



PDF Download  
3736252.3742644.pdf  
11 January 2026  
Total Citations: 0  
Total Downloads: 353

 Latest updates: <https://dl.acm.org/doi/10.1145/3736252.3742644>

RESEARCH-ARTICLE

## Near-feasible Fair Allocations in Two-sided Markets

**JAVIER CEMBRANO**, Max Planck Institute for Informatics, Saarbrücken, Saarland, Germany

**ANDRÉS MORAGA**, Pontifical Catholic University of Chile, Santiago, RM, Chile

**VÍCTOR VERDUGO**, Pontifical Catholic University of Chile, Santiago, RM, Chile

Open Access Support provided by:

Max Planck Institute for Informatics

Pontifical Catholic University of Chile

Published: 07 July 2025

[Citation in BibTeX format](#)

EC '25: 26th ACM Conference on Economics and Computation

July 7 - 10, 2025  
CA, Stanford, USA

Conference Sponsors:  
SIGECOM

# Near-feasible Fair Allocations in Two-sided Markets\*

JAVIER CEMBRANO, Max Planck Institut für Informatik, Germany

ANDRÉS MORAGA, Pontificia Universidad Católica de Chile, Chile

VICTOR VERDUGO, Pontificia Universidad Católica de Chile, Chile

We study resource allocation in two-sided markets from a fundamental perspective and introduce a general modeling and algorithmic framework to effectively incorporate the complex and multidimensional aspects of fairness. Our main technical contribution is to show the existence of a range of near-feasible resource allocations parameterized in different model primitives to give flexibility when balancing the different policymaking requirements, allowing policy designers to fix these values according to the specific application. To construct our near-feasible allocations, we start from a fractional resource allocation and perform an iterative rounding procedure to get an integer allocation. We show a simple yet flexible and strong sufficient condition for the target feasibility deviations to guarantee that the rounding procedure succeeds, exhibiting the underlying trade-offs between market capacities, agents' demand, and fairness. To showcase our framework's modeling and algorithmic capabilities, we consider three prominent market design problems: school allocation, stable matching with couples, and political apportionment. In each of them, we obtain strengthened guarantees on the existence of near-feasible allocations capturing the corresponding fairness notions, such as proportionality, envy-freeness, and stability.

CCS Concepts: • **Theory of computation** → **Rounding techniques; Algorithmic mechanism design; Mathematics of computing** → *Approximation algorithms*.

Additional Key Words and Phrases: Fair Allocation, Rounding Algorithms

## ACM Reference Format:

Javier Cembrano, Andrés Moraga, and Victor Verdugo. 2025. Near-feasible Fair Allocations in Two-sided Markets. In *The 26th ACM Conference on Economics and Computation (EC '25), July 7–10, 2025, Stanford, CA, USA*. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3736252.3742644>

## 1 Introduction

Resource allocation is a fundamental task that lies at the intersection of economics, computer science, and operations, where the objective is to develop efficient policies to distribute scarce resources among various agents or entities. This naturally results in a two-sided market: One side consists of the agents, while the other side encompasses the different resources that need to be allocated, all while adhering to market capacities and demands. Many real-world applications involve indivisible resources, a large number of entities, and combinatorial constraints, making the design of efficient resource allocations increasingly complex from both a modeling and computational perspectives. This two-sided paradigm applies to various situations, including the allocation of students to schools, the composition of political representative bodies, the allocation of doctors to hospitals, organ

\*A full version of this paper can be found on arXiv:2506.01178.

Authors' Contact Information: Javier Cembrano, [jcembran@mpi-inf.mpg.de](mailto:jcembran@mpi-inf.mpg.de), Max Planck Institut für Informatik, Saarbrücken, Germany; Andrés Moraga, [andresmoraga0@uc.cl](mailto:andresmoraga0@uc.cl), Pontificia Universidad Católica de Chile, Santiago, Chile; Victor Verdugo, [victor.verdugo@uc.cl](mailto:victor.verdugo@uc.cl), Pontificia Universidad Católica de Chile, Santiago, Chile.



This work is licensed under a Creative Commons Attribution 4.0 International License.

*EC '25, July 7–10, 2025, Stanford, CA, USA*

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1943-1/2025/07

<https://doi.org/10.1145/3736252.3742644>

donation systems, job markets, and online marketplaces, among many others; see, e.g., Balinski and Young [2010], Haeringer [2018], Immorlica et al. [2023], Roth [2018].

While efficiency is a key goal in market design, ensuring fairness in resource allocation is arguably one of the most critical factors in assessing the quality of a policy. It is essential to evaluate how effectively the policy treats different groups within a population, ensuring that resources, opportunities, and outcomes are distributed equitably. Historical inequalities and biases have often led to disparities among various social, economic, or demographic dimensions, prompting policymakers to enhance policies with proven fairness and efficiency guarantees to rectify these imbalances. Typically, these efforts rely on market-specific characteristics to design effective policies, taking advantage of the rich structure of the problem space. However, many existing models and algorithms lack the robustness to address new fairness considerations.

Computational challenges stem from multiple sources. On one hand, complex combinatorial constraints result in NP-hard problems, requiring a trade-off between efficiency and computational effort. Additionally, incorporating fairness requirements across the different dimensions not only adds an extra layer of computational difficulty but also often leads to infeasible scenarios due to deep existing impossibilities in reconciling efficiency with fairness. To navigate these challenges, implementing *near-feasible* solutions, i.e., solutions that may slightly violate some constraints of the problem, is an effective way to overcome these policymaking difficulties. From an optimization standpoint, the goal is to provide provable guarantees regarding how close these solutions are to the set of feasible policies.

### 1.1 Our Contribution and Techniques

In this work, we study resource allocation in two-sided markets from a fundamental perspective and introduce a general modeling and algorithmic framework to effectively incorporate the complex and multidimensional aspects of fairness. Within this framework, we provide a rounding theorem to compute near-feasible allocations with strong and flexible approximation guarantees. In the following, we summarize our contributions and discuss the implications of our results.

**A Multidimensional Resource Allocation Framework.** Our first contribution is the introduction of a general framework to model resource allocation problems in two-sided markets. In this model, there is a set of agents and indivisible resources with finite capacities to be allocated across the agents. Multiple dimensions characterize the set of agents; each agent may belong to a group within each dimension. This model feature allows for a multidimensional representation of the agents, e.g., age, gender, ethnicity, and socio-economic or demographic aspects. The market is provided with target utilities, depending on the dimensions and the groups, that capture the different fairness requirements. Then, the feasible points of an integer program define the set of feasible resource allocations, though this set may be empty in some cases. In our near-feasible resource allocations, the utilities for each group are approximated within a given factor while requiring an extra amount of resources; however, this additional amount is regulated in two ways: by bounding the maximum extra capacity per resource and controlling the total additional market capacity. In Section 2, we provide the formal definition of our multidimensional model, its instances, and our notion of near-feasible resource allocations.

**A Rounding Theorem with Flexible Guarantees.** In our multidimensional resource allocation model, the set of feasible resource allocations is represented by the feasible points of an integer program. From a fundamental point of view, the instances of our multidimensional resource allocation problem can be lifted to get a corresponding (weighted) hypergraph encoding the agents, resource bundles, dimensions, and groups. Naturally, strict requirements on the allocation quality

might yield infeasible regions for the corresponding integer program. Our main technical contribution establishes the existence of near-feasible resource allocations, with adjustable approximation guarantees. By tuning model parameters, policymakers can balance and enforce deviation bounds tailored to specific applications.

To construct our near-feasible allocations, we start from a fractional solution of the natural linear relaxation and apply an iterative rounding procedure to derive a feasible solution for the integer program. We show a simple yet flexible and strong sufficient condition for the target feasibility deviations to guarantee that the iterative rounding procedure computes a near-feasible resource allocation (Theorem 3.1). This condition highlights the underlying trade-offs among market capacities, agents' demand, and utility distribution across dimensions. In Section 3, we present the formal statement of our main theorem, as well as the rounding algorithm and its analysis.

**Computing Near-feasible Fair Allocations.** To showcase our framework's modeling and algorithmic capabilities, we consider three prominent market design problems: school allocation, stable matching with couples, and political apportionment. In each of them, ruled by the sufficient condition of our main rounding theorem, we obtain a range of strengthened guarantees on the existence of near-feasible allocations under general fairness requirements notions, including the proportionality and envy-freeness objectives.

*School allocation.* In Section 4.1, we propose a general multidimensional allocation model which, in particular, captures the school allocation setting introduced by Procaccia et al. [2024] as single-dimensional instances. We model group fairness by a convex optimization-driven approach, and using our rounding Theorem 3.1 starting from a feasible fractional solution, we get a general guarantee on the existence of near-feasible multidimensional resource allocations (Theorem 4.1). For the case of proportional fairness, as a corollary, we can directly accommodate the policy-maker priorities by trading off the utility approximation and maximum resource capacity augmentation to get slight constant deviations, which differentiates our result from previous single-dimensional work on group utilities and total allocation excess [A. et al., 2024, Procaccia et al., 2024].

Then, we introduce a new multidimensional notion of group envy-freeness where the ratio between the total utility of a group for its allocation and its total utility for any other group's allocation should not be smaller than the ratio between the sizes of these groups. For this natural notion, we provide near-feasible resource allocations with approximation guarantees that do not depend on the number of agents, thus breaking the impossibility found by Procaccia et al. [2024] for their more stringent single-dimensional envy-freeness notion (Theorem 4.3).

*Stable matching with couples.* While stable matchings are guaranteed to exist in the basic single-demand setting [Gale and Shapley, 1962], this is not the case for pairs and, more generally, multi-demand settings. To illustrate how our framework can also accommodate stability requirements, we show in Section 4.2 how by using our rounding Theorem 3.1, we can directly recover a recent guarantee by Nguyen and Vohra [2018] for the existence of near-feasible stable allocations. We further extend this setting to handle group fairness requirements and allocation stability to guarantee the existence of near-feasible, stable, and fair allocations (Theorem 4.4).

Our result allows us to incorporate both stability and fairness while keeping the deviations from resource capacities bounded by small constants. For example, in the single-dimensional case, one can guarantee deviations from group fairness of at most five by increasing the resource capacities by at most four, while keeping the deviation from the total aggregate capacity at two. Our new near-feasibility results in Section 4.2 concern the interaction between stability and fairness, thereby contributing to the efforts on designing two-sided markets in a multidimensional environment

under more involved stability concepts. We believe our algorithmic framework will help provide near-feasible allocations in further stable matching settings under complex fairness requirements.

*Political apportionment.* In the multidimensional apportionment problem, introduced by Balinski and Demange [1989a,b] for the case of two dimensions and extended by Cembrano et al. [2022] to an arbitrary number of dimensions, the goal is to allocate the seats of a representative body proportionally across several dimensions. Classic apportionment methods, e.g., divisor methods or Hamilton’s method, aim to assign seats across districts proportionally to their population or across parties proportionally to their electoral support [Balinski and Young, 2010, Pukelsheim, 2017].

However, in addition to political parties and geographical divisions, natural dimensions include demographics of the elected members such as gender or ethnicity; see, e.g., Arnosti et al. [2024], Cembrano et al. [2024], Mathieu and Verdugo [2024]. In Section 4.3, using our rounding theorem, we improve over the result by Cembrano et al. [2022] to get enhanced near-feasibility guarantees for multidimensional apportionment (Theorem 4.5). Remarkably, we can further bound the total deviation from the house size, while the rounding algorithm by Cembrano et al. only controls the deviations on each dimension.

## 1.2 Further Related Work

From an algorithmic perspective, our approach to finding near-feasible allocations closely relates to the classic discrepancy minimization problem. In its basic form, there is a fractional vector  $x$  with entries in  $[0, 1]$  such that  $Ax = b$  for some binary matrix  $A$  and an integer vector  $b$ ; the goal is to find a binary rounding  $y$  minimizing the maximum additive deviation  $\|Ax - Ay\|_\infty$ . Since the seminal iterative-rounding work by Beck and Fiala [1981], the problem has been extensively studied, including different norms to measure deviations, probabilistic guarantees, general combinatorial constraints, and online variants [Bansal, 2010, Bansal et al., 2019, Bansal and Garg, 2017, Bansal et al., 2020, Bukh, 2016, Lovett and Meka, 2015, Rothvoss, 2017]. Iterative rounding has been a successful tool for designing discrepancy and near-feasibility algorithms, and we refer to the book by Lau et al. [2011] for a primer on classic applications.

Very recently, there has been a series of works on near-feasible stable allocations; related to our work is the one by Nguyen et al. [2021] on group-stability and Nguyen and Vohra [2019] on proportionality, where they provide the existence of near-feasible allocations via rounding methods; in Section 5, we briefly discuss the differences between our general fairness approach and the aforementioned works. In a related line, there are recent works about optimal capacity design for school matching [Bobbio et al., 2023], refugee settlement [Ahani et al., 2021, Andersson and Ehlers, 2020, Delacrétaz et al., 2016] and healthcare rationing [Aziz and Brandl, 2021, Pathak et al., 2021]. In general, the study of fairness in allocation problems has a rich history, and the concepts of proportionality and envy-freeness have been key design objectives; we refer to the book by Moulin [2004] for an extensive treatment of the fair division theory.

## 2 Multidimensional Capacitated Resource Allocation

We let  $\mathbb{R}_+$  (resp.  $\mathbb{R}_{++}$ ) denote the non-negative (resp. strictly positive) reals,  $\mathbb{N}$  denote the strictly positive integers, and  $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ . We write  $[n]$  as a shortcut for  $\{1, 2, \dots, n\}$  and  $[n]_0$  as a shortcut for  $\{0, 1, \dots, n\}$ . In the *multidimensional capacitated resource allocation* problem, or MCRA for short, an instance is structured in the following way:

- (I) **Agents, resources, and groups.** We have a set  $A$  of *entities* or *agents*, a set  $R$  of *resources*, and a set  $A' \subseteq A$  of *binding agents*, i.e., agents that *must* be allocated a resource. The agents are organized in groups according to  $d$  *dimensions*, namely, for each dimension  $\ell \in [d]$  we have  $k_\ell$

different groups  $G_{\ell,1}, \dots, G_{\ell,k_\ell} \subseteq A$  of agents such that for every  $i, j \in [k_\ell]$  with  $i \neq j$ , we have  $G_{\ell,i} \cap G_{\ell,j} = \emptyset$ , i.e., an agent can be part of at most one group on each dimension.

- (II) **Demands and capacities.** Each agent  $a \in A$  has a *demand*  $\omega_a \in \mathbb{N}$  for the number of resources that should receive, and an  $\omega_a$ -*bundle* is any function  $q: R \rightarrow \mathbb{N}_0$  such that  $\sum_{r \in R} q(r) = \omega_a$ . We denote by  $\mathcal{T}_a$  the set of  $\omega_a$ -bundles, and we write  $\omega^*$  for the maximum demand of an agent, i.e.,  $\max_{a \in A} \omega_a$ . We have a resource *capacity function*  $c: R \rightarrow \mathbb{N}$ .

An instance of MCRA is determined by a tuple  $\mathcal{I} = (A, A', R, G, \omega, c)$  ruled by (I)-(II). Given an instance of MCRA, we let  $\mathcal{E} = \{(a, q) : a \in A, q \in \mathcal{T}_a\}$  denote the feasible agent-bundle pairs. A *resource allocation* is a mapping  $x: \mathcal{E} \rightarrow \{0, 1\}$  that satisfies the following conditions:

$$\sum_{q \in \mathcal{T}_a} x(a, q) = 1 \quad \text{for every } a \in A', \quad (1)$$

$$\sum_{q \in \mathcal{T}_a} x(a, q) \leq 1 \quad \text{for every } a \in A \setminus A', \quad (2)$$

$$\sum_{a \in A} \sum_{q \in \mathcal{T}_a} q(r) \cdot x(a, q) \leq c(r) \quad \text{for every } r \in R. \quad (3)$$

The set of constraints (1) ensures that each agent in  $A'$  is allocated exactly one bundle, whereas (2) ensures that every other agent is allocated at most one bundle. The set of constraints (3) ensures that at most  $c(r)$  units of resource  $r$  are allocated. We say that  $x$  is a *fractional resource allocation* when it satisfies (1)-(3) and  $x(a, q) \in [0, 1]$  for each  $(a, q) \in \mathcal{E}$ , i.e., integrality is relaxed. For a fractional resource allocation  $x$ , we let  $\mathcal{A}(x) = \{a \in A : |\{q \in \mathcal{T}_a : x(a, q) > 0\}| \geq 2\}$  denote the agents with more than one bundle (fractionally) allocated.

## 2.1 Utilities and Near-feasible Allocations

Throughout this work, we consider agent-dependent *utility functions* suitable for our applications. For a collection of utility functions  $u_a: \mathcal{T}_a \rightarrow \mathbb{R}_+$  for each  $a \in A$ , and a (integral or fractional) mapping  $x: \mathcal{E} \rightarrow [0, 1]$ , we let  $U_{\ell,i}(x) = \sum_{a \in G_{\ell,i}} \sum_{q \in \mathcal{T}_a} u_a(q) \cdot x(a, q)$  denote the total utility of group  $G_{\ell,i}$ , for each  $\ell \in [d]$  and  $i \in [k_\ell]$ . We further let  $U_{\ell,i}^* = \max\{u_a(q) : a \in G_{\ell,i} \text{ and } q \in \mathcal{T}_a\}$  denote the maximum utility of an agent in  $G_{\ell,i}$ . For notation simplicity, we denote the collection of utility functions  $(u_a)_{a \in A}$  by  $u$ .

**Near-feasible allocations.** Let  $\alpha \in \mathbb{N}_0^d$  and  $\delta, \Delta \in \mathbb{N}_0$ . Given an instance  $\mathcal{I} = (A, A', R, G, \omega, c)$  of MCRA, a fractional resource allocation  $x$ , and a collection of utility functions  $u$ , a mapping  $y: \mathcal{E} \rightarrow \{0, 1\}$  is an  $(\alpha, \delta, \Delta)$ -*approximation of  $x$  with respect to  $u$*  if it satisfies (1)-(2) and

$$|U_{\ell,i}(y) - U_{\ell,i}(x)| < \alpha_\ell \cdot U_{\ell,i}^* \quad \text{for every } \ell \in [d], i \in [k_\ell], \quad (4)$$

$$|\sum_{a \in A} \sum_{q \in \mathcal{T}_a} q(r) \cdot (y(a, q) - x(a, q))| < \delta \quad \text{for every } r \in R, \quad (5)$$

$$|\sum_{a \in A} \sum_{q \in \mathcal{T}_a} \omega_a \cdot (y(a, q) - x(a, q))| < \omega^* \cdot \Delta. \quad (6)$$

In such near-feasible allocation, the utility of each group  $G_{\ell,i}$  deviates strictly less than  $\alpha_\ell \cdot U_{\ell,i}^*$  from the utilities in  $x$ , the deviation on the assigned agents to each resource are strictly less than  $\delta$ , and the deviation on the total number of allocated resources is strictly less than  $\omega^* \cdot \Delta$ . As we will illustrate in Section 4, in some applications the specific structure of the problem guarantees that some of the left-hand sides of these inequalities are integer values, so the maximum possible deviations become  $\alpha_\ell \cdot U_{\ell,i}^* - 1$ ,  $\delta - 1$ , or  $\omega^* \cdot \Delta - 1$ , respectively. Finally, we say that a mapping  $y: \mathcal{E} \rightarrow \{0, 1\}$  is a *rounding* of a fractional resource allocation  $x$  if  $y(a, q) = 0$  whenever  $x(a, q) = 0$  and  $y(a, q) = 1$  whenever  $x(a, q) = 1$ .

## 3 A Rounding Theorem

In this section, we present our main technical result. On an intuitive level, the idea is to start from a fractional resource allocation satisfying a certain fairness notion, which is known to be kept

upon rounding. Our result then states the existence of a resource allocation obtained by rounding this fractional allocation that guarantees small deviations from the resource capacities and from arbitrarily defined utilities. While having small deviations from capacities constitutes a natural goal in this setting, the definition of the group utilities will allow our model to capture fairness in different applications. Remarkably, our result gives a high degree of flexibility for choosing the maximum deviations, allowing policy designers to fix these values according to the specific application.

**THEOREM 3.1.** *Let  $\mathcal{I} = (A, A', R, G, \omega, c)$  be an instance of MCRA, let  $x$  be a fractional resource allocation for this instance, and let  $u_a: \mathcal{T}_a \rightarrow \mathbb{R}_+$  be a utility function for each agent  $a \in A$ . Fix  $\psi = 1$  if  $\mathcal{A}(x) \neq \emptyset$  or  $d \leq 1$ , and  $\psi \in \{0, 1\}$  arbitrarily otherwise. Let  $\alpha \in \mathbb{N}_0^d$  and  $\delta \in \mathbb{N}_0$  be such that*

$$\frac{\mathbb{1}_{\psi=1}}{2} + \sum_{\ell \in [d]} \frac{1}{\alpha_\ell + 1} + \frac{\omega^*}{\delta + 1} \leq 1 \tag{7}$$

and let  $\Delta \in \mathbb{N}_0$  be such that one of the following holds:

- (i)  $\Delta \geq 2$  and  $\psi = 1$ ; or
- (ii) inequality (7) holds strictly and  $\Delta \geq 1 / (1 - (\frac{\mathbb{1}_{\psi=1}}{2} + \sum_{\ell \in [d]} \frac{1}{\alpha_\ell + 1} + \frac{\omega^*}{\delta + 1})) - 1$ .

Then, there exists a rounding  $y$  of  $x$  that is an  $(\alpha, \delta, \Delta)$ -approximation of  $x$  with respect to  $u$ . Furthermore,  $y$  can be found in time polynomial in  $|A|, k_1, \dots, k_d$ , and  $|R|^{\omega^*}$ .

The parameter  $\psi$  in the theorem captures whether we impose or not, in a certain setting, constraints over the fractional allocation associated with a single agent. If some agent has two or more associated fractional variables under allocation  $x$ , i.e.,  $\mathcal{A}(x) \neq \emptyset$ , this has to be the case to ensure that a fractional allocation is produced; otherwise, we have the freedom to impose these constraints and bound the total deviation  $\Delta$  or not impose them and get smaller violations for groups and individual resources.

To prove this theorem, we construct an iterative rounding algorithm that starts from a fractional resource allocation and iteratively rounds its components while ensuring that no constraint is violated by too much. We first introduce some notation we use for its description and analysis. For an instance  $\mathcal{I} = (A, A', R, G, \omega, c)$ , a subset  $\mathcal{F} \subseteq \mathcal{E}$ , and a mapping  $x: \mathcal{F} \rightarrow (0, 1)$ , we write

$$\begin{aligned} \mathcal{F}(a) &= \{(a', q) \in \mathcal{F} : a' = a\} && \text{for every } a \in A, \\ \mathcal{F}(\ell, i) &= \{(a, q) \in \mathcal{F} : a \in G_{\ell, i}\} && \text{for every } \ell \in [d], i \in [k_\ell], \\ \mathcal{F}(r) &= \{(a, q) \in \mathcal{F} : q(r) \geq 1\} && \text{for every } r \in R. \end{aligned}$$

These concepts allow for a natural hypergraph interpretation, where the vertices are all agents, resources, and groups, and each tuple  $(a, q) \in \mathcal{F}$  is associated with a hyperedge comprising agent  $a$ , all groups to which  $a$  belongs, and all resources in  $q$ . With this point of view, the sets  $\mathcal{F}(a)$ ,  $\mathcal{F}(\ell, i)$ , and  $\mathcal{F}(r)$  correspond to the incident hyperedges of a given vertex. Although this is not a proper hypergraph as a bundle may contain several copies of a resource, we sometimes use this interpretation for informal explanations and discussions. We often refer to a pair  $(a, q) \in \mathcal{F}$  simply as  $e$ , in particular, whenever specification about  $a$  or  $q$  is not needed. In slight abuse of notation, we write  $u(e)$  instead of  $u_a(q)$  when  $e = (a, q)$ . We also write  $A(\mathcal{F}) = \{a \in A : \mathcal{F}(a) \geq 1\}$ ,  $R(\mathcal{F}) = \{r \in R : \mathcal{F}(r) \geq 1\}$ , and  $G_\ell(\mathcal{F}) = \{i \in [k_\ell] : \mathcal{F}(\ell, i) \geq 1\}$  for each  $\ell \in [d]$ , to refer to the agents, resources, and groups in the  $\ell$ th dimension with incident hyperedges.

We now introduce the linear program solved by our algorithm in each iteration. We let  $\mathcal{I} = (A, A', R, G, \omega, c)$  be an instance,  $u$  be a collection of utility functions,  $\tilde{x}: \mathcal{F} \rightarrow (0, 1)$  be a mapping on  $\mathcal{F} \subseteq \mathcal{E}$ ,  $\tilde{A} \subseteq A$  be a subset of binding agents,  $\tilde{G} \subseteq \{(\ell, i) : \ell \in [d], i \in [k_\ell]\}$  be a subset of

binding groups,  $\tilde{R} \subseteq R$  be a subset of binding resources, and  $\chi \in \{0, 1\}$  be a binary value. We consider the following linear program  $\text{LP}(\mathcal{I}, u, \tilde{x}, \tilde{A}, \tilde{G}, \tilde{R}, \chi)$  with variables  $y: \mathcal{F} \rightarrow [0, 1]$ :

$$\sum_{e \in \mathcal{F}(a)} y(e) = 1 \quad \text{for every } a \in \tilde{A}, \quad (8)$$

$$\sum_{e \in \mathcal{F}(a)} y(e) \leq 1 \quad \text{for every } a \in A(\mathcal{F}) \setminus \tilde{A}, \quad (9)$$

$$\sum_{e \in \mathcal{F}(\ell, i)} u(e) \cdot y(e) = \sum_{e \in \mathcal{F}(\ell, i)} u(e) \cdot \tilde{x}(e) \quad \text{for every } \ell \in [d], i \in \tilde{G}_\ell, \quad (10)$$

$$\sum_{(a, q) \in \mathcal{F}(r)} q(r) \cdot y(a, q) = \sum_{(a, q) \in \mathcal{F}(r)} q(r) \cdot \tilde{x}(a, q) \quad \text{for every } r \in \tilde{R}, \quad (11)$$

$$\sum_{(a, q) \in \mathcal{F}} \omega_a \cdot y(a, q) = \sum_{(a, q) \in \mathcal{F}} \omega_a \cdot \tilde{x}(a, q) \quad \text{if } \chi = 1, \quad (12)$$

$$0 \leq y(e) \leq 1 \quad \text{for every } e \in \mathcal{F}. \quad (13)$$

In each iteration of our algorithm,  $\mathcal{F}$  will represent the remaining tuples  $(a, q)$  with an associated fractional value  $\tilde{x}(a, q)$  in the allocation. Solving this linear program will then ensure to maintain a total allocation of 1 for all agents in  $\tilde{A}$  and at most one for the other agents, the same total utility as in  $\tilde{x}$  for every group  $G_{\ell, i}$  with  $i \in \tilde{G}_\ell$ , the same number of allocated agents for every resource in  $\tilde{R}$ , and the same total number of allocated agents when  $\chi = 1$ . The definitions of these sets of agents, groups, and resources in each iteration will be made to balance two simultaneous objectives: (i) that agents, groups, and resources with many associated fractional values do not deviate from their current allocation, and (ii) that there are more variables than linearly independent equality constraints so that the algorithm makes progress and eventually terminates. Similarly, the binary value  $\chi$  will be set to zero only when there are few fractional values, to ensure that the total deviation from the number of resources allocated initially is kept under control.

Our rounding algorithm, formally described in Algorithm 1, starts from a given fractional allocation  $x$  (which we also call  $x^0$ ) and iteratively solves  $\text{LP}(\mathcal{I}, u, \tilde{x}, \tilde{A}, \tilde{G}, \tilde{R}, \chi)$  for  $\tilde{x}$  being the fractional entries of  $x$ , each time only considering the necessary constraints to prevent deviations beyond the allowed ones. At step  $t$ , the algorithm starts from an allocation  $x^t$  and constructs an allocation  $x^{t+1}$  by leaving the integral entries of  $x^t$  unchanged and taking, for the fractional entries  $\tilde{x}^t$ , an extreme point of this linear program. This program considers, for instance, constraint (10) for each group  $G_{\ell, i}$  such that there are  $\alpha_\ell + 1$  or more fractional values  $\tilde{x}^t(a, q)$  for some  $a \in G_{\ell, i}$  and  $q \in \mathcal{T}_a$ ; as the other groups  $G_{\ell, i}$  have at most  $\alpha_\ell$  such fractional values, they will deviate less than  $\alpha_\ell \cdot U_{\ell, i}^*$  from their previous utility. Similarly, the program considers constraint (8) for all agents with two or more values of  $\tilde{x}^t(a, q)$  summing up to 1, it considers constraint (11) for all resources  $r \in R$  with  $\delta + 1$  or more units associated with fractional values of  $\tilde{x}^t$ , and it considers constraint (12) only if there are  $\Delta + 1$  or more tuples  $(a, q)$  with agent  $a$  not being considered for constraint (8). We will prove that this construction ensures having more variables than linearly independent equality constraints in each iteration. Thus, we either fix a new variable to an integer value or an inequality constraint becomes tight.

When analyzing the algorithm, we usually fix its input and directly refer to the objects constructed during its execution with the names defined therein. For an iteration  $t \in [T - 1]_0$ , we further let

$$C^t = |\tilde{A}^t| + \sum_{\ell \in [d]} |\tilde{G}_\ell^t| + |\tilde{R}^t| + \chi^t$$

denote the number of equality constraints of  $\text{LP}(\mathcal{I}, u, \tilde{x}^t, \tilde{A}^t, \tilde{G}^t, \tilde{R}^t, \chi^t)$ . The following lemma states the property that ensures that the algorithm terminates: The linear program solved in each iteration has more variables than linearly independent equality constraints.

**LEMMA 3.2.** *Let  $\mathcal{I} = (A, A', R, G, \omega, c)$  be an instance,  $x$  a fractional resource allocation for it,  $u$  a collection of utility functions, and  $\alpha, \delta, \Delta, \psi$  such that the conditions in the statement of Theorem 3.1 are*

**ALGORITHM 1:** Iterative rounding for MCRA

**Input:** instance  $\mathcal{I} = (A, A', R, G, \omega, c)$ , utility functions  $u$ , fractional resource allocation  $x$ , maximum deviations  $\alpha \in \mathbb{N}_0^d$  and  $\delta, \Delta \in \mathbb{N}_0$

**Output:** resource allocation  $y$

$x^0 \leftarrow x$ ;

$t \leftarrow 0$ ;

**while true do**

$\mathcal{F}^t \leftarrow \{(a, q) \in \mathcal{E} : 0 < x^t(a, q) < 1\}$ ;

$\tilde{x}^t(e) \leftarrow x^t(e)$  for every  $e \in \mathcal{F}^t$ ;

$\tilde{A}^t \leftarrow \{a \in \mathcal{A}(x^t) : \sum_{e \in \mathcal{F}^t(a)} x^t(e) = 1\}$ ;

$\tilde{G}_\ell^t \leftarrow \{(\ell, i) : i \in [k_\ell] \text{ s.t. } |\mathcal{F}^t(\ell, i)| \geq \alpha_\ell + 1\}$  for every  $\ell \in [d]$ ;

$\tilde{R}^t \leftarrow \{r \in R : \sum_{(a, q) \in \mathcal{F}^t(r)} q(r) \geq \delta + 1\}$ ;

$\chi^t \leftarrow 1$  if  $|A(\mathcal{F}^t) \setminus \tilde{A}^t| \geq \Delta + 1$ , **else**  $\chi^t \leftarrow 0$ ;

**if**  $\tilde{A}^t = \tilde{R}^t = \emptyset$ ,  $\tilde{G}_\ell^t = \emptyset$  for every  $\ell \in [d]$ , **and**  $\chi^t = 0$  **then**

**break**

**end**

$x^{t+1}(e) \leftarrow x^t(e)$  for every  $e \in \mathcal{E} \setminus \mathcal{F}^t$ ;

$y^* \leftarrow$  extreme point of  $\text{LP}(\mathcal{I}, u, \tilde{x}^t, \tilde{A}^t, \tilde{G}^t, \tilde{R}^t, \chi^t)$ ;

$x^{t+1}(e) \leftarrow y^*(e)$  for every  $e \in \mathcal{F}^t$ ;

$t \leftarrow t + 1$

**end**

$T \leftarrow t$ ;

$x^{T+1}(e) \leftarrow x^T(e)$  for every  $e \in \mathcal{E} \setminus \mathcal{F}^T$ ;

fix  $x^{T+1}(e) \in \{\lfloor x^T(e) \rfloor, \lceil x^T(e) \rceil\}$  for every  $e \in \mathcal{F}^T$  such that  $\sum_{e \in \mathcal{F}^T(a)} x^{T+1}(e) \leq 1$  for each  $a \in A(\mathcal{F}^T)$ ;

**return**  $x^{T+1}$

satisfied. Then, for every iteration  $t \in [T - 1]_0$  it holds that  $C^t \leq |\mathcal{F}^t|$ . Furthermore, if the inequality is tight, then the following properties hold:

- (i) If  $\psi = 1$ , then  $\tilde{A}^t = A(\mathcal{F}^t)$ ;
- (ii)  $\tilde{G}_\ell^t = G_\ell(\mathcal{F}^t)$  and  $\bigcup_{i \in G_\ell(\mathcal{F}^t)} G_{\ell, i} = A(\mathcal{F}^t)$  for every  $\ell \in [d]$ ;
- (iii)  $\tilde{R}^t = R(\mathcal{F}^t)$  and  $\sum_{r \in R(\mathcal{F}^t)} q(r) = \omega^*$  for every  $(a, q) \in \mathcal{F}^t$ .

To show Theorem 3.1, we prove that the outcome  $y$  of Algorithm 1 for any instance and fractional allocation satisfies all conditions stated in the theorem. That the algorithm terminates follows from Lemma 3.2, as we show that the properties stated therein when the inequality is tight contradict the linear independence of the set of equality constraints. The running time is obtained by bounding the size of the linear program solved in each step and observing that its size also constitutes an upper bound for the number of iterations of the algorithm. That  $y$  satisfies the claimed notion of approximation is the most demanding part of the proof, requiring an understanding of the number and structure of the fractional variables upon termination of the iterative rounding procedure. This proof of this theorem is deferred to the full version; in the rest of the paper, we showcase its capabilities through specific applications.

## 4 Applications

In this section, we show how to use our rounding Theorem 3.1 to get approximation guarantees for several resource allocation problems that fall in our MCRA setting. In Section 4.1, we provide novel approximation guarantees for MCRA instances under group fairness constraints by rounding an optimal solution of a convex program used to model fairness across groups. Then, in Section 4.2,

we show how our framework can also be combined with stability requirements to obtain approximation guarantees for near-feasible stable allocation under group fairness constraints. Finally, in Section 4.3, we provide enhanced guarantees for near-feasible allocations in multidimensional political apportionment.

#### 4.1 Group-fairness in Resource Allocation

We consider a fairly general resource allocation setting where every agent  $a$  has a certain demand of  $\omega_a$  resources and a certain utility for each bundle, and the goal is to assign exactly one bundle to each agent. Formally, an instance of the *assignment-MCRA* problem is a tuple  $\mathcal{I} = (A, R, E, G, \omega, c, u)$ , where  $A, R, G, \omega$ , and  $c$  are structured in the same way as in the MCRA, we have a utility function  $u_a: \mathcal{T}_a \rightarrow \mathbb{R}_+$  for each agent  $a \in A$ , and for every pair  $(a, r) \in E \subseteq A \times R$  we say that  $r$  is *acceptable* for  $a$ . We continue to denote  $\mathcal{E} = \{(a, q) : a \in A, q \in \mathcal{T}_a\}$  and  $\omega^* = \max_{a \in A} \omega_a$ . For a given instance and an agent  $a \in A$ , we let  $\mathcal{T}_{a,E}$  be the set of  $\omega_a$ -bundles  $q \in \mathcal{T}_a$  such that  $(a, r) \in E$  for every  $r \in R$  with  $q(r) \geq 1$ , i.e., bundles made of acceptable resources for  $a$ . To distinguish the notation from that of the previous section, we write  $\mathcal{M} = \{(a, q) : a \in A, q \in \mathcal{T}_{a,E}\}$  for the set of feasible agent-bundle pairs in this context. A mapping  $x$  on  $\mathcal{M}$  is a (fractional) resource allocation for  $\mathcal{I}$  if its natural extension  $x'$  on  $\mathcal{E}$ , where  $x'(a, q) = x(a, q)$  for  $(a, q) \in \mathcal{M}$  and  $x'(a, q) = 0$  for  $(a, q) \in \mathcal{E} \setminus \mathcal{M}$ , is a (fractional) resource allocation for the instance  $(A, A, R, G, \omega, c)$  of MCRA, i.e.,  $A' = A$  so every agent is allocated exactly one bundle. We call an instance of assignment-MCRA *fractionally feasible* if it admits at least one fractional resource allocation.

**Group fairness.** Similarly to A. et al. [2024], we model the fairness requirements across groups by following an optimization-driven approach; namely, our goal is to find a resource allocation that maximizes the sum of a certain objective function of each group's utility. Formally, given a non-decreasing concave function  $f: \mathbb{R}_+ \rightarrow \mathbb{R}_+$  and an instance  $\mathcal{I}$  of assignment-MCRA, consider the following maximization problem:

$$\max \left\{ \sum_{\ell \in [d]} \sum_{i \in [k_\ell]} f(U_{\ell,i}(x)) : x \text{ is a fractional resource allocation} \right\}. \quad [\text{Fair}]$$

Recall that  $U_{\ell,i}(x)$  corresponds to the utility of group  $G_{\ell,i}$ , and therefore, in [Fair], the goal is to find a fractional resource allocation that maximizes the total utility across the groups. The program [Fair] can capture a broad family of natural fairness notions by setting the appropriate function  $f$ , like the classic utilitarian objective with  $f(z) = z$  and the celebrated *proportionality* objective by using  $f(z) = \ln(z)$ ; see. e.g., Young [2020]. In general, we say that an optimal solution for [Fair] is *fair with respect to  $f$*  and denote it by  $x^f$ . We remark that this optimal fractional allocation can be computed using state-of-the-art routines for convex optimization due to the concavity of the objective function; see, e.g., Bubeck et al. [2015].

**Near-feasible fair allocations.** Given an instance of assignment-MCRA and a non-decreasing concave function  $f: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ , we say that a mapping  $y: \mathcal{M} \rightarrow \{0, 1\}$  is an  $(\alpha, \delta, \Delta^+)$ -*approximately fair allocation with respect to  $f$*  if the following holds:

$$\sum_{q \in \mathcal{T}_{a,E}} y(a, q) = 1 \quad \text{for every } a \in A, \quad (14)$$

$$|U_{\ell,i}(y) - U_{\ell,i}(x^f)| < \alpha_\ell \cdot U_{\ell,i}^* \quad \text{for every } \ell \in [d], i \in [k_\ell], \quad (15)$$

$$\sum_{a \in A} \sum_{q \in \mathcal{T}_{a,E}} q(r) \cdot y(a, q) - c(r) \leq \delta \quad \text{for every } r \in R, \quad (16)$$

$$\sum_{r \in R} \max\{0, \sum_{a \in A} \sum_{q \in \mathcal{T}_{a,E}} q(r) \cdot y(a, q) - c(r)\} \leq \Delta^+, \quad (17)$$

where  $\alpha \in \mathbb{N}_0^d$ , and  $\delta, \Delta^+ \in \mathbb{N}_0$ . In such near-feasible allocation, the utility of each group  $G_{\ell,i}$  deviates strictly less than  $\alpha_\ell \cdot U_{\ell,i}^*$  from the utilities in the fair with respect to  $f$  solution  $x^f$ , the

capacity of each resource is exceeded by at most  $\delta$ , and the total excess with respect to the resource capacities is at most  $\Delta^+$ . We remark that since in the assignment-MCRA every agent is binding, the total deviation from  $x^f$  on the number of allocated resources is always equal to zero; instead, we aggregate the excess usage of each resource as a sensible parameter in this setting. This notion of approximately fair allocations is closely related to those by Procaccia et al. [2024] and A. et al. [2024], but more general: The former work does not take the individual deviations from resource capacities into account, while the latter imposes the group utilities in the fair fractional allocation as lower bounds only. Using our rounding Theorem 3.1 along with a bound on the number of fractional variables of the initial fair allocation, we get the following guarantees for near-feasible allocations.

**THEOREM 4.1.** *Let  $\mathcal{I}$  be a fractionally feasible instance of assignment-MCRA and  $f: \mathbb{R}_+ \rightarrow \mathbb{R}_+$  a non-decreasing concave function. Let  $\alpha \in \mathbb{N}_0^d$  and  $\delta \in \mathbb{N}_0$  be such that*

$$\sum_{\ell \in [d]} \frac{1}{\alpha_\ell + 1} + \frac{\omega^*}{\delta + 2} \leq \frac{1}{2},$$

*and let  $\Delta^+ = \min\{(\omega^* - 1)|A| + \omega^*|R| + (\omega^* + 1) \sum_{\ell \in [d]} k_\ell, \delta|R|\}$ . Then, there exists an  $(\alpha, \delta, \Delta^+)$ -approximately fair allocation for  $\mathcal{I}$  with respect to  $f$ . Furthermore, this allocation can be found in time polynomial in  $|A|, |R|^{\omega^*}$ , and  $\sum_{\ell \in [d]} k_\ell$ .*

**Consequences for proportional fairness.** In what follows, we discuss the consequences of our approximation guarantees in Theorem 4.1 for the relevant case of proportional fairness, i.e., when we take  $f(z) = \ln(z)$  in [Fair]. We recall that in this case, the optimality conditions guarantee that any optimal fractional resource allocation  $x^f$  of [Fair] satisfies

$$\sum_{\ell \in [d]} \sum_{i \in [k_\ell]} \frac{U_{\ell,i}(x)}{U_{\ell,i}(x^f)} \leq \sum_{\ell \in [d]} k_\ell \text{ for every fractional resource allocation } x,$$

which is the classic proportional fairness notion [Young, 2020]. Recently, Procaccia et al. [2024] studied the design of school allocation policies with provable proportionality guarantees in the presence of groups on the students' side. Their setting is captured by our assignment-MCRA framework when the agents are the students, the resources are the schools, there is a single dimension ( $d = 1$ ) with  $k_1$  many groups, and every student  $a$  is assigned to precisely one single school ( $\omega_a = 1$ ) among their acceptable schools  $\mathcal{T}_{a,E}$ . Formally, for  $d = 1$  and  $k_1 = k$ , we say that  $y: \mathcal{M} \rightarrow \{0, 1\}$  is an  $(\alpha, \delta, \Delta^+)$ -approximately proportional allocation if it satisfies (14), (16), (17), and  $U_i(y) \geq U_i(x)/k - \alpha \cdot U_i^*$  for every  $i \in [k]$  and every fractional resource allocation  $x$ , where we have omitted the dimension subindex.

Thanks to Theorem 4.1, we can trade off the values of  $\alpha$  and  $\delta$  to accommodate the policy-maker priorities in terms of utility approximation and maximum resource capacity augmentation for each resource (i.e., school capacities). For instance, when  $\omega^* = 1$ , both maximum deviations can be set to small constants; some pairs in the Pareto frontier defined by Theorem 4.1 are  $(\alpha, \delta) \in \{(2, 4), (3, 2), (5, 1)\}$ . In practice, slight constant deviations from the capacity of each school constitute a natural goal, and this differentiates our result from previous work in this setting, where the focus was restricted to the deviations from group utilities and the total excess of allocated resources [A. et al., 2024, Procaccia et al., 2024].

While Theorem 4.1 has no direct implications for the case where deviations  $\alpha = 0$  from group utilities or  $\delta = 0$  from school capacities are sought, it is not hard to see that these deviations can be achieved by simply rounding all fractional entries of our initial fair fractional allocation  $x^f$  up or down, respectively. This yields a non-constant deviation with respect to the other objective, potentially up to the order of  $\Delta^+$ . We show that this cannot be avoided: If we require  $\alpha = 0$  deviations

from proportionality, we need to accept non-constant deviations  $\delta$  from the schools' capacities; if we require  $\delta = 0$  deviation from the schools' capacities, we need to accept non-constant deviations  $\alpha$  from proportionality.

**PROPOSITION 4.2.** *For every  $\alpha \in \mathbb{N}_0^d$  and  $\delta, \Delta^+ \in \mathbb{N}_0$ , there exist instances  $\mathcal{I}, \mathcal{I}'$  of assignment-MCRA with  $d = 1$  such that  $\mathcal{I}$  does not admit a  $(0, \delta, \Delta^+)$ -approximately proportional allocation and  $\mathcal{I}'$  does not admit an  $(\alpha, 0, \Delta^+)$ -approximately proportional allocation.*

For  $\mathcal{I}$ , we take an instance where all schools have capacity one and all students belong to a different group and have positive utility for a single school. Proportionality then implies a large violation of the capacity of this school. For  $\mathcal{I}'$ , we consider an example that Procaccia et al. [2024] used to prove a lower bound, linear in the number of groups, on the smallest possible simultaneous deviation from proportionality and aggregate capacity that an allocation can achieve. It is based on cycles of students from two alternating groups and schools of alternating quality, so that forbidding capacity violations leaves a group with zero utility and the other group with a large utility. This proposition implies that, in a sense, our result providing constant (but non-zero) deviations from both group utilities and school capacities is the best we can aim for. The search for the best-possible constants is, however, a natural direction for future work.

We remark that, in addition to providing more flexibility for the maximum deviations, our framework directly handles agent-dependent utility functions, bundles consisting of more than a single resource, and multiple dimensions for the groups, which arise naturally, e.g., when seeking fairness across overlapping groups. Our result opens the door to designing proportional allocation policies under several socio-demographic dimensions, enhancing the policy's fairness guarantees.

**A new notion of envy-freeness.** In the same context of school redistricting ( $d = 1, \omega^* = 1$ ), Procaccia et al. [2024] also studied the existence of approximately *envy-free* allocations. In their work, an allocation is  $\alpha$ -envy-free if for every pair of groups  $i_1, i_2 \in [k]$  there is no alternative allocation where (i) agents in  $i_1$  are allocated a subset of those resources allocated to agents in  $i_2$  in the original allocation, and (ii) the utility of group  $i_1$  increases by more than  $\alpha \cdot U_{i_1}^*$  with respect to the original allocation. Note that this notion only makes sense if the demands of all agents are the same and the utilities of agents in the same group are the same, which are modeling assumptions in Procaccia et al. [2024].<sup>1</sup> However, under this envy-freeness notion, they proved a strong impossibility as there is a family of instances for which the deviation grows linearly in the number of agents.

To get around this impossibility, we introduce a relaxed version of envy-freeness, where the ratio between the total utility of group  $G_{\ell,i}$  for its allocation and its total utility for the allocation of the group  $G_{\ell,j}$  should not be smaller than the ratio between the sizes of these groups. Note that when groups have unit size, this is equivalent to the classic notion of envy-freeness in fair division of indivisible goods; e.g., [Moulin, 2004]. We still consider common demands, i.e.,  $\omega_a = \omega^*$  for every  $a \in A$ . We thus denote the (common) set of possible bundles by  $\mathcal{T}$  for simplicity. We say that an instance of assignment-MCRA is *group-homogeneous* if  $\omega_a = \omega^*$  for every  $a \in A$  and, for every dimension  $\ell \in [d]$ , every  $i \in [k_\ell]$ , and every  $a, b \in G_{\ell,i}$ , we have  $u_a(q) = u_b(q)$  for every  $q \in \mathcal{T}$ . We denote this common (group) utility function by  $u_{\ell,i}$ . For a group-homogeneous instance of assignment-MCRA, we say that  $y: \mathcal{M} \rightarrow \{0, 1\}$  (resp.  $[0, 1]$ ) is an  $(\alpha, \delta)$ -approximately envy-free allocation (resp. fractional) if it satisfies (14), (16), and

$$\frac{|G_{\ell,i}|}{|G_{\ell,j}|} \sum_{b \in G_{\ell,j}} \sum_{q \in \mathcal{T}} u_{\ell,i}(q) \cdot y(b, q) - U_{\ell,i}(y) < \alpha_\ell \cdot U_{\ell,i}^* \quad (18)$$

<sup>1</sup>In fact, they assume that the utilities of all agents are the same.

for every  $\ell \in [d]$  and  $i, j \in [k_\ell]$ , where  $\alpha \in \mathbb{N}^d$  and  $\delta, \Delta \in \mathbb{N}$ . That is, the maximum deviation from our notion of envy freeness for each group  $G_{\ell,i}$  is captured by  $\alpha_\ell$  times the maximum utility of this group for a single resource, and the meaning of the maximum deviation  $\delta$  from resource capacities is the same as in approximately fair allocations. We omit the total excess  $\Delta^+$  because only the trivial bound  $\delta|R|$  on this deviation will remain in this case. We obtain the following result.

**THEOREM 4.3.** *Let  $\mathcal{I}$  be a group-homogeneous and fractionally feasible instance of assignment-MCRA, and let  $\alpha \in \mathbb{N}_0^d$  and  $\delta \in \mathbb{N}_0$  be such that*

$$\sum_{\ell \in [d]} \frac{2(k_\ell - 1)}{\alpha_\ell + 1} + \frac{\omega^*}{\delta + 1} \leq \frac{1}{2}.$$

*Then, there exists an  $(\alpha, \delta)$ -approximately envy-free allocation for  $\mathcal{I}$ .*

These approximately envy-free allocations can be found by applying an iterative rounding procedure analogous to the one described in Algorithm 1, starting from an envy-free fractional allocation. The difference with our original rounding algorithm is the number of constraints we must maintain in each step, which becomes quadratic in the number of groups in each dimension due to the envy-freeness constraints. We remark that the condition in Theorem 4.3 holds for  $\alpha_\ell \in \Theta(dk_\ell)$  for each  $\ell \in [d]$  and  $\delta \in \Theta(1)$ , thus breaking the linear dependence in  $|A|$  from the more stringent version by Procaccia et al. [2024].

## 4.2 Stable Matchings with Couples

In this section, unlike the previous one, the non-existence of fair resource allocations does not come from group fairness constraints but from stability; while stable matchings are guaranteed to exist in the basic single-demand setting [Gale and Shapley, 1962], this is not the case for multi-demand agents. To illustrate how our framework can also accommodate stability requirements, we show how our approach exploits the fact that, for agents with demand up to  $\omega^* = 2$ , the existence of fractional allocations satisfying stability is guaranteed, and rounding these allocations preserves this property [Nguyen and Vohra, 2018]. Using our rounding Theorem 3.1, we can directly recover a recent guarantee by Nguyen and Vohra [2018] for the existence of near-feasible stable allocations, and we further extend this setting to handle group fairness requirements and allocation stability to guarantee the existence of near-feasible, stable, and fair allocations (Theorem 4.4).

**Instances and stability.** An instance of the *couples-MCRA* problem is a tuple  $(A, R, E, G, c, u, \succ)$ , where  $A, R, E, G, c$ , and  $u$  are structured in the same way as in assignment-MCRA; see Section 4.1 for the details. We omit the parameter  $\omega$  from the instance description as we fix  $\omega^* = 2$  and partition the agents accordingly into  $A = A_1 \dot{\cup} A_2$ , where  $\omega_a = 1$  for every  $a \in A_1$  (*single agents*) and  $\omega_a = 2$  for every  $a \in A_2$  (*couples*). Following the notation introduced in Section 4.1,  $\mathcal{T}_{a,E}$  denotes the acceptable bundles for each agent  $a \in A$ , and  $\mathcal{M} = \{(a, q) : a \in A, q \in \mathcal{T}_{a,E}\}$  is the set of feasible agent-bundle pairs. Finally,  $\succ$  represents a set of linear orders:  $\succ_r$  is a linear order over  $\{a \in A : (a, q) \in \mathcal{M}\}$  representing the preferences of each resource  $r \in R$  and  $\succ_a$  is a linear order over  $\mathcal{T}_{a,E}$  representing the preferences of each agent  $a \in A$ .

A mapping  $x$  on  $\mathcal{M}$  is a (fractional) resource allocation for  $\mathcal{I}$  if its natural extension  $x'$  on  $\mathcal{E} = \{(a, q) : a \in A, q \in \mathcal{T}_a\}$ , given by  $x'(a, q) = x(a, q)$  for every  $(a, q) \in \mathcal{M}$  and  $x'(a, q) = 0$  for every  $(a, q) \in \mathcal{E} \setminus \mathcal{M}$ , is a (fractional) resource allocation for the instance  $(A, \emptyset, R, G, \omega, c)$  of MCRA, where we have replaced  $A'$  by  $\emptyset$ . Note that every instance of couples-MCRA admits a fractional allocation due to the absence of binding agents; in particular, the trivial allocation where no agent is assigned is always feasible. A resource allocation  $x$  on  $\mathcal{M}$  is *blocked* when some of the following three situations happen:

- (1) There is a pair  $(a, r)$  with  $a \in A_1$  and  $x(a, q) = 1$  such that  $a$  prefers  $r$  to  $q$  (1) and  $r$  either has remaining capacity or prefers  $a$  over some other agent allocated in  $x$ .
- (2) There is a pair  $(a, r)$  with  $a \in A_2$  and  $x(a, q) = 1$  such that  $a$  prefers  $r$  over  $q$  and  $r$  either has remaining capacity for the couple  $a$  or prefers  $a$  over some other agent(s) allocated in  $x$ .
- (3) There is a couple  $a \in A_2$  and two different resources  $r, r'$  such that  $a$  would prefer to be assigned to the resources  $r$  and  $r'$ , one each, over their current allocation in  $x$ , and each of the resources has remaining capacity or prefers the corresponding element of the couple  $a$  over some other agent allocated in  $x$ .

A resource allocation  $x$  is *stable for the capacities  $c$*  if it is not blocked.

**Fractional allocations and fairness.** We now extend our notion of group fairness and approximately fair allocations from Section 4.1 to the case where we also require stability, provided the maximum demand is  $\omega^* = 2$ . Consider the following linear program introduced by Nguyen and Vohra [2018] with variables  $x: \mathcal{M} \rightarrow [0, 1]$ :

$$\begin{aligned} \sum_{(a,q) \in \mathcal{M}} q(r) \cdot x(a, q) &\leq c(r) && \text{for every } r \in R, \\ \sum_{q \in \mathcal{T}_{a,E}} x(a, q) &\leq 1 && \text{for every } a \in A. \end{aligned} \quad [\text{LP-Stable}]$$

Nguyen and Vohra [2018] showed that the polytope given by [LP-Stable] has a dominating extreme point  $x: \mathcal{M} \rightarrow [0, 1]$  and that any rounding of such point is a stable allocation. Since our use of dominating extreme points is restricted to this property, we refer to Nguyen and Vohra [2018] for the precise definition in this setting. In what follows, we exploit this fact to extend our result regarding approximate group fairness and find allocations that are both approximately fair and stable. Similarly to Section 4.1, given a non-decreasing concave function  $f: \mathbb{R}_+ \rightarrow \mathbb{R}_+$  and an instance  $\mathcal{I}$  of couples-MCRA, we consider the maximization problem

$$\max \left\{ \sum_{\ell \in [d]} \sum_{i \in [k_\ell]} f(U_{\ell,i}(x)) : x \text{ is a dominating extreme point for [LP-Stable]} \right\}, \quad [\text{Fair-Stable}]$$

where we have now restricted our feasible set to dominating extreme points of the polytope given by [LP-Stable] so that rounding any such point will give a stable allocation. We say that an optimal solution for [Fair-Stable] is *stable and fair with respect to  $f$*  and denote it by  $x^f$ .

In terms of computational efficiency, it is unknown whether the problem of finding a dominating extreme point of a polytope defined by hypergraph constraints such as [LP-Stable] can be solved in polynomial time. This has only been answered in the positive for special cases that do not capture the polytope [LP-Stable], e.g., Chandrasekaran et al. [2025], Faenza et al. [2025], and thus we do not know whether the optimization program [Fair-Stable] can be solved efficiently. We remark that this is the case even without the group fairness, i.e., in the setting of Nguyen and Vohra [2018].

**Near-feasible stable and fair allocations.** Given an instance of couples-MCRA and a non-decreasing concave function  $f: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ , we say that  $y: \mathcal{M} \rightarrow \{0, 1\}$  is an  $(\alpha, \delta, \Delta)$ -*approximately fair stable allocation with respect to  $f$* , for  $\alpha \in \mathbb{N}_0^d$  and  $\delta, \Delta \in \mathbb{N}$ , if the following holds:

$$\sum_{q \in \mathcal{T}_{a,E}} y(a, q) \leq 1 \quad \text{for every } a \in A, \quad (19)$$

$$|U_{\ell,i}(y) - U_{\ell,i}(x^f)| < \alpha_\ell \cdot U_{\ell,i}^* \quad \text{for every } \ell \in [d], i \in [k_\ell], \quad (20)$$

$$\sum_{a \in A} \sum_{q \in \mathcal{T}_{a,E}} q(r) \cdot y(a, q) - c(r) \leq \delta \quad \text{for every } r \in R, \quad (21)$$

$$\sum_{a \in A} \sum_{q \in \mathcal{T}_{a,E}} \omega_a \cdot y(a, q) - \sum_{r \in R} c(r) \leq 2\Delta, \quad (22)$$

and  $y$  is stable for capacities  $c': R \rightarrow \mathbb{R}_+$  given by  $c'(r) = \sum_{a \in A} \sum_{q \in \mathcal{T}_{a,E}} q(r) \cdot y(a, q)$ , i.e., by the actual number of assigned agents. We remark that the capacity violation conditions (21) and (22) can alternatively be expressed in absolute values with respect to the stable and fair allocation  $x^f$ ,

which may be more suitable for cases with lower bounds on the capacities; we stick to our setting with upper bounds for consistency. A direct application of Theorem 3.1 yields the following result.

**THEOREM 4.4.** *Let  $\mathcal{I}$  be a fractionally feasible instance of couples-MCRA and  $f: \mathbb{R}_+ \rightarrow \mathbb{R}_+$  a non-decreasing concave function. Let  $\alpha \in \mathbb{N}_0^d$  and  $\delta \in \mathbb{N}_0$  be such that*

$$\sum_{\ell \in [d]} \frac{1}{\alpha_\ell + 1} + \frac{2}{\delta + 2} \leq \frac{1}{2}.$$

*Then, there exists an  $(\alpha, \delta, 2)$ -approximately fair stable allocation for  $\mathcal{I}$  with respect to  $f$ .*

We remark that in the case without fairness constraints ( $d = 0$ ), we recover the bounds  $\delta = \Delta = 2$  from Nguyen and Vohra [2018]. Our result, however, allows to incorporate these constraints while keeping the deviations from resource capacities bounded by small constants. For example, in the case of a one-dimensional partition of the agents, one can guarantee deviations from group fairness of at most  $\alpha_1 = 5$  by increasing the deviations  $\delta$  to 4, while keeping the deviation from the aggregate capacity at  $\Delta = 2$ .

### 4.3 Multidimensional Apportionment

In the *multidimensional apportionment problem*, introduced by Balinski and Demange [1989a,b] for the case of two dimensions and extended by Cembrano et al. [2022] to an arbitrary number of dimensions  $d \geq 2$ , the goal is to allocate the seats of a representative body proportionally across several dimensions. Classic apportionment methods, e.g., divisor methods or Hamilton’s method, aim to assign seats to districts proportionally to their population or political parties proportionally to their electoral support [Balinski and Young, 2010, Pukelsheim, 2017]. However, one may consider several natural dimensions to decide the seat allocation, both by requiring proportionality across them with respect to their electoral support and by incorporating bounds on the number of seats that groups within them should receive. In addition to political parties and geographical divisions, natural dimensions include demographics of the elected members such as gender or ethnicity; see, e.g., Arnosti et al. [2024], Cembrano et al. [2024], Mathieu and Verdugo [2024]. Using our rounding Theorem 3.1, we improve over the result by Cembrano et al. [2022] to get enhanced near-feasibility guarantees for multidimensional apportionment. Remarkably, we can further bound the total deviation from the house size, while the rounding algorithm by Cembrano et al. only controls the deviations on each dimension.

**Signpost sequences and rounding rules.** The core idea behind proportional apportionment, and in particular divisor methods, is to *scale and round*. Then, to formally introduce the multidimensional apportionment problem, we need to define the idea of a *rounding rule*, which in turn requires the definition of a *signpost sequence*. A signpost sequence is a function  $s: \mathbb{N}_0 \rightarrow \mathbb{R}_+$  such that  $s(0) = 0$ ,  $s(t) \in [t - 1, t]$  for every  $t \in \mathbb{N}$ , and  $s(t + 1) > s(t)$  for every  $t \in \mathbb{N}$ . Given a signpost sequence  $s$ , the rounding rule  $\llbracket \cdot \rrbracket_s$  is defined as follows: For every  $q \in \mathbb{R}_+$ , we let  $\llbracket q \rrbracket_s = \{t\}$  if  $s(t) < q < s(t + 1)$ , and  $\{t - 1, t\}$  if  $q = s(t)$ . In simple terms, any value  $q \in [t - 1, t]$  gets rounded up if  $q > s(t)$ , down if  $q < s(t)$ , and we allow both  $t - 1$  and  $t$  as possible roundings is  $q = s(t)$ .

**Instances and near-feasible apportionments.** An instance of the *multidimensional apportionment* problem, or MA for short, is a tuple  $\mathcal{I} = (G, E, V, b, B, c)$ , where  $G$  is structured as in MCRA,  $E \subseteq \prod_{\ell=1}^d G_{\ell,i}$ ,  $V \in \mathbb{N}^E$ ,  $b_{\ell,i}, B_{\ell,i} \in \mathbb{N}_0$  for every  $\ell \in [d]$  and  $i \in [k_\ell]$ , and  $c \in \mathbb{N}$ . Each  $\ell \in [d]$  represents a dimension according to which the candidates are grouped for the election, e.g., their political parties, districts, gender, or ethnicity. The set  $E$  contains the tuples of groups receiving a strictly positive number of votes, which are specified in the tensor  $V$ . The values  $b_{\ell,i}, B_{\ell,i}$  represent a lower and upper bound on the number of seats that should be assigned to group  $G_{\ell,i}$ , respectively,

and  $c$  is the size of the house. We say that a dimension  $\ell \in [d]$  is *binding* if  $b_{\ell,i} = B_{\ell,i}$  for every  $i \in [k_\ell]$ , i.e., if there are hard bounds on the number of seats that all groups in this dimension should receive.

In this section, a mapping  $x : E \times [c] \rightarrow \{0, 1\}$  (resp.  $[0, 1]$ ) is an *apportionment* (resp. fractional) if  $b_{\ell,i} \leq \sum_{e \in E: e_\ell = i} \sum_{t \in [c]} x(e, t) \leq B_{\ell,i}$  for every  $\ell \in [d]$  and  $i \in [k_\ell]$ , and  $\sum_{e \in E} \sum_{t \in [c]} x(e, t) = c$ , where the first condition ensures the group bounds and the second one guarantees to fulfill the house size. As usual, we call an instance *fractionally feasible* if it admits a fractional apportionment. For an instance of MA, and given  $\alpha \in \mathbb{N}_0^d$  and  $\Delta \in \mathbb{N}_0$ , a mapping  $y : E \times [c] \rightarrow \{0, 1\}$  is an  $(\alpha, \Delta)$ -*approximately proportional apportionment* if there exist values  $\lambda_{\ell,i} > 0$  for every  $\ell \in [d]$  and every  $i \in [k_\ell]$  such that

$$\sum_{t \in [c]} x(e, t) = \lceil \sum_{\ell \in [d]} \lambda_{\ell, e_\ell} \rceil \quad \text{for every } e \in E, \quad (23)$$

$$b_{\ell,i} - \alpha_\ell \leq \sum_{e \in E: e_\ell = i} \sum_{t \in [c]} x(e, t) \leq B_{\ell,i} + \alpha_\ell \quad \text{for every } \ell \in [d], i \in [k_\ell], \quad (24)$$

$$|\sum_{e \in E} \sum_{t \in [c]} x(e, t) - c| \leq \Delta. \quad (25)$$

In addition to approximately respecting the group bounds and the house size, we have now imposed the natural notion of proportionality in this setting: Each group  $G_{\ell,i}$  has an associated *multiplier*  $\lambda_{\ell,i}$  and each tuple is assigned a number of seats given by its number of votes scaled by all the multipliers associated with groups in the tuple. The existence of multipliers such that all deviations are zero is guaranteed when  $d \in \{1, 2\}$  [Balinski and Demange, 1989a,b, Gaffke and Pukelsheim, 2008], but not when  $d \geq 3$  [Cembrano et al., 2022]. However, we exploit the fact that fractional apportionments are guaranteed to exist and that proportionality is kept upon rounding.

**Improved guarantee for near-feasible apportionments.** As a starting point for applying our rounding theorem, we use the following linear program, introduced by Cembrano et al. [2022]:

$$\begin{aligned} \min \quad & \sum_{e \in E} \sum_{t \in [c]} x(e, t) \ln(s(t)/V_e) && \text{[LP-MA]} \\ \text{s.t.} \quad & b_{\ell,i} \leq \sum_{e \in E: e_\ell = i} \sum_{t \in [c]} x(e, t) \leq B_{\ell,i} \quad \text{for every } \ell \in [d], i \in [k_\ell], \\ & \sum_{e \in E} \sum_{t \in [c]} x(e, t) = c, \\ & x(e, t) \in [0, 1] \quad \text{for every } e \in E, t \in [c]. \end{aligned}$$

Using a primal-dual analysis, it can be shown that [LP-MA] fully characterizes proportional apportionments, in the sense that a mapping  $x$  is a proportional apportionment if and only if it is an optimal solution for it [Cembrano et al., 2022, Theorem 1]. While this linear program does not have, in general, an optimal integral solution, we can still round an optimal solution of [LP-MA] using Theorem 3.1 and maintain the proportionality condition (23) to obtain improved approximation guarantees for near-feasible apportionments. The proof consists of a direct application of Theorem 3.1 for an MCRA instance with all utilities equal to 1. Note that this integer instantiation of the utilities implies the terms  $1/(\alpha_\ell + 1)$  from the main theorem are replaced by  $1/(\alpha_\ell + 2)$ , which translates into smaller deviations.

**THEOREM 4.5.** *Let  $\mathcal{I}$  be an instance of MA. Let  $\alpha \in \mathbb{N}_0^d$  be such that  $\sum_{\ell \in [d]} 1/(\alpha_\ell + 2) \leq 1$  and  $\Delta \in \mathbb{N}_0$  defined by  $\Delta = \min\{\min\{\Delta_\ell : \ell \in [d]\}, \lceil 1/(1 - \sum_{\ell \in [d]} 1/(\alpha_\ell + 2)) - 2 \rceil\}$ , where  $\Delta_\ell = \alpha_\ell k_\ell$  for every binding dimension  $\ell \in [d]$  and  $\Delta_\ell = (\alpha_\ell + 1)k_\ell - 1$  for every non-binding dimension. Then, there exists an  $(\alpha, \Delta)$ -approximate proportional apportionment for  $\mathcal{I}$ . Furthermore, this solution can be found in time polynomial in  $|E|$  and  $c$ .*

We remark that our guarantees in Theorem 4.5 improve over the result by Cembrano et al. [2022] in that we can bound the total deviation  $\Delta$  from the house size, while the rounding algorithm by Cembrano et al. only allows trading off the deviations on each dimension. On the one hand, we get

small deviations from the house size if the sum  $\sum_{\ell \in [d]} 1/(\alpha_\ell + 2)$  is not too close to 1; for example, for  $d = 3$ , taking  $\alpha \in \{(0, 6, 6), (1, 2, 4), (2, 2, 2)\}$  allows taking  $\Delta = \lceil 1/(1 - \sum_{\ell \in [d]} 1/(\alpha_\ell + 2)) - 2 \rceil = 2$ . On the other hand, we also get slight deviations from the house size when a dimension  $\ell$  has few groups; for example, the case with  $k_\ell = 2$  (which may arise, for example, if the corresponding dimension is gender) yields deviations from the house size of at most  $2\alpha_\ell$  if the dimension is binding and at most  $2(\alpha_\ell + 1) - 1$  otherwise. Note that, in particular, the existence of a binding dimension  $\ell$  with deviations  $\alpha_\ell = 0$  implies exactly fulfilling the house size.

We finally remark that, in principle, our rounding algorithm could be further applied to more general apportionment settings in which several representative bodies are to be elected, but this would require novel structural results capturing proportionality in this general case.

## 5 Discussion and Final Remarks

Our work develops a general iterative rounding framework for resource allocation in two-sided markets that produces near-feasible allocations while controlling violations in both resource capacities and fairness targets. Fed with an appropriate fractional allocation and tuning a small set of deviation parameters, our approach gives a unified and flexible tool for handling fairness in several resource allocation problems, including school allocations, stable matchings, and political apportionment. In doing so, it not only recovers guarantees from these specialized settings but also provides new ones for previously omitted objectives and extends prior guarantees to cases involving multi-demand agents and overlapping, multidimensional group structures. The resulting near-feasible solutions provide robust performance guarantees, even when exact feasibility is theoretically impossible or computationally prohibitive. Furthermore, the flexibility in choosing the deviation parameters enables policy designers to tailor the trade-offs between efficiency, fairness, and resource augmentation to the needs of a specific application.

In Section 4.1, we propose a general multidimensional assignment model to handle general group fairness requirements by a convex optimization-driven approach. Using our rounding Theorem 3.1, we get a range of flexible guarantees on the existence of near-feasible and fair solutions. In particular, we introduce a relaxed notion of group envy-freeness that can escape existing impossibilities to accommodate efficient and near-feasible allocations; we believe this new notion might be of independent interest and deserves further study.

The implications in Section 4.2 of our main rounding Theorem 3.1 contribute to a recent line on the computation of near-feasible allocation under stability and complex constraints. Nguyen and Vohra [2019] consider a stable matching problem where the hospitals classify their acceptable set of doctors according to types. While they model proportionality employing lower or upper bounds on the proportion of doctors of each type to be assigned for each hospital, our approach in Section 4.2 incorporates general fairness considerations as the objective function in a convex optimization program. Theorem 4.4 ensures controlled deviations that can be tailored to real-world constraints, giving flexibility to policy designers in implementing stability and optimizing trade-offs between efficiency and fairness. We also believe our framework can be further exploited under different stability concepts, e.g., the group-stability notion considered by Nguyen et al. [2021] to match families and localities with contracts and budget constraints.

Beyond the application settings showcased in this work, our framework could be further explored in other domains, such as makespan scheduling [Feldman et al., 2025, Saha and Srinivasan, 2018] and fair algorithms for clustering [Bera et al., 2019, Makarychev and Vakilian, 2021].

## Acknowledgments

This work was partially funded by the ANID (Chile) through Grant FONDECYT 1241846.

## References

- Santhini K. A., Kamesh Munagala, Meghana Nasre, Govind S Sankar, et al. 2024. Approximation Algorithms for School Assignment: Group Fairness and Multi-criteria Optimization. *arXiv preprint arXiv:2403.15623* (2024).
- Narges Ahani, Tommy Andersson, Alessandro Martinello, Alexander Teytelboym, and Andrew C Trapp. 2021. Placement optimization in refugee resettlement. *Operations Research* 69, 5 (2021), 1468–1486.
- Tommy Andersson and Lars Ehlers. 2020. Assigning refugees to landlords in Sweden: Efficient, stable, and maximum matchings. *The Scandinavian Journal of Economics* 122, 3 (2020), 937–965.
- Nick Arnosti, Carlos Bonet, and Jay Sethuraman. 2024. Explainable Affirmative Action. In *Proceedings of the 25th ACM Conference on Economics and Computation*. 310–310.
- Haris Aziz and Florian Brandl. 2021. Efficient, fair, and incentive-compatible healthcare rationing. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 103–104.
- Michel Balinski and Gabrielle Demange. 1989a. Algorithms for proportional matrices in reals and integers. *Mathematical Programming* 45, 1-3 (1989), 193–210.
- Michel Balinski and Gabrielle Demange. 1989b. An Axiomatic Approach to Proportionality Between Matrices. *Mathematics of Operations Research* 14, 4 (1989), 700–719.
- Michel Balinski and H. Peyton Young. 2010. *Fair representation: Meeting the ideal of one man, one vote*. Brookings Institution Press.
- Nikhil Bansal. 2010. Constructive algorithms for discrepancy minimization. In *Proceedings of the 51st Annual Symposium on Foundations of Computer Science*. 3–10.
- Nikhil Bansal, Daniel Dadush, and Shashwat Garg. 2019. An algorithm for Komlós conjecture matching Banaszczyk’s bound. *SIAM J. Comput.* 48, 2 (2019), 534–553.
- Nikhil Bansal and Shashwat Garg. 2017. Algorithmic discrepancy beyond partial coloring. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 914–926.
- Nikhil Bansal, Haotian Jiang, Sahil Singla, and Makrand Sinha. 2020. Online vector balancing and geometric discrepancy. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*. 1139–1152.
- József Beck and Tibor Fiala. 1981. Integer-making theorems. *Discrete Applied Mathematics* 3, 1 (1981), 1–8.
- Suman Bera, Deeparnab Chakrabarty, Nicolas Flores, and Maryam Negahbani. 2019. Fair algorithms for clustering. *Advances in Neural Information Processing Systems* 32 (2019).
- Federico Bobbio, Margarida Carvalho, Andrea Lodi, Ignacio Rios, and Alfredo Torricco. 2023. Capacity Planning in Stable Matching: An Application to School Choice. In *Proceedings of the 24th ACM Conference on Economics and Computation*. 295–295.
- Sébastien Bubeck et al. 2015. Convex optimization: Algorithms and complexity. *Foundations and Trends in Machine Learning* 8, 3-4 (2015), 231–357.
- Boris Bukh. 2016. An improvement of the Beck–Fiala theorem. *Combinatorics, Probability and Computing* 25, 3 (2016), 380–398.
- Javier Cembrano, José Correa, Gonzalo Díaz, and Victor Verdugo. 2024. Proportionality in Multiple Dimensions to Design Electoral Systems. *arXiv preprint arXiv:2410.03304* (2024).
- Javier Cembrano, José Correa, and Victor Verdugo. 2022. Multidimensional political apportionment. *Proceedings of the National Academy of Sciences* 119, 15 (2022), e2109305119.
- Karthekeyan Chandrasekaran, Yuri Faenza, Chengyue He, and Jay Sethuraman. 2025. Scarf’s Algorithm on Arborescence Hypergraphs. *To appear in ICALP* (2025).
- David Delacrétaz, Scott Duke Kominers, and Alexander Teytelboym. 2016. Refugee resettlement. *University of Oxford Department of Economics Working Paper* (2016).
- Yuri Faenza, Chengyue He, and Jay Sethuraman. 2025. Scarf’s algorithm and stable marriages. *Mathematics of Operations Research* (2025).
- Michal Feldman, Jugal Garg, Vishnu V Narayan, and Tomasz Ponitka. 2025. Proportionally Fair Makespan Approximation. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 13839–13846.
- Norbert Gaffke and Friedrich Pukelsheim. 2008. Divisor methods for proportional representation systems: An optimization approach to vector and matrix apportionment problems. *Mathematical Social Sciences* 56, 2 (2008), 166–184.
- David Gale and Lloyd S Shapley. 1962. College admissions and the stability of marriage. *The American Mathematical Monthly* 69, 1 (1962), 9–15.
- Guillaume Haeringer. 2018. *Market design: auctions and matching*. MIT Press.
- Nicole Immorlica, Federico Echenique, and Vijay V Vazirani. 2023. *Online and Matching-based Market Design*. Cambridge University Press.
- Lap Chi Lau, Ramamoorthi Ravi, and Mohit Singh. 2011. *Iterative methods in combinatorial optimization*. Vol. 46. Cambridge University Press.

- Shachar Lovett and Raghu Meka. 2015. Constructive discrepancy minimization by walking on the edges. *SIAM J. Comput.* 44, 5 (2015), 1573–1582.
- Yury Makarychev and Ali Vakilian. 2021. Approximation algorithms for socially fair clustering. In *Conference on Learning Theory*. 3246–3264.
- Claire Mathieu and Victor Verdugo. 2024. Apportionment with parity constraints. *Mathematical Programming* 203, 1 (2024), 135–168.
- Hervé Moulin. 2004. *Fair division and collective welfare*. MIT press.
- Hai Nguyen, Thành Nguyen, and Alexander Teytelboym. 2021. Stability in matching markets with complex constraints. *Management Science* 67, 12 (2021), 7438–7454.
- Thành Nguyen and Rakesh Vohra. 2018. Near-Feasible Stable Matchings with Couples. *American Economic Review* 108, 11 (2018), 3154–69.
- Thành Nguyen and Rakesh Vohra. 2019. Stable matching with proportionality constraints. *Operations Research* 67, 6 (2019), 1503–1519.
- Parag A Pathak, Tayfun Sönmez, M Utku Ünver, and M Bumin Yenmez. 2021. Fair allocation of vaccines, ventilators and antiviral treatments: leaving no ethical value behind in health care rationing. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 785–786.
- Ariel Procaccia, Isaac Robinson, and Jamie Tucker-Foltz. 2024. School Redistricting: Wiping Unfairness Off the Map. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms*. 2704–2724.
- Friedrich Pukelsheim. 2017. *Proportional Representation*. Springer International Publishing.
- Alvin E Roth. 2018. Marketplaces, markets, and market design. *American Economic Review* 108, 7 (2018), 1609–1658.
- Thomas Rothvoss. 2017. Constructive discrepancy minimization for convex sets. *SIAM J. Comput.* 46, 1 (2017), 224–234.
- Barna Saha and Aravind Srinivasan. 2018. A new approximation technique for resource-allocation problems. *Random Structures & Algorithms* 52, 4 (2018), 680–715.
- H Peyton Young. 2020. *Equity: in theory and practice*. Princeton University Press.