Efficiently Computing Equilibria in Budget-Aggregation Games

Patrick Becker
TU Munich
Munich, Germany
patrick.becker@tum.de

Alexander Fries
TU Munich
Munich, Germany
fries.alexander@web.de

Matthias Greger
Paris Dauphine - PSL
Paris, France
matthias.greger@dauphine.psl.eu

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

ABSTRACT

Budget aggregation deals with the social choice problem of distributing an exogenously given budget among a set of public projects, given agents' preferences. Taking a game-theoretic perspective, we study *budget-aggregation games* where each agent has virtual decision power over some fraction of the budget. We investigate the structure and show efficient computability of Nash equilibria for various common preference models in this setting. In particular, we show that equilibria for Leontief utilities can be found in polynomial time, solving an open problem from Brandt et al. [15], and give an explicit polynomial-time algorithm for computing equilibria for ℓ_1 preferences.

KEYWORDS

Budget aggregation; Equilibrium computation; Nash equilibrium

ACM Reference Format:
Patrick Becker, Alexander Fries, Matthias Greger, and Erel Segal-Halevi. 2026. Efficiently Computing Equilibria in Budget-Aggregation Games. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 10 pages. https://doi.org/10.65109/CCMB1463

1 INTRODUCTION

Participatory budgeting [6, 18] is a democratic process in which citizens are allowed to vote on how the funds of their city or state should be distributed. It is justified both by ethical arguments (people have a fundamental right to decide how their money is used) and by practical arguments (letting people decide on the budget strengthens cohesion in the community). This approach can be taken to the extreme: instead of just letting citizens vote on which projects they approve (as is commonly done today), one can give each citizen their own (virtual) share of the budget (e.g. $1/n$ of the total budget, in case there are n citizens with equal entitlements), and allow them to decide how to allocate that share among the potential projects. This can be an attractive approach, as it gives full power to the citizens and ensures some degree of fairness by splitting that power equally among them. Furthermore, each citizen is able to directly observe her influence on the outcome by means of the distribution of her individual budget.

However, this approach raises severe strategic problems. For example, if a citizen thought that projects x and y should each receive 0.5 of the total budget, but later observed that most other citizens allocated their shares to x, then she would have preferred to spend her entire budget on y, leaving her regretful since her budget was spent suboptimally in hindsight. This bad user experience can be avoided if we have the possibility to coordinate individual budgets by advising each citizen in advance how to distribute their own budget such that for each citizen, the advised distribution is optimal, given that all other citizens follow the provided advice.

In other words, we would like to compute a *Nash equilibrium (NE)* of the game in which the players are citizens, and the set of strategies of each player is the set of possible distributions of their share of the budget. We call this game the *budget-aggregation game*. An NE of the budget-aggregation game hits a sweet spot between giving complete autonomy to the citizens to distribute their shares of the budget and ensuring that they are indeed satisfied with their own choice. More specifically, every agent receives at least their maximin share, a property studied in fair division [see, e.g., 14]. In words, an agent receives the highest possible utility they can secure with their share of the budget under the least favorable allocation of the remaining shares.

Our contributions. We propose a new class of normal-form games, called budget-aggregation games, which model (equal-resource) public funding situations. We position these games within the broader framework of normal-form games and survey existing results on the computation of NEs. On top of that, our paper analyzes NEs for the most common utility functions that have been considered in participatory budgeting and related fields. We prove theorems on the structure of NEs and develop algorithms to efficiently compute them. We focus on the following utility models:

- For *linear utilities*, we show that an NE can be computed by solving a linear program (Section 4).
- For *Leontief utilities*, we show that an NE can be computed in polynomial time using a variant of the ellipsoid method, solving an open problem from Brandt et al. [15] (Section 5).
- For *binary symmetric separable utilities*, we explore strategic similarities to Leontief utilities to show that an NE can be found in polynomial time. (Section 6).
- For ℓ_1 *disutilities*, we again present an intuitive polynomial-time algorithm to compute an NE (Section 7).

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. https://doi.org/10.65109/CCMB1463

Finally, we extend the basic model by allowing agents to have different weights. We show that, for most of the above utility models (namely linear, Leontief and ℓ_1), an NE can still be computed in pseudo-polynomial time (Section 8).

The full version of this paper [8] includes complete proofs of Lemmas 5.4, 5.6 and 8.2 and Theorem 5.3 as well as additional examples and a discussion of related literature on market equilibria.

2 RELATED WORK

About 75 years ago, two groundbreaking papers by John Nash [59, 60] proved the existence of an NE in general normal-form games. Apart from its intrinsic value, the proof technique of applying fixed-point theorems, especially Kakutani's fixed-point theorem (which Nash attributed to David Gale), paved the way for other equilibrium existence theorems like Debreu's social equilibrium [32] or general equilibrium in the Arrow-Debreu model [3, 58].

An important generalization of Nash's theorem involves the notion of a *concave game* [69]. A concave game is defined by a collection of convex sets $(S_i)_{i\in N}$, where S_i represents the set of possible strategies of agent i, and a collection $(u_i)_{i\in N}$ of individual utility functions u_i that map each strategy profile to a real number; u_i can be any function that is continuous in all strategies, and concave in the strategy of agent i.[1] Nash's game is a special case in which every S_i is a standard simplex in some Euclidean space (representing the set of all lotteries over pure actions), and every u_i is a linear combination of products of n terms. Our budget-aggregation game is a special case in which every S_i is (still) a simplex, but the utility functions are substantially different, due to the interpretation as a distribution of exogenously given budget rather than a lottery.

2.1 Computing an equilibrium is hard

Like most fixed-point-based proofs, the ones by Nash [59, 60] and Rosen [69] do not yield efficient algorithms for finding an equilibrium. Moreover, there are 3-player games with integer payoffs with a unique (mixed) NE in which all probabilities are irrational numbers [10, 60], and 4-player games with a unique NE in which all probabilities are *irradical* numbers (cannot be expressed using integer roots) [62]. This shows that no finite algorithm can compute an exact NE in general games.

Even computing an ϵ-approximate NE probably cannot be done in polynomial time, as it is PPAD-hard even for two players. This was first proven for ϵ that shrinks exponentially in the number of actions [21], then for ϵ that shrinks polynomially [22], and later even for some constant $\epsilon > 0$ [72]. Computing an equilibrium is PPAD-complete, also for some other classes of concave games [63].

However, there are other classes of games for which an NE can be computed efficiently. We survey some of them below.

2.2 Nash games

For the original game studied by Nash (where the strategy set of each agent is the set of lotteries over a finite set of actions), many practical algorithms for computing an NE were designed, such as the ones by Lemke and Howson [54] and Porter et al. [68], but they do not have a polynomial-time guarantee.

Polynomial-time algorithms for ϵ-approximate NE (where no agent can increase her utility by more than ϵ by deviating, where the utilities are normalized to $[0, 1]$) are known for $\epsilon \approx 0.34$ for two players [73], $\epsilon \approx 0.5$ for n-player polymatrix games, and $\epsilon > 0.6$ for general n-player games with $n \geq 3$ [34]. Some special classes of games in which an exact NE can be computed in polynomial time are two-player zero-sum games [74], anonymous games in which each player has a fixed small number of strategies and Lipschitz-continuous utility functions [23, 29, 30], graphical games in which the graph (representing the influence between players) is a tree [49] or a general graph with degree 2 [36], win-lose games with at most 2 winning positions per action of each player [25], and two-player games in which the payoff matrices have a bounded rank [2, 57].

2.3 Aggregative games

An *aggregative game* is a game in which the utility of each player i depends only on the strategy of i, and on some aggregate function of all strategies. Our game corresponds to the special case in which that aggregate function is a sum. Aggregative games were initially studied for one-dimensional strategy sets (each agent chooses a number) and one-dimensional aggregators. Under certain conditions, such games admit a *potential function*,[2] and thus always have an NE [27, 35, 48, 52]. There are also efficient algorithms for computing an NE when players have discrete strategy sets [50].

For aggregative games with continuous strategy sets, Babichenko [7] proves that, for every $\epsilon > 0$, there exists a best-reply dynamics that reaches an $O(\lambda\epsilon)$-approximate NE in time (number of steps) $O(n \log n/\epsilon^2)$, where λ is the Lipschitz constant of the agents' utility functions. This result is applicable to a special case of our budget-aggregation game in which there are only $m = 2$ issues, as in that case, the strategy of each player is completely determined by the amount he puts on the first issue. As far as we know, it is open whether an *exact* NE for $m = 2$ can be computed in polynomial time for this general class of games.

Cummings et al. [28] study aggregative games with a multi-dimensional aggregator (as in our game), but the players' strategy sets are finite (unlike our game). For this class, they present an algorithm that computes an NE in time polynomial in the number of players and actions, but exponential in the number of dimensions of the aggregator.

When both players' strategy sets and the aggregator are multi-dimensional, one can still prove that, under the required assumptions, they have a potential function, and therefore have an NE [47]. However, computing the NE becomes harder. The literature focuses on numeric, distributed algorithms that converge asymptotically to an NE [e.g., 64, 76], see Li et al. [55] for a recent survey. We could not find an algorithm with a polynomial run-time guarantee.

Our contribution is identifying a natural and practical class of aggregative games, for which an exact NE can be computed efficiently. In the following, different types of aggregative games are covered.

[1]The definition in Rosen [69] is even more general and allows *coupled constraints*, that is, the set of strategies available to each agent may depend on the choices made by other agents. We do not need this generalization here.

[2]Roughly, a potential function is a function of the strategy profile, whose value increases when any player makes an improving move. A weaker definition only requires the potential to increase when any player makes a best response.

2.3.1 Congestion games. These games are a subclass of aggregative games with multidimensional aggregators, where each coordinate of the aggregator represents the congestion on one of the resources. Usually, the players' strategy sets are *discrete*; each player can choose a subset of resources to use (lotteries are not allowed). Equilibria without lotteries are called *pure* Nash equilibria (PNE).

Even though Nash's original existence proof is not valid in general for discrete strategy sets, congestion games (CGs) always have a potential function, so every sequence of individual improvements must terminate in a PNE [70]. However, in general, this sequence might be exponentially long. Specifically, computing a PNE in general CGs is PLS-complete, which means that it probably cannot be done by a polynomial-time algorithm [38].[3]

Some special classes of CGs in which a PNE can be computed or approximated in polynomial time are:

- Symmetric network CGs: all players have the same set of possible strategies, and this set is the set of paths between a source node and a target node in a given network [38].
- Nonatomic CGs: CGs in which there is a continuum of players, where each player's effect on the outcome is negligible (in contrast to our setting, where each player has a substantial effect on the outcome). Finding a PNE in a nonatomic CG can be reduced to solving a convex optimization problem, which can be done in weakly polynomial time. In the special case of a *network* nonatomic CG with Lipschitz-continuous utility functions, there is an FPTAS (an algorithm computing an ϵ-approximate PNE in time strongly-polynomial in the problem size and $1/\epsilon$). [38].
- CGs in which the strategy space of each player consists of the bases of a matroid over the set of resources [1].
- For symmetric (non-network) CGs, there is an FPTAS based on simple greedy dynamics [24].
- For general CGs, when the delay functions are polynomials of bounded degree, there are dynamics based on clever sequences of greedy steps that obtain a constant-factor approximate PNE [20].

The CG class most closely related to our budget-aggregation game is the class of *splittable (atomic) CGs* [61, 71], in which each agent controls a divisible load and can split it arbitrarily among different paths in a network. The players' utilities are decreasing functions of the total load allocated to edges that they use: $u_i(\delta) = -\sum_{x\in A} \delta_{i,x} \cdot c_{i,x}(\delta_x)$, where $c_{i,x}$ is the congestion function of edge x for player i. An important special case is when the congestion functions are affine, $c_{i,x}(\delta_x) = a_{i,x} \cdot \delta_x + b_{i,x}$. Klimm and Warode [51] prove that, even with affine congestion functions, computing a PNE is PPAD-complete. Still, polynomial-time algorithms are known for some special cases. If the congestion functions are affine and player-independent, an approximate PNE can be computed by convex programming for general graphs [26]; for specific classes of graphs called "well-designed", an exact PNE can be computed by a combinatorial algorithm [44]. When strategy sets are singletons — the strategy set of each player is a set of edges, rather than a set of paths from a source to a target — there is a polynomial-time algorithm for affine player-specific congestion functions [43]; for convex player-independent congestion functions, there are algorithms that are polynomial when either the number of players or the number of edges is fixed [9].

Note that the utility functions are different than our linear utility functions in that the (negative) utility that a player derives from an edge is multiplied by the player's own contribution to that edge. In other words, the dependence of $u_i(\delta)$ on $\delta_{i,x}$ is quadratic, as $\delta_{i,x}$ both multiplies $c_{i,x}$ and contributes to δ_x.

[3] PLS is similar to PPAD in that both contain problems in which a solution is guaranteed to exist — both are subsets of TFNP. It is believed that none of them contains the other, and that none of them is contained in P.

2.3.2 Budget-aggregation games with approval preferences. The need to study NEs in budget-aggregation games is also reflected by the fact that many fairness-motivated solution concepts in participatory budgeting and multiwinner voting [see, e.g., 53] are related to NEs, especially when agents have approval preferences, i.e., linear utilities with binary valuations, also called dichotomous utilities [12]. For the uncapped setting of budget-aggregation, when projects can receive arbitrary amounts of the budget, Brandl et al. [13] coined the term "decomposability" for distributions of the overall budget that admit a decomposition into distributions of the individual budgets where each agent spends her budget optimally, i.e., only on approved projects. The standard model in participatory budgeting assumes that projects have fixed costs. For such problems, common rules like the *method of equal shares* [65, 66] or *Phragmén's rule* [46, 67] rely on the same idea that agents successively allocate their budgets to fund approved projects. Most of these works take a "mechanism design perspective" where agents report their preferences, which are then aggregated into a collective outcome, instead of individual budget distributions. A notable exception is *donor coordination* [see, e.g., 39] where the budget is initially owned by the agents. Brandt et al. [15, 17] proved that, with Leontief utilities, the unique equilibrium distribution is the unique distribution that maximizes the product of agents' utilities where agents are weighted by their contributions.

For the special case of multiwinner approval voting (where all projects have the same costs and correspond to candidates), Haret et al. [42] recently defined a *budgeting game*, which is very similar to our budget-aggregation game. In their game, as in ours, each player has a fixed budget and can choose a distribution of this budget among m issues. However, their issues are *discrete*; each issue is *active* if its total allotment is at least 1 and *inactive* otherwise. The (discrete) utility of a player is the number of active issues that the player approves. They prove that their budgeting game always has an NE, but do not elaborate on its computational aspects.

It is also worth noting that returning NEs already ensures some degree of fairness as they satisfy *individual fair share* [4, 12] for arbitrary preference models (see Proposition 3.4).

3 MODEL

Let N be a set of n agents and A be a set of m (pure) alternatives. The outcome space $\Delta(A)$ is formed by all distributions $\delta \in [0,1]^m$ with $\sum_{x\in A} \delta_x = 1$. For a distribution δ, we denote its support by $\text{supp}(\delta) := \{x \in A : \delta_x > 0\}$. In a budget-aggregation game, the budget (w.l.o.g., normalized to 1) is equally divided among the agents, and each of them is allowed to distribute their share among the alternatives. The set of strategies available to each payer is thus $S_i = \{\delta_i \in [0,1]^m$ such that $\sum_{x\in A} \delta_{i,x} = 1/n\}$, which is an $(m-1)$-dimensional simplex in $\mathbb{R}^m$.

Definition 3.1. A budget-aggregation game is defined as the tuple $(N, A, (u_i)_{i\in N})$ consisting of

- a set of agents N with $n = |N|$,
- a set of alternatives (issues) A,
- and n individual utility functions $u_i\colon \Delta(A) \to \mathbb{R}$ for all $i \in N$.

The strategy of agent i, denoted by $\delta_i = (\delta_{i,x})_{x\in A}$, states that agent i allocates $\delta_{i,x}$ to alternative x; these strategies are called *individual distributions*. We denote the *overall distribution* by $\delta := \sum_{i\in N} \delta_i$. Note that agents only care about the overall distribution of the budget, which is why the domain of u_i is $\Delta(A)$ instead of $S_1 \times \cdots \times S_n$. Furthermore, let $S_{-i} := S_1 \times \ldots \times S_{i-1} \times S_{i+1} \times \ldots \times S_n$ be the space of reduced strategy profiles where we exclude agent i.

Definition 3.2. A strategy profile $(\delta_i^*)_{i\in N}$ is called a *Nash equilibrium* if for all $i \in N$,

$$u_i(\delta^*) = \max_{\delta_i\in S_i} u_i\left(\sum_{j\in N\setminus i} \delta_j^* + \delta_i\right).$$

The sum $\delta^* = \sum_{i\in N} \delta_i^*$ is called an *equilibrium distribution.*

Interestingly, NEs already guarantee some degree of fairness, as they satisfy a natural fairness notion called *individual fair share.*

Definition 3.3. A strategy profile $(\delta_i)_{i\in N}$ satisfies *individual fair share* if for all $i \in N$,

$$u_i\left(\sum_{j\in N} \delta_j\right) \geq \max_{\delta_i'\in S_i} \min_{(\delta_j')_{j\neq i}\in S_{-i}} u_i\left(\sum_{j\in N} \delta_j'\right).$$

In words, individual fair share ensures that each agent obtains at least as much utility as she could guarantee on her own, by allocating her budget independently of the others.

Proposition 3.4. *Every Nash equilibrium $(\delta_i^*)_{i\in N}$ satisfies individual fair share.*

Proof. For every NE $(\delta_i^*)_{i\in N}$,

$$\begin{aligned} u_i(\delta^*) &= \max_{\delta_i\in S_i} u_i\left(\sum_{j\in N\setminus i} \delta_j^* + \delta_i\right) \\ &\geq \max_{\delta_i\in S_i} \min_{(\delta_j)_{j\neq i}\in S_{-i}} u_i\left(\sum_{j\in N\setminus i} \delta_j + \delta_i\right). \end{aligned}$$

Thus, $(\delta_i^*)_{i\in N}$ satisfies individual fair share. □

We are interested in understanding the structure and computing NEs for various utility models. In the following sections, we will consider linear utilities (Section 4), Leontief utilities (Section 5), binary symmetric separable utilities (Section 6), and l_1 disutilities (Section 7).

4 LINEAR UTILITIES

Linear utilities, also known as *additive* or *von Neumann-Morgenstern utility functions*, are the most common extension of valuations for a set of pure alternatives A to preferences over $\Delta(A)$. When agents' preferences satisfy four seemingly mild conditions, von Neumann and Morgenstern [75] showed that they can be represented by linear utility functions. Formally, i has linear utilities if for all $\delta \in \Delta(A)$,

$$u_i(\delta) = \sum_{x\in A} v_{i,x}\delta_x.$$

with nonnegative valuations $(v_{i,x})_{x\in A}$ for alternatives.[4]

From an economic perspective, alternatives admit an interpretation as substitutes, as agent i is indifferent to how the budget is divided between two alternatives x, y with $v_{i,x} = v_{i,y}$. Due to this property, the set of NEs admits an easily understandable structure. Let $A_i^{\max} := \arg\max_{x\in A} v_{i,x}$ be the set of alternatives highest valued by agent i.

Proposition 4.1. *For linear utilities, the set of Nash equilibria consists of all $(\delta_i)_{i\in N}$ with $\mathrm{supp}(\delta_i) \subseteq A_i^{\max}$ for all $i \in N$.*[5]

Proof. By Definition 3.2, $(\delta_i^*)_{i\in N}$ is an NE if and only if

$$u_i(\delta^*) = \max_{\delta_i\in S_i} u_i(\delta^* - \delta_i^* + \delta_i)$$

for all $i \in N$. By linearity of u_i, this is equivalent to

$$\frac{n-1}{n} \cdot u_i\left(\frac{n}{n-1} \cdot (\delta^* - \delta_i^*)\right) + \frac{1}{n} \cdot \max_{\delta_i\in S_i} u_i(n \cdot \delta_i).$$

The term $u_i(n \cdot \delta_i)$ is maximal if and only if $\mathrm{supp}(\delta_i) \subseteq A_i^{\max}$, which proves the statement. □

It is straightforward to see that the set of NEs is convex. Its simple structure allows for the computation of NEs with additional properties, e.g., the ones that maximize utilitarian welfare among equilibrium distributions.[6]

Furthermore, given a distribution δ, it is easy to check whether it can be decomposed into an NE by solving the following linear feasibility program with variables $d_{i,x}$ (for $i \in N$ and $x \in A$).

$$\begin{aligned} \delta_x &= \sum_{i\in N} d_{i,x} && \text{for all } x \in A; \\ 1/n &= \sum_{x\in A_i^{\max}} d_{i,x} && \text{for all } i \in N; \\ d_{i,x} &\geq 0 && \text{for all } i \in N, x \in A. \end{aligned}$$

5 LEONTIEF UTILITIES

Brandt et al. [15, 17] introduced *Leontief preferences* for contexts where alternatives share characteristics of complements — meaning that agents only benefit when they receive all of them (in a specific ratio), rather than being able to substitute one for another. For each agent i, let $A_i^+ := \{x \in A\colon v_{i,x} > 0\}$ denote the set of alternatives the agent values positively. Leontief preferences are represented by

$$u_i(\delta) = \min_{x\in A_i^+} \delta_x/v_{i,x}.$$

These authors characterize the set of all NEs with the help of the concept of *critical alternatives.*

[4] Agents do not have different valuations for individual distributions $\delta_{i,x}$, $\delta_{j,x}$ to the same alternative x as they only care about the total amount δ_x allocated to x.

[5] If we interpret the overall distribution as a probability distribution, Proposition 4.1 implies that a distribution is in equilibrium if and only if it is the outcome of a random dictatorship, where a player is selected uniformly at random, and then selects one of her best outcomes. Random dictatorship is sometimes used as a benchmark of fairness; see, e.g., Bogomolnaia et al. [12].

[6] Such equilibrium distributions might still be Pareto dominated by other (non-equilibrium) distributions [see, e.g., 13].

Definition 5.1 (Brandt et al. [17]). Given a distribution $\delta \in \Delta$, agent i's set of critical alternatives is defined as

$$T_{\delta,i} \coloneqq \arg\min_{x \in A_i^+} \frac{\delta_x}{v_{i,x}}.$$

An alternative is critical in the sense that decreasing its share also decreases the utility of agent i.

Lemma 5.2 (Brandt et al. [17], Lemma 1). *In a budget-aggregation game where all agents have Leontief utilities, a profile $(\delta_i)_{i \in N}$ is an NE if and only if $\text{supp}(\delta_i) \subseteq T_{\delta,i}$ for all $i \in N$ (where $\delta := \sum_{i \in N} \delta_i$).*

Brandt et al. further showed that all NEs result in the same overall distribution, which can be computed by solving the convex program obtained from maximizing the product of agents' utilities.

They leave the question open as to whether the equilibrium distribution can be computed exactly in polynomial time. The remainder of this section is dedicated to answering that question in the affirmative.

Theorem 5.3. *For Leontief utilities, a Nash equilibrium (and the corresponding unique equilibrium distribution) can be computed in time polynomial in the binary encoding length of the input.*

Note that the equilibrium distribution does not change when individual valuations are rescaled. Similarly, rescaling contributions preserves the share each alternative receives. Thus, for the sake of simplicity, we assume throughout this section that all valuations are natural numbers. We prove that the equilibrium distribution can be computed in time $\text{poly}(n, m, \log_2(v_{\max}))$ where $v_{\max} \coloneqq \max_{i \in N, x \in A} v_{i,x}$, i.e., the run-time is polynomial in the binary encoding length of the input.

As a first step, we prove an upper bound on the binary encoding length of the equilibrium distribution δ^* and its corresponding utility profile (which we denote by u^*).

Lemma 5.4. *If the agents' valuations $v_{i,x}$ are natural numbers, then the equilibrium distribution δ^* and its utility profile u^* are rational-valued. Moreover, their binary encoding length is bounded by a polynomial function of the binary encoding length of $v_{i,x}$ and n.*

Proof sketch. Given the critical alternatives of each agent under the equilibrium distribution, we can set up a linear program to compute it. Thus, the equilibrium distribution and its utility profile are rational-valued.

Moreover, knowing that agents only contribute to their critical alternatives, the shares of alternatives x, y that are critical to some agent i can be linearly related to each other through her valuations $v_{i,x}$ and $v_{i,y}$. Solving the resulting system of linear equations, one can prove an upper bound on the binary encoding length of each δ^*_x and, consequently, also on u^*. □

Lemma 5.4 cannot be used directly for computing the equilibrium distribution in polynomial time, since the proof requires us to know $T_{\delta^*,i}$; iterating over all possible $T_{\delta^*,i}$ would require exponential time. To prove polynomial-time computability, we leverage the following lemma, which is Theorem 13 of Jain [45]:

Lemma 5.5 (Jain [45]). *Let S be a convex set given by a strong separation oracle, and $\phi > 0$ an integer.*

There is an oracle-polynomial time and ϕ-linear time algorithm which does one of the following:

(1) Concludes that there is no point in S with binary encoding length at most ϕ, or
(2) produces a point in S with binary encoding length at most $P(n) \cdot \phi$, where $P(n)$ is a polynomial.

We apply Lemma 5.5 as follows. For every positive rational number z_0, we define a convex set $S(z_0) \subseteq \mathbb{R}^{n+m}$, where the variables are u_i for $i \in N$ and d_x for $x \in A$:

$$\begin{aligned}
\prod_{i=1}^{n} {u_i}^{1/n} &\geq z_0; && \\
d_x &\geq u_i \cdot v_{i,x} && \text{for all } i \in N, x \in A; \\
\sum_{x \in A} d_x &= 1; && \\
u_i &\geq 0 && \text{for all } i \in N; \\
d_x &\geq 0 && \text{for all } x \in A;
\end{aligned}$$

Intuitively, the set $S(z_0)$ represents all pairs (δ, u) such that δ is a feasible distribution, u is a lower bound on the corresponding utility profile, and the geometric mean of all utilities is at least z_0. A *strong separation oracle for* $S(z_0)$ is a function that accepts as input a rational vector $\mathbf{y}' = (u'_1, \ldots, u'_n, d'_1, \ldots, d'_m)$. It should return one of two outcomes: either an assertion that $\mathbf{y}' \in S(z_0)$, or a hyperplane that separates $\mathbf{y}'$ from $S(z_0)$ (that is, a rational vector $\mathbf{c}$ such that $\mathbf{c}^\top \mathbf{y}' < \mathbf{c}^\top \mathbf{y}$ for all $\mathbf{y} \in S(z_0)$).

Lemma 5.6. *For every rational $z_0 > 0$, there is a polynomial-time strong separation oracle for the convex set $S(z_0)$.*

Proof sketch. The main challenge is to find a separating hyperplane for vectors outside of $S(z_0)$ for which only the non-linear constraint from above is violated. Using a similar idea as Jain [45], we apply the inequality of weighted arithmetic mean and weighted geometric mean to find such a hyperplane. □

Proof sketch of Theorem 5.3. Lemma 5.6 allows us to apply Lemma 5.5 to $S(z_0)$. We can now compute the equilibrium distribution by applying binary search to z_0 and find an NE by turning the linear program described in the proof of Lemma 5.4 into a linear feasibility problem where the variables u_i and d_x are replaced by their respective values in the equilibrium distribution. □

Together with the uniqueness of the overall distribution, this implies that the set of NEs is convex. Theorem 5.3 is also applicable to other utility models where NEs coincide with the ones for Leontief utilities, e.g., Cobb-Douglas utility functions [17].

6 BINARY SYMMETRIC SEPARABLE UTILS

In this section, we investigate NEs for very general utility models when agents have separable preferences. Separability requires that agents receive utility independently from each alternative. More formally, fixing the outcome on a subset of all alternatives has no influence on an agent's preferences over how to allocate the remaining budget among the other alternatives. The interested reader is advised to consult Blackorby et al. [11], Debreu [33] for an axiomatic characterization of separable preferences.

Here, we say that preferences are separable if they can be represented by a utility function of the form

$$u_i(\delta) = \sum_{x \in A} g_{i,x}(\delta_x)$$

where $g_{i,x}: \mathbb{R}_{\geq 0} \to \mathbb{R}$ are continuous and weakly increasing.

Linear utilities represent the most common class of separable preferences. Contrarily, Leontief preferences are not separable as the minimum is taken over all alternatives with positive valuations.

We consider the class of *binary symmetric* separable preferences where each agent i has a non-empty set of approved alternatives $A_i \subseteq A$ but does not further distinguish between the alternatives.[7] Furthermore, her utility is assumed to *strictly* increase when the budget on some $x \in A_i$ increases.

In terms of utility functions, such preferences are represented by

$$u_i(\delta) = \sum_{x \in A_i} g_i(\delta_x)$$

where $g_i: [0,1] \to \mathbb{R}$ are continuous and strictly increasing.

For example, setting $g_i(\delta_x) = \delta_x$ for all $i \in N$ corresponds to linear preferences with binary valuations $v_{i,x} \in \{0,1\}$ [see, e.g., 12]. Choosing $g_i(\delta_x) = \ln(\delta_x)$ leads to Cobb-Douglas utilities with binary weights. However, note that our model allows the g_i's to differ among agents.

The following example shows that, for such general preferences, NEs do not always exist.

Example 6.1. Suppose $n = 2$ and $A = \{a, b\}$, with $g_1(\cdot) = \sqrt{\cdot}$, i.e., $u_1(\delta) = \sqrt{\delta_a} + \sqrt{\delta_b}$ and $g_2(\cdot) = \cdot^2$, i.e., $u_2(\delta) = \delta_a^2 + \delta_b^2$.

If $\delta_a > \delta_b$, then $\delta_a > 1/2$, hence both agents contribute to a. By strict concavity of g_1, Agent 1 has an incentive to move some budget from a to b. The case $\delta_a < \delta_b$ can be treated analogously.

If $\delta_a = \delta_b$, then by strict convexity of g_2, Agent 2 has an incentive to move her entire budget to an alternative $x \in \{a, b\}$ with $\delta_{1,x} \geq 1/4$, resulting in $\delta_a \neq \delta_b$. Thus, there cannot exist an NE.

Intuitively, Agent 1 wants the entire budget to be distributed as equally as possible, whereas Agent 2 prefers to focus on one alternative. Therefore, we treat separately the two cases where all g_i's are strictly concave (Section 6.1) and strictly convex (Section 6.2).

6.1 Strictly concave utilities

In this section, all g_i's are assumed to be strictly concave, that is, $g_i(\lambda\delta + (1-\lambda)\delta') > \lambda g_i(\delta) + (1-\lambda)g_i(\delta')$ for all different $\delta, \delta' \in \Delta$ and $\lambda \in (0,1)$.

Theorem 6.2 (Brandt et al. [17], Proposition 9). *For binary symmetric separable and strictly concave utilities, the set of NEs coincides with that for Leontief utilities with the same valuations ($v_{i,x} \in \{0,1\}$).*

For binary valuations, Brandt et al. [17] stated two different possibilities to compute an overall equilibrium distribution via linear programming. This distribution can be decomposed into an NE using the same linear program described at the end of the proof of Theorem 5.3.

Corollary 6.3. *For binary symmetric separable and strictly concave utilities, an NE can be computed via linear programming.*

[7] This is what we mean by the term "symmetric" in our context.

Furthermore, the concept of critical alternatives can readily be transferred in order to check whether a given strategy profile constitutes an NE.

6.2 Strictly convex utilities

In this section, all g_i's are assumed to be strictly convex, that is, $\lambda g_i(\delta) + (1-\lambda)g_i(\delta') > g_i(\lambda\delta + (1-\lambda)\delta')$ for all different $\delta, \delta' \in \Delta$ and $\lambda \in (0,1)$. Note that in this case, each utility function $u_i(\delta)$, seen as an m-dimensional real-valued function, is convex, as it is a sum of strictly convex functions. Furthermore, it is strictly convex if the domain is restricted to the $|A_i|$-dimensional subspace spanned by agent i's approved alternatives.

Let $M_{\delta,i} \coloneqq \arg\max_{x \in A_i} \delta_x$ be the set of alternatives with maximum share under δ that are approved by agent i.

Lemma 6.4. *For binary symmetric separable and strictly convex utilities, the set of Nash equilibria consists of all $(\delta_i^*)_{i \in N}$ with $\text{supp}(\delta_i^*) = M_{\delta^*,i}$ and $|\text{supp}(\delta_i^*)| = |M_{\delta^*,i}| = 1$ for all $i \in N$.*

Proof. The distribution $(\delta_i^*)_{i \in N}$ is an NE if and only if every agent i's distribution δ_i^* is a best response and therefore a solution of the optimization problem

$$\max_{\delta_i \in S_i} u_i(\delta^* - \delta_i^* + \delta_i).$$

The objective function is strictly convex on this domain, and the feasible region is the $(|A_i| - 1)$-dimensional simplex spanned by those strategies where agent i places her entire budget on a single approved alternative. So every maximizer is a vertex of this simplex, which proves that $|\text{supp}(\delta_i)| = 1$. By adding $1/n$ to an alternative $x \in A_i$ in $\delta^* - \delta_i^*$, agent i's utility increases by $g(\delta_x^* - \delta_{i,x}^* + 1/n) - g(\delta_x^* - \delta_{i,x}^*)$. This increase is maximal if and only if x is an alternative with maximal share under the distribution $\delta^* - \delta_i^*$, i.e., $x \in M_{\delta^* - \delta_i^*, i}$ since g is strictly convex. This proves that $(\delta_i^*)_{i \in N}$ is an NE if and only if, every agent plays a pure strategy, i.e., $|\text{supp}(\delta_i^*)| = 1$, and $\text{supp}(\delta_i^*) \subseteq M_{\delta^* - \delta_i^*, i}$ for all $i \in N$. Thus, $|M_{\delta^*,i}| = 1$ and $\text{supp}(\delta_i^*) = M_{\delta^*,i}$ for all $i \in N$.

For the other direction, it remains to show that if $\text{supp}(\delta_i) = M_{\delta,i}$ and $|M_{\delta,i}| = 1$ for all $i \in N$, it holds that $\text{supp}(\delta_i) \subseteq M_{\delta - \delta_i, i}$ for all $i \in N$. Therefore, assume $x \in \text{supp}(\delta_i)$ and $y \in A \setminus \{x\}$ for agent i. Then, since all agents play pure strategies, all δ_x are some multiple of $1/n$ and $\delta_x \geq \delta_y + 1/n$. Thus,

$$\delta_y - \delta_{i,y} = \delta_y \leq \delta_x - 1/n = \delta_x - \delta_{i,x}$$

and $\text{supp}(\delta_i) \subseteq M_{\delta - \delta_i, i}$. □

Corollary 6.5. *For a Nash equilibrium $(\delta_i^*)_{i \in N}$ and two alternatives $x, y \in A$, $\delta_x^* = \delta_y^*$ implies $|A_i \cap \{x, y\}| = 1$ for all $i \in N$ with $\text{supp}(\delta_i^*) \in \{x, y\}$.*

Lemma 6.4 implies that checking whether a given distribution δ admits an NE $(\delta_i)_{i \in N}$ can be done in polynomial time. In a first step, verify that $|M_{\delta,i}| = 1$ for all $i \in N$. Second, check that the decomposition $(\delta_i)_{i \in N}$ with $\delta_{i,x} = 1/n$ for $x \in M_{\delta,i}$ (and $\delta_{i,x} = 0$ otherwise) is a valid decomposition of δ.

Theorem 6.6. *For binary symmetric separable and strictly convex utilities, a Nash equilibrium always exists and can be found in polynomial time.*

Proof. We use the following greedy approach to construct an NE $(\delta_i^*)_{i\in N}$.[8] Let $N' \subseteq N$ denote the set of agents who have not yet allocated their budgets. In each step, select an alternative $x^* \in \arg\max_{x\in A} \sum_{i\in N':x\in A_i} 1/n$, i.e., x^* maximizes the aggregated share of the remaining agents who approve it. Each of these agents then allocates her entire budget to x^*, and is excluded from N'. The procedure is repeated until no agents remain.

This process runs in polynomial time in n and m as it terminates after at most n steps with a complete strategy profile (δ_i^*) where in each step, a suitable alternative x^* can be found by counting approvals of the remaining agents in N'. By construction, $|\text{supp}\, \delta_i^*| = 1$ for all $i \in N$. Furthermore, $\text{supp}(\delta_i^*) \subseteq M_{\delta^*,i}$. Otherwise, some agent i approves an alternative x with $\delta_x^* > \delta^*_{\text{supp}(\delta_i^*)}$ and would have contributed to x instead under the described greedy approach. Finally, if $|M_{\delta^*,i}| > 1$ for some agent i, then there exist alternatives $x, y \in A_i$ with $\delta_x^* = \delta_y^*$ and w.l.o.g. $\text{supp}(\delta_i^*) = x$, contradicting Corollary 6.5.

Thus, $(\delta_i^*)_{i\in N}$ constitutes an NE by Lemma 6.4. □

The following example shows that the set of NEs is not convex.

Example 6.7. Consider a profile with a single agent i and two alternatives x and y that are both approved by the agent, i.e., $A_i = A$. Using Lemma 6.4, it is easy to see that there exist exactly two equilibria, which are given by the two distributions that allocate the entire budget to one of the alternatives.

In fact, the number of NE is finite and can be bounded explicitly.

Proposition 6.8. *The size of the set of Nash equilibria is bounded by* $\min\{m!, \prod_{i\in N}|A_i|\}$. *These bounds are tight.*

The bounds correspond to the number of possible strict orders of the alternatives and the number of pure strategy profiles. The proof, including instances that show the tightness of these bounds, can be found in the full version of this paper [8].

7 ℓ_1 DISUTILITIES

A widely studied utility model in the context of budget aggregation is based on ℓ_1 preferences, also referred to as ℓ_1 disutilities [39, 41, 56]. The existing literature mostly focuses on the axiomatic investigation of mechanisms [16, 37], in particular, moving phantom mechanisms [19, 31, 40].

For ℓ_1 preferences, each agent is assumed to have a peak allocation — a preferred distribution of the budget — and her utility is determined by the negative ℓ_1 distance between any given allocation and her peak. Formally, let $p_i = (p_{i,x})_{x\in A} \in \Delta(A)$ denote the peak of agent i. Then, her utility for a given allocation $\delta \in \Delta(A)$ is defined as

$$u_i(\delta) = -\|\delta - p_i\|_1.$$

Alternatively, the utility can be interpreted in terms of overlap utilities, as shown by Goel et al. [41]. In the overlap utility model, the utility of agent i is defined as the sum of the overlaps between the allocation δ and the agent's peak p_i:

$$u_i(\delta) = \sum_{x\in A} \min(\delta_x, p_{i,x}).$$

[8] This algorithm is equivalent to an algorithm called *"maximum payment rule"* [5], which was recently presented as a rule for approximately-fair budget distribution for agents with binary linear utilities.

Goel et al. [41] demonstrate that ℓ_1 disutilities are equivalent (up to an affine transformation) to overlap utilities, making both models interchangeable in our budget-aggregation context. Note that these preferences are separable, but not binary symmetric separable.

These utilities capture the behavior of agents who derive value from receiving resources up to a certain satiation point, beyond which additional resources provide no further benefit. The utility is linear in the allocation up to the agent's peak, and constant thereafter — reflecting satiation once the desired threshold is met.

An alternative x is *oversupplied* for agent i at δ if $\delta_x > p_{i,x}$. Denote by $M_i^>(\delta)$ the set of such alternatives for a given δ. Similarly, we say x is *undersupplied* for agent i at δ if $\delta_x < p_{i,x}$. Let $M_i^<(\delta)$ be the set of undersupplied alternatives and additionally define $M_i^\leq(\delta)$ as the set of alternatives for which $\delta_x \leq p_{i,x}$. We show that a strategy profile constitutes an NE if and only if no agent contributes to one of her oversupplied alternatives.

Lemma 7.1. *For ℓ_1 preferences, a strategy profile $(\delta_i^*)_{i\in N}$ is a Nash equilibrium if and only if $supp(\delta_i^*) \subseteq M_i^\leq(\delta^*)$ for all $i \in N$.*

Proof. ⇒: Assume that there is some agent i with $\delta_{i,x}^* > 0$ for some alternative x which is oversupplied, i.e. $\delta_x^* > p_{i,x}$. Redistributing $\delta_{i,x}^*$ to some undersupplied alternative, say x', strictly increases the utility of agent i, showing that the strategy profile $(\delta_i^*)_{i\in N}$ is not an NE.

⇐: Suppose $(\delta_i^*)_{i\in N}$ satisfies $supp(\delta_i^*) \subseteq M_i^\leq(\delta^*)$ for all $i \in N$. Since the marginal utility decrease in x equals the marginal utility increase in x' for all $x, x' \in M_i^\leq(\delta^*)$, agent i cannot strictly increase her utility by redistributing her budget δ_i. Therefore, $u_i(\delta^*) = \max_{\delta_i'\in S_i} u_i(\delta^* - \delta_i^* + \delta_i')$, establishing that $(\delta_i^*)_{i\in N}$ is an NE. □

The existence of an NE for ℓ_1 disutilities follows from the more general existence result by Debreu [32] for quasi-concave utilities. Given our above characterization, it is not surprising that NEs do not yield a unique overall distribution in the case of ℓ_1 preferences, since the characterization depends solely on the support of the agents' individual distributions. An illustrative example and the proof of the following proposition can be found in the full version of this paper [8].

Proposition 7.2. *For ℓ_1 preferences, the set of equilibrium distributions is not convex.*

Although the set of NEs may be non-convex — demonstrating potential structural complexity — there nonetheless exists a polynomial-time algorithm that computes an NE.

Theorem 7.3. *For ℓ_1 preferences, a Nash equilibrium can be computed in polynomial time.*

Proof. We present the following polynomial-time algorithm. Let each agent i have a budget, denoted by b_i. Initially, this budget is set to $b_i = 1/n$ for all $i \in N$. Iterate over all alternatives $x \in A$ in an arbitrary order. For a given alternative x, order all agents in a descending order of their peaks for this alternative, i.e., $p_{i_1,x} \geq p_{i_2,x} \geq \ldots \geq p_{i_n,x}$. When it is agent i's turn, she allocates $\min(b_i, p_{i,x} - \delta_x)$ on alternative x. Afterwards, b_i is decreased and δ_x is increased by this amount. Once agent i has spent her entire share ($b_i = 0$), she is deleted from the set of considered agents. The ordering of the agents ensures that x does not become oversupplied to an agent who already contributed to x later on, i.e., no agent contributes to an oversupplied alternative. When no agent wants to contribute to alternative x, we go to the next alternative x' and reorder the remaining agents according to their peaks ($p_{i_1,x'} \geq p_{i_2,x'} \ldots \geq p_{i_n,x'}$). Then, we let them again sequentially contribute to this alternative.

We prove that every agent spends her entire budget after iterating over all alternatives. Assume for contradiction that there is some agent who has still some budget left after she had the possibility to contribute to every alternative. This would mean that $p_{i,x} \leq \delta_x$ for all $x \in A$ in the final partial distribution δ, which is impossible due to the fact that $p_i \in \Delta(A)$ but $\sum_{x \in A} \delta_x < 1$. All in all, the suggested construction returns a strategy profile that constitutes an NE.

Iterating over all m alternatives and, for every alternative, sorting at most n agents according to their peaks leads to a run-time complexity of $O(mn \log n)$. □

This construction provides an efficient algorithm for obtaining an NE given a set of optimal distributions of the agents. The full version of this paper [8] contains a linear feasibility program to decide whether a given overall distribution can be represented as individual distributions that constitute an NE.

Theorem 7.4. *For ℓ_1 preferences, one can check in polynomial time whether an overall distribution can be decomposed into an NE.*

8 WEIGHTED AGENTS

In practice, one might want to assign different shares to the agents, e.g., when they represent groups of citizens of different sizes, or citizens in different tax brackets.

We define a *weighted budget-aggregation game* as a tuple $(N, A, (u_i)_{i \in N}, (w_i)_{i \in N})$, where N, A and $(u_i)_{i \in N}$ are as in Definition 3.1, and w_i is the weight of agent i, with $\sum_{i \in N} w_i = 1$. The strategy set of each agent i is given by

$$S_i(w_i) := \{\delta_i \in [0,1]^m \text{ such that } \sum_{x \in A} \delta_{i,x} = w_i\}.$$

The definition of an NE transfers directly from Definition 3.2, except that S_i depends on w_i as stated above.

For some classes of utility functions, the weighted setting can be reduced to the unweighted setting by agent cloning if all weights are rational numbers, resulting in a pseudo-polynomial time algorithm. We formalize the agent cloning process below.

Definition 8.1. Given a weighted budget-aggregation game $(N, A, (u_i)_{i \in N}, (w_i)_{i \in N})$ where all w_i are rational numbers, let D be their least common denominator, such that $w_i = q_i/D$ for all $i \in N$, and all q_i are integers. The corresponding *cloned game* is an (unweighted) budget-aggregation game in which the set of agents N^c contains q_i clones of each agent i (with the same utility function u_i), so that $|N^c| = \sum_{i \in N} q_i = D$.

The following lemma connects NEs of the cloned game to NEs of the original weighted budget-aggregation game for all considered classes of utility functions.

Lemma 8.2. *Let $(N, A, (u_i)_{i \in N}, (w_i)_{i \in N})$ be a weighted budget-aggregation game in which all w_i are rational numbers, and all u_i are either linear, Leontief, binary symmetric separable and strictly concave/strictly convex, or ℓ_1. Let $(\delta_{i^c})_{i^c \in N^c}$ be a Nash equilibrium of the corresponding cloned game. Define a strategy profile $(\delta_i)_{i \in N}$ such that δ_i is the sum of δ_{i^c} for all i^c which are clones of i. Then $(\delta_i)_{i \in N}$ is a Nash equilibrium of the weighted game.*

Proof sketch. Assuming that $(\delta_i)_{i \in N}$ is no NE in the weighted game due to a beneficial deviation of agent i, it can be shown that for each utility model, some clone of agent i has a beneficial deviation in the cloned game. □

Theorem 8.3. *A Nash equilibrium can be computed in pseudo-polynomial time for weighted budget-aggregation games in which all utilities are either linear, Leontief, binary symmetric separable and strictly concave/strictly convex, or ℓ_1.*

Proof. Given a weighted budget-aggregation game with weights $w_i = q_i/D$, construct the corresponding cloned game, which is an unweighted game. This game has D agents; hence, an NE can be computed in time polynomial in m and D. By Lemma 8.2, this NE corresponds to an NE for the weighted game; the run-time is pseudo-polynomial in the parameters of that game. □

We believe Lemma 8.2 can be extended to other classes of utilities. It would be interesting to characterize the set of utility functions for which Lemma 8.2 remains true.

Moreover, it is open whether an NE in a weighted budget-aggregation game can be computed in time polynomial in n and m (rather than pseudo-polynomial), e.g., for Leontief preferences.

9 CONCLUSION

In this paper, we introduced budget-aggregation games to formalize a game-theoretic distributed approach for dividing a budget among a set of public projects and investigated NEs of the resulting games. Although finding NEs is known to be hard for general classes of utility models, our work shows that they can be computed in polynomial time for many common utility models, namely, linear, Leontief, and ℓ_1 preferences as well as for classes of binary symmetric and separable preferences.

It would be interesting to see whether there exists a class of utility functions containing all of the previously considered preference models that still allow for an efficient computation of NEs. As a first step, one might want to extend binary symmetric separable utilities to allow for valuations other than 0 or 1, or consider preferences based on ℓ_p metrics other than ℓ_1.

Other possible directions for future work include the selection of specific equilibria, e.g., by incorporating additional axioms like Pareto efficiency or investigating limits of equilibrium dynamics, and the investigation of the price of anarchy.

ACKNOWLEDGMENTS

We thank Yakov Babichenko, Felix Brandt, Warut Suksompong, and the anonymous reviewers of AAMAS 2026 for helpful comments. Patrick is supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/14-1. Erel is supported by the Israel Science Foundation grant no. 1092/24.

REFERENCES

[1] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. 2008. On the impact of combinatorial structure on congestion games. *Journal of the ACM (JACM)* 55, 6 (2008), 1–22.

[2] Bharat Adsul, Jugal Garg, Ruta Mehta, and Milind Sohoni. 2011. Rank-1 bimatrix games: A homeomorphism and a polynomial time algorithm. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC)*. 195–204.

[3] Kenneth J. Arrow and Gerard Debreu. 1954. Existence of an Equilibrium for a Competitive Economy. *Econometrica* 22, 3 (1954), 265–290.

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

[5] Haris Aziz, Patrick Lederer, Xinhang Lu, Mashbat Suzuki, and Jeremy Vollen. 2025. Approximately Fair and Population Consistent Budget Division via Simple Payment Schemes. In *Proceedings of the 26h ACM Conference on Economics and Computation (ACM-EC)*. 349.

[6] Haris Aziz and Nisarg Shah. 2021. Participatory Budgeting: Models and Approaches. In *Pathways between Social Science and Computational Social Science: Theories, Methods and Interpretations*, T. Rudas and P. Gábor (Eds.). Springer.

[7] Yakov Babichenko. 2018. Fast convergence of best-reply dynamics in aggregative games. *Mathematics of Operations Research* 43, 1 (2018), 333–346.

[8] Patrick Becker, Alexander Fries, Matthias Greger, and Erel Segal-Halevi. 2026. *Efficiently Computing Equilibria in Budget-Aggregation Games*. Technical Report. https://arxiv.org/abs/2509.08767.

[9] Umang Bhaskar and Phani Raj Lolakapuri. 2018. Equilibrium computation in atomic splittable routing games. In *26th Annual European Symposium on Algorithms (ESA 2018)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 58–1.

[10] Vittorio Bilò and Marios Mavronicolas. 2014. Complexity of rational and irrational Nash equilibria. *Theory of Computing Systems* 54, 3 (2014), 491–527.

[11] Charles Blackorby, Daniel Primont, and R. Robert Russell. 1978. *Duality, Separability and Functional Structure: Theory and Economic Applications*. New York: North-Holland.

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

[13] 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.

[14] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. 2016. Introduction to Computational Social Choice. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, J. Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, Chapter 1.

[15] Felix Brandt, Matthias Greger, Erel Segal-Halevi, and Warut Suksompong. 2023. Balanced Donor Coordination. In *Proceedings of the 24th ACM Conference on Economics and Computation (ACM-EC)*. 299.

[16] Felix Brandt, Matthias Greger, Erel Segal-Halevi, and Warut Suksompong. 2024. Optimal Budget Aggregation with Single-Peaked Preferences. In *Proceedings of the 25th ACM Conference on Economics and Computation (ACM-EC)*. 49.

[17] Felix Brandt, Matthias Greger, Erel Segal-Halevi, and Warut Suksompong. 2025. Coordinating Charitable Donations with Leontief Preferences. *Journal of Economic Theory* 230 (2025), 106096.

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

[19] Ioannis Caragiannis, George Christodoulou, and Nicos Protopapas. 2024. Truthful aggregation of budget proposals with proportionality guarantees. *Artificial Intelligence* 335 (2024), 104178.

[20] Ioannis Caragiannis, Angelo Fanelli, Nick Gravin, and Alexander Skopalik. 2015. Approximate pure Nash equilibria in weighted congestion games: Existence, efficient computation, and structure. *ACM Transactions on Economics and Computation* 3, 1 (2015), 2:1–2:32.

[21] Xi Chen and Xiaotie Deng. 2006. Settling the complexity of two-player Nash equilibrium. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 261–272.

[22] Xi Chen, Xiaotie Deng, and Shang-hua Teng. 2006. Computing Nash equilibria: Approximation and smoothed complexity. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 603–612.

[23] Yu Cheng, Ilias Diakonikolas, and Alistair Stewart. 2017. Playing anonymous games using simple strategies. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 616–631.

[24] Steve Chien and Alistair Sinclair. 2011. Convergence to approximate Nash equilibria in congestion games. *Games and Economic Behavior* 71, 2 (2011), 315–327.

[25] Bruno Codenotti, Mauro Leoncini, and Giovanni Resta. 2006. Efficient computation of Nash equilibria for very sparse win-lose bimatrix games. In *Proceedings of the 14th Conference on Annual European Symposium (ESA)*. 232–243.

[26] Roberto Cominetti, José R Correa, and Nicolás E Stier-Moses. 2009. The impact of oligopolistic competition in networks. *Operations Research* 57, 6 (2009), 1421–1437.

[27] R. Cornes and R. Harley. 2012. Fully Aggregative Games. *Economics Letters* 116 (2012), 631–633.

[28] Rachel Cummings, Michael Kearns, Aaron Roth, and Zhiwei Steven Wu. 2015. Privacy and Truthful Equilibrium Selection for Aggregative Games. In *Proceedings of the 11th International Conference on Web and Internet Economics (WINE)*. 286–299.

[29] Constantinos Daskalakis and Christos Papadimitriou. 2007. Computing equilibria in anonymous games. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 83–93.

[30] Constantinos Daskalakis and Christos H. Papadimitriou. 2015. Approximate Nash equilibria in anonymous games. *Journal of Economic Theory* 156 (2015), 207–245.

[31] Mark de Berg, Rupert Freeman, Ulrike Schmidt-Kraepelin, and Markus Utke. 2024. *Truthful Budget Aggregation: Beyond Moving-Phantom Mechanisms*. Technical Report. https://arxiv.org/pdf/2405.20303.

[32] Gérard Debreu. 1952. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences (PNAS)* 38, 10 (1952), 886–893.

[33] G. Debreu. 1959. *Topological Methods in Cardinal Utility Theory*. Cowles Foundation Discussion Paper 76. Cowles Foundation at Yale University.

[34] Argyrios Deligkas, John Fearnley, Rahul Savani, and Paul Spirakis. 2017. Computing approximate Nash equilibria in polymatrix games. *Algorithmica* 77, 2 (2017), 487–514.

[35] Pradeep Dubey, Ori Haimanko, and Andriy Zapechelnyuk. 2006. Strategic complements and substitutes, and potential games. *Games and Economic Behavior* 54, 1 (2006), 77–94.

[36] Edith Elkind, Leslie Ann Goldberg, and Paul Goldberg. 2006. Nash equilibria in graphical games on trees revisited. In *Proceedings of the 7th ACM Conference on Electronic Commerce (ACM-EC)*. 100–109.

[37] Edith Elkind, Matthias Greger, Patrick Lederer, Warut Suksompong, and Nicholas Teh. 2026. Settling the score: Portioning with cardinal preferences. *Artificial Intelligence* 352 (2026), 104487.

[38] Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. 2004. The complexity of pure Nash equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*. 604–612.

[39] Rupert Freeman, D. M. Pennock, Dominik Peters, and J. Wortman Vaughan. 2019. Truthful aggregation of budget proposals. In *Proceedings of the 2019 ACM Conference on Economics and Computation (ACM-EC)*. 751–752.

[40] Rupert Freeman, David M. Pennock, Dominik Peters, and Jennifer Wortman Vaughan. 2021. Truthful aggregation of budget proposals. *Journal of Economic Theory* 193 (2021), 105234.

[41] Ashish Goel, Anilesh K. Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. 2019. Knapsack Voting for Participatory Budgeting. *ACM Transactions on Economics and Computation* 7, 2 (2019), 8:1–8:27.

[42] Adrian Haret, Sophie Klumper, Jan Maly, and Guido Schäfer. 2024. Committees and Equilibria: Multiwinner Approval Voting Through the Lens of Budgeting Games. In *Proceedings of the 25th ACM Conference on Economics and Computation (ACM-EC)*. 51–70.

[43] Tobias Harks and Veerle Timmermans. 2022. Equilibrium computation in resource allocation games. *Mathematical Programming* 194, 1 (2022), 1–34.

[44] Chien-Chung Huang. 2013. Collusion in atomic splittable routing games. *Theory of Computing Systems* 52, 4 (2013), 763–801.

[45] Kamal Jain. 2007. A Polynomial Time Algorithm for Computing an Arrow–Debreu Market Equilibrium for Linear Utilities. *SIAM J. Comput.* 37, 1 (2007), 303–318.

[46] S. Janson. 2016. *Phragmén's and Thiele's election methods*. Technical Report arXiv:1611.08826 [math.HO]. arXiv.org.

[47] M.K. Jensen. 2010. Aggregative Games and Best-Reply Potentials. *Economic Theory* 43 (2010), 45–66.

[48] Martin Jensen. 2018. Aggregative games. In *Handbook of Game Theory and Industrial Organization, Volume I*. Edward Elgar Publishing, 66–92.

[49] Michael Kearns, Michael L. Littman, and Satinder Singh. 2001. Graphical models for game theory. *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence* (2001), 253–260.

[50] Michael Kearns and Yishay Mansour. 2002. Efficient Nash computation in large population games with bounded influence. In *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence*. 259–266.

[51] Max Klimm and Philipp Warode. 2025. Complexity and Parametric Computation of Equilibria in Atomic Splittable Congestion Games via Weighted Block Laplacians. *SIAM J. Comput.* 54, 5 (2025), 1241–1293.

[52] Nikolai S. Kukushkin. 1994. A fixed-point theorem for decreasing mappings. *Economics Letters* 46, 1 (1994), 23–26.

[53] Martin Lackner and Piotr Skowron. 2023. *Multi-winner voting with approval preferences*. Springer Nature.

[54] Carlton E. Lemke and J. T. Howson. 1964. Equilibrium points of bimatrix games. *SIAM J. Appl. Math.* 12, 2 (1964), 413–423.

[55] Huaqing Li, Jun Li, Liang Ran, Lifeng Zheng, and Tingwen Huang. 2025. A Survey of Distributed Algorithms for Aggregative Games. *IEEE/CAA Journal of Automatica Sinica* 12, 5 (2025), 859–871.

[56] Tobias Lindner. 2011. *Zur Manipulierbarkeit der Allokation öffentlicher Güter: Theoretische Analyse und Simulationsergebnisse*. Ph.D. Dissertation. Karlsruhe Institute of Technology.

[57] Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. 2003. Playing large games using simple strategies. In *Proceedings of the 4th ACM conference on Electronic Commerce*. 36–41.

[58] Lionel W. McKenzie. 1959. On the Existence of General Equilibrium for a Competitive Market. *Econometrica* 27, 1 (1959), 54–71.

[59] John F. Nash. 1950. Equilibrium Points in n-Person Games. *Proceedings of the National Academy of Sciences (PNAS)* 36 (1950), 48–49.

[60] John F. Nash. 1951. Non-cooperative games. *Annals of Mathematics* 54, 2 (1951), 286–295.

[61] Ariel Orda, Raphael Rom, and Nahum Shimkin. 2002. Competitive routing in multiuser communication networks. *IEEE/ACM Transactions on networking* 1, 5 (2002), 510–521.

[62] Edan Orzech and Martin Rinard. 2025. *Nash equilibria with irradical probabilities*. Technical Report. https://arxiv.org/abs/2507.09422.

[63] Christos Papadimitriou, Emmanouil-Vasileios Vlatakis-Gkaragkounis, and Manolis Zampetakis. 2023. The Computational Complexity of Multi-player Concave Games and Kakutani Fixed Points. In *Proceedings of the 24th ACM Conference on Economics and Computation*. 1045–1045.

[64] Francesca Parise, Sergio Grammatico, Basilio Gentile, and John Lygeros. 2020. Distributed convergence to Nash equilibria in network and average aggregative games. *Automatica* 117 (2020), 108959.

[65] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. 2021. Proportional Participatory Budgeting with Additive Utilities. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*. 12726–12737.

[66] Dominik Peters and Piotr Skowron. 2020. Proportionality and the Limits of Welfarism. In *Proceedings of the 21nd ACM Conference on Economics and Computation (ACM-EC)*. 793–794.

[67] Edvard Phragmén. 1899. Till frågan om en proportionell valmetod. *Statsvetenskaplig Tidskrift* 2, 2 (1899), 297–305.

[68] Ryan Porter, Eugene Nudelman, and Yoav Shoham. 2008. Simple search methods for finding a Nash equilibrium. *Games and Economic Behavior* 63, 2 (2008), 642–662.

[69] J. B. Rosen. 1965. Existence and uniqueness of equilibrium points for concave N-person games. *Econometrica* 33, 3 (1965), 520–534.

[70] Robert W. Rosenthal. 1973. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory* 2, 1 (1973), 65–67.

[71] Tim Roughgarden and Florian Schoppmann. 2015. Local smoothness and the price of anarchy in splittable congestion games. *Journal of Economic Theory* 156 (2015), 317–342.

[72] Aviad Rubinstein. 2015. Inapproximability of Nash equilibrium. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*. 409–418.

[73] Haralampos Tsaknakis and Paul G. Spirakis. 2007. An optimization approach for approximate Nash equilibria. In *Internet and Network Economics*. Springer, 42–56.

[74] John von Neumann. 1928. Zur Theorie der Gesellschaftsspiele. *Math. Ann.* 100, 1 (1928), 295–320.

[75] John von Neumann and Oskar Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.

[76] Yongqiang Wang and Angelia Nedić. 2024. Differentially Private Distributed Algorithms for Aggregative Games With Guaranteed Convergence. *IEEE Trans. Automat. Control* 69, 8 (2024), 5168–5183.