18 [→ p. 2]

Yehuda Vardi and D. Lee. 1993. From image deblurring to optimal investments: Maximum likelihood solutions for positive linear inverse problems. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 55, 3 (1993), 569–598. doi:10.1111/j.2517-6161.1993.tb01925.x [→ p. 3]

Yehuda Vardi, Larry A Shepp, and Linda Kaufman. 1985. A statistical model for positron emission tomography. *Journal of the American statistical Association* 80, 389 (1985), 8–20. doi:10.1080/01621459.1985.10477119 [→ p. 3]

Vijay V. Vazirani. 2012. The notion of a rational convex program, and an algorithm for the Arrow-Debreu Nash bargaining game. *Journal of the ACM (JACM)* 59, 2 (2012), 1–36. doi:10.1145/2160158.2160160 [→ p. 5]

Vijay V. Vazirani and Mihalis Yannakakis. 2011. Market equilibrium under separable, piecewise-linear, concave utilities. *Journal of the ACM (JACM)* 58, 3 (2011), 1–25. doi:10.1145/1970392.1970394 [→ p.]

Nisheeth K. Vishnoi. 2021. *Algorithms for Convex Optimization*. Cambridge University Press. doi:10.1017/9781108699211 https://convex-optimization.github.io/. [→ p. 7]

Fang Wu and Li Zhang. 2007. Proportional response dynamics leads to market equilibrium. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC)*. 354–363. doi:10.1145/1250790.1250844 [→ p. 2, 6]

Li Zhang. 2011. Proportional response dynamics in the Fisher market. *Theoretical Computer Science* 412, 24 (2011), 2691–2698. doi:10.1016/j.tcs.2010.06.021 [→ p. 2, 6]

Renbo Zhao. 2023. Convergence rate analysis of the multiplicative gradient method for PET-type problems. *Operations Research Letters* 51, 1 (2023), 26–32. doi:10.1016/j.orl.2022.11.010 [→ p. 2, 3, 6]

Renbo Zhao and Robert M. Freund. 2023. Analysis of the Frank–Wolfe method for convex composite optimization involving a logarithmically-homogeneous barrier. *Mathematical Programming* 199, 1 (2023), 123–163. doi:10.1007/s10107-022-01820-9 [→ p.]