Noname manuscript No.
(will be inserted by the editor)

A Family of Convex Models to Achieve Fairness through Dispersion Control

Abhay Singh Bhadoriya · Deepjyoti Deka · Kaarthik Sundar

Received: date / Accepted: date

Abstract Controlling the dispersion of a subset of decision variables in an optimization problem is crucial for enforcing fairness or load-balancing across a wide range of applications. Building on the well-known equivalence of finite-dimensional norms, the note develops a family of parameterized convex models that regulate the dispersion of a vector of decision-variable values through its coefficient of variation. Each model contains a single parameter that takes a value in the interval $[0, 1]$. When the parameter is set to zero, the model imposes only a trivial constraint on the optimization problem; when set to one, it enforces equality of all the decision variables. As the parameter varies, the coefficient of variation is provably bounded above by a monotonic function of that parameter. The note also presents theoretical results that relate the space of feasible solutions to all the models. Finally it compares the models' solution quality on a variant of the assignment problem that regulates the dispersion in the assignment costs.

Keywords Dispersion control · Fairness · Convex optimization · Parametric family

Mathematics Subject Classification (2020) 90C05 · 90C25 · 91B32 · 90B50

A. Bhadoriya
E-mail: abhaysbh@amazon.com

D. Deka
MIT Energy Institute, MIT, Cambridge, MA. E-mail: deepj87@mit.edu

K. Sundar
Los Alamos National Laboratory, Los Alamos, NM. E-mail: kaarthik@lanl.gov

1 Introduction

Fairness has deep roots in game theory, where foundational bargaining solutions were rigorously characterized via axiomatic frameworks. Nash's seminal work formalized the bargaining problem through a set of axioms–Pareto efficiency, symmetry, invariance to affine transformations, and independence of irrelevant alternatives–which uniquely identify the Nash bargaining solution [22]. Subsequent refinements, such as the Kalai-Smorodinsky solution, altered these axioms to prioritize proportionality among agents [15]. In economics and communication theory [28, 26], these game-theoretic fairness axioms guide the design of utility functions so that decentralized, self-interested agents naturally converge to equitable outcomes. By carefully embedding fairness criteria into each agent's payoff, one ensures that the collective solution respects core notions of equity and parity, even when each participant acts independently.

However, in many application domains, problems are centrally optimized: a single decision maker minimizes a global cost or maximizes overall efficiency while imposing load-balancing requirements to distribute burdens–battery discharge, travel distance, computational load [10, 14, 4]–uniformly across agents. Here, "fairness" is more accurately viewed as dispersion control, wherein the vector of decision-variable values must exhibit limited spread under a centralized objective. In contemporary optimization, efforts have been made to extend beyond achieving minimal cost or maximal utility by explicitly regulating the dispersion of selected decision variables. By bounding or penalizing the spread through metrics such as variance, range, coefficient of variation, or more sophisticated equity indices [28], one can embed fairness or load-balancing constraints directly into existing optimization models. However, many dispersion metrics do not admit convex reformulations; directly bounding them thus renders the optimization problem non-convex and computationally challenging [28]. Nevertheless, the dispersion-control paradigm is ubiquitous: ensuring equitable service levels in cloud-computing clusters [10], allocating shared spectrum in wireless networks [20], operating power grids during extreme events [27], among others. Controlling dispersion thus provides a unifying formalism that couples classical efficiency criteria with societally critical notions of parity, equity, and fairness.

1.1 Literature Review

In the literature, two principal modelling paradigms have emerged for incorporating fairness into decision-making problems that traditionally focus on maximizing efficiency.

The first paradigm replaces the efficiency objective with an explicitly fairness-oriented objective–for example, the Rawlsian max–min objective [23], α-fairness [21] (a concave objective with α as a parameter), proportional fairness [26] (a nonlinear objective), or ℓ_p-norm dispersion objectives [2] (a convex objective with p as a parameter). In this framework, parity is pursued as the primary goal. However, a well-documented limitation [27] is that the parameters–α in the case of α-fairness and p in the case of ℓ_p-norm objectives–do not have a direct or monotonic relationship with the fairness level achieved in the optimal solution. In other words, adjusting these parameters does not necessarily increase or decrease fairness.

Fairness levels are typically assessed a posteriori using indices such as the Gini coefficient [11], Atkinson index [1], Hoover index [18], and Jain's index [13] etc. A substantial body of work has sought to enforce fairness by directly minimizing or maximizing such indices, or by reducing dispersion measures such as the coefficient of variation or the range of decision variables [28]. These approaches, however, generally suffer from one of two drawbacks: (i) the resulting model is non-convex, or (ii) the model produces a single fair solution that compromises efficiency–an outcome often unacceptable in applications where cost considerations are critical. Despite these issues, this paradigm remains widely used, given the limited availability of alternative frameworks for fairness in centralized optimization.

The second, more recent paradigm–which motivates the present study–retains an efficiency-centric objective while constraining dispersion. Examples include the ε-fairness second-order cone approach [27] or the imposition of explicit upper bounds on fairness indices such as the Gini coefficient [2], Jain's index, coefficient of variation or range [28]. Here, fairness is modelled by enforcing a lower (or upper) bound on the minimum (or maximum) acceptable fairness level while preserving the efficiency objective. This approach does not always yield a convex optimization problem, as noted in [28], and even when convexity is preserved, computing a suitable lower bound for fairness often requires trial-and-error. Two notable exceptions are the ε-fairness second-order cone method [27], which achieves this using a single second-order cone constraint, and the direct bounding of the Gini coefficient [4], which reformulates

the problem into a set of linear inequalities at the cost of introducing $O(n^2)$ auxiliary variables. For a comprehensive review of both paradigms, the reader is referred to the survey in [28]. This note examines ε-fairness and extends it to a family of convex models. The appeal of ε-fairness lies in its desirable properties: it enforces fairness through a single parametric convex constraint with $\varepsilon \in [0, 1]$ as the parameter, and the parameter ε is monotonic in the level of fairness. Moreover, it naturally facilitates the study of the Pareto front of optimal solutions balancing efficiency and fairness in a bi-objective setting [4].

We now present the main contributions of this note in the context of this literature.

1.2 Contributions

The contributions of the note are as follows: we extend the recently introduced model of ε-fairness [27]–that controls the dispersion of decision variables via its coefficient of variation [8]–to a countably infinite family of convex models. Each model is derived from the first principles through the classical equivalence of finite-dimensional norms [12], namely between the ℓ_p-norm and the ℓ_1-norm for any $p \geqslant 2$. All models are governed by a single parameter, $\varepsilon \in [0, 1]$: when $\varepsilon = 0$, no constraint on the dispersion is imposed, and when $\varepsilon = 1$, the constraint forces the dispersion to be zero. In particular, for $p = 2$, the constraint reduces to a second-order cone [27]; for $p = \infty$, it reduces to a set of linear inequalities; and for $p \in (2, \infty)$, it reduces to the ℓ_p-norm cone. In addition, we present theoretical results that characterize the relationships among these convex models for all values of ε, including their relative strength, inclusion properties, and points of equivalence. Finally, we compare and contrast the different models in terms of computation time and their effectiveness on a variant of the assignment problem where the dispersion of the assignment costs to complete the jobs across different agents is controlled.

2 Model Development

We begin by introducing notation and recalling a classical result on the equivalence of norms in finite-dimensional spaces [12]. Let $\boldsymbol{x} \in \mathbb{R}^n_{\geqslant 0}$ denote an n-dimensional non-negative vector, whose dispersion we wish to regulate. For

any integer $p \geqslant 1$, define its ℓ^p-norm by

$$\|\boldsymbol{x}\|_p \;=\; \Big(\sum_{i=1}^{n} x_i^p\Big)^{1/p}, \tag{1}$$

Here we omit the absolute value in (1) since $x_i \geqslant 0$ for all i. Also define

$$D_p \;\triangleq\; n^{1-1/p} - 1. \tag{2}$$

Then, for any pair of integers p_1 and p_2, such that $0 < p_1 < p_2$, the norm-equivalence result [12] states

$$\mathcal{NE}(p_1, p_2): \;\; \|\boldsymbol{x}\|_{p_2} \;\leqslant\; \|\boldsymbol{x}\|_{p_1} \;\leqslant\; \left(\frac{D_{p_2}+1}{D_{p_1}+1}\right) \|\boldsymbol{x}\|_{p_2}. \tag{3}$$

Specializing to the case $(p_1, p_2) = (1, p)$ with $p \geqslant 2$, we obtain

$$\mathcal{NE}(1, p): \;\; \|\boldsymbol{x}\|_p \;\leqslant\; \|\boldsymbol{x}\|_1 \;\leqslant\; (D_p + 1)\, \|\boldsymbol{x}\|_p. \tag{4}$$

Moreover:

- The lower bound $\|\boldsymbol{x}\|_p = \|\boldsymbol{x}\|_1$ is attained if and only if all but one component of $\boldsymbol{x}$ are zero (the most *unfair* distribution of $\boldsymbol{x}$).
- The upper bound $\|\boldsymbol{x}\|_1 = (D_p + 1)\, \|\boldsymbol{x}\|_p$ is attained if and only if all components of $\boldsymbol{x}$ are equal (the most *fair* distribution of $\boldsymbol{x}$).

This intuition naturally leads to the following definition, extending the concept of ε-fairness in [27].

Definition 2.1 (ε, p)-fairness – For any integer $p \geqslant 2$, $\boldsymbol{x} \in \mathbb{R}^n_{\geqslant 0}$ and $\varepsilon \in [0, 1]$, we say that $\boldsymbol{x}$ is (ε, p)-*fair* if

$$\big(1 + \varepsilon\, D_p\big)\, \|\boldsymbol{x}\|_p \;=\; \|\boldsymbol{x}\|_1.$$

Furthermore, we say that $\boldsymbol{x}$ is *at least* (ε, p)-*fair* if

$$\big(1 + \varepsilon\, D_p\big)\, \|\boldsymbol{x}\|_p \;\leqslant\; \|\boldsymbol{x}\|_1. \tag{5}$$

Additionally, we denote the set of at least (ε, p)-fair vectors as

$$\mathcal{X}(\varepsilon, p) \;\triangleq\; \big\{\boldsymbol{x} \in \mathbb{R}^n_{\geqslant 0} : (1 + \varepsilon\, D_p)\, \|\boldsymbol{x}\|_p \leqslant \|\boldsymbol{x}\|_1\big\}. \tag{6}$$

In the above definition, the factor $(1 + \varepsilon D_p)$ can be interpreted as the convex combination of 1 and $(D_p + 1)$, with ε as the non-negative multiplier. Furthermore, the constraint in (5) is convex for every $p \geqslant 2$. In particular, for $p = 2$,

(5) reduces to a second-order cone [27]; for $2 < p < \infty$, it defines an ℓ_p-norm cone, and for $p = \infty$, it reduces to a collection of linear inequalities as follows:

$$\|\boldsymbol{x}\|_1 \;\geqslant\; (1 + \varepsilon D_\infty)\,\|\boldsymbol{x}\|_\infty \iff \sum_{i=1}^{n} x_i \;\geqslant\; (1 + n\varepsilon - \varepsilon)\cdot x_i \quad \forall\, i. \tag{7}$$

For any fixed ε, the Definition 2.1 yields a countably infinite family of convex models (by varying the integer $p \geqslant 2$) for enforcing fairness in the distribution of $\boldsymbol{x}$. We now present a proposition that directly follows from the definition of $\mathcal{X}(\varepsilon, p)$ in (6).

Proposition 2.1 *For any integer $p \geqslant 2$ and and $\varepsilon_1, \varepsilon_2 \in [0, 1]$ such that $\varepsilon_1 > \varepsilon_2$, $\mathcal{X}(\varepsilon_1, p) \subset \mathcal{X}(\varepsilon_2, p)$.* □

The following normalization argument to simplifies notation for the main theoretical results in the subsequent section.

Remark 2.1 Normalization – Since all finite-dimensional norms are positively homogeneous of degree 1, the inequality in (6), i.e.,

$$\left(1 + \varepsilon\, D_p\right) \|\boldsymbol{x}\|_p \;\leqslant\; \|\boldsymbol{x}\|_1$$

is invariant under scaling of $\boldsymbol{x}$. Thus, it is convenient to normalize any non-zero vector[1] by its ℓ_1-norm. To that end, define

$$\Delta_n \;\triangleq\; \left\{ \boldsymbol{x} \in \mathbb{R}^n_{\geqslant 0} : \sum_{i=1}^{n} x_i = 1 \right\}, \tag{8}$$

the standard probability simplex in $\mathbb{R}^n$. Normalizing any $\boldsymbol{x} \in \mathcal{X}(\varepsilon, p)$ with respect to its ℓ_1-norm then yields

$$\mathcal{Y}(\varepsilon, p) \;\triangleq\; \left\{ \boldsymbol{x} \in \Delta_n : (1 + \varepsilon\, D_p)\,\|\boldsymbol{x}\|_p \leqslant 1 \right\}. \tag{9}$$

Observe that $\mathcal{Y}(\varepsilon, p)$ is exactly the image of $\mathcal{X}(\varepsilon, p)$ under the map $\boldsymbol{x} \mapsto \boldsymbol{x}/(\boldsymbol{e}^\intercal \boldsymbol{x})$; here $\boldsymbol{e}$ is a vector of 1s. Hence, any theoretical result for $\mathcal{X}(\varepsilon, p)$ can be transferred immediately to $\mathcal{Y}(\varepsilon, p)$ and vice-versa. All theoretical properties of (ε, p)-fairness in the subsequent section will be proved for the normalized set $\mathcal{Y}(\varepsilon, p)$.

[1] The zero vector, $\mathbf{0}$, cannot be normalized and is trivially (ε, p)-fair for any $p \geqslant 2$ and $\varepsilon \in [0, 1]$. Hence, throughout the rest of the note, we assume that $\boldsymbol{x}$ is always a non-zero vector.

3 Theoretical Properties

We first establish the relationship between the (ε, p)-fairness of Definition 2.1 and the coefficient of variation of the vector $\boldsymbol{x}$. The coefficient of variation of $\boldsymbol{x}$ is a statistical measure of the dispersion of $\boldsymbol{x}$. The result in the subsequent section 3.1 provides the intuition for how varying ε in (5) for any fixed integer $p \geqslant 2$ allows the control of dispersion of $\boldsymbol{x}$.

3.1 Bound of Coefficient of Variation

The first property we establish is that the (ε, p)-fairness constraint (5) implies an upper bound on the coefficient of variation of $\boldsymbol{x}$, and that this bound tightens as ε increases.

Proposition 3.1 *For any $\varepsilon \in [0,1]$, an integer $p \geqslant 2$, and $\boldsymbol{x} \in \mathcal{Y}(\varepsilon, p)$ defined in (9), the coefficient of variation of $\boldsymbol{x}$*

$$\mathrm{CV}(\boldsymbol{x}) \triangleq \frac{\sqrt{\frac{1}{n}\sum_{i=1}^{n}(x_i-\mu)^2}}{\mu}, \qquad \textit{where} \qquad \mu = \frac{1}{n},$$

satisfies the following bound:

$$\left(\mathrm{CV}(\boldsymbol{x})\right)^2 \;\leqslant\; B_p(\varepsilon) \;\triangleq\; \frac{(D_p+1)^2}{(1+\varepsilon\, D_p)^2} - 1\,. \tag{10}$$

Moreover, $B_p(\varepsilon)$ is strictly decreasing in ε on $[0,1]$, with $B_p(0) = (D_p+1)^2 - 1$ and $B_p(1) = 0$.

Proof For any $p \geqslant 2$, using norm equivalence $\mathcal{NE}(2,p)$ in (3) and (9), we have

$$\|\boldsymbol{x}\|_2 \;\leqslant\; \frac{D_p+1}{\sqrt{n}}\|\boldsymbol{x}\|_p \;\leqslant\; \frac{D_p+1}{\sqrt{n}}\left(\frac{1}{1+\varepsilon\, D_p}\right). \tag{11}$$

On the other hand, since $n\mu = 1$ and observing that

$$\frac{1}{n}\sum_{i=1}^{n} x_i^2 = \frac{\|\boldsymbol{x}\|_2^2}{n}, \quad \frac{1}{n}\sum_{i=1}^{n}(x_i-\mu)^2 = \frac{\|\boldsymbol{x}\|_2^2}{n} - \mu^2,$$

we obtain

$$\left(\mathrm{CV}(\boldsymbol{x})\right)^2 = \frac{\|\boldsymbol{x}\|_2^2 - n\mu^2}{n\mu^2} = n\,\|\boldsymbol{x}\|_2^2 - 1.$$

Substituting the bound on $\|\boldsymbol{x}\|_2$ from (11) yields

$$n\,\|\boldsymbol{x}\|_2^2 - 1 \;\leqslant\; \frac{(D_p+1)^2}{(1+\varepsilon\,D_p)^2} - 1.$$

Hence,

$$(\mathrm{CV}(\boldsymbol{x}))^2 \;\leqslant\; \frac{(D_p+1)^2}{(1+\varepsilon\,D_p)^2} - 1 = B_p(\varepsilon).$$

Finally, since $D_p > 0$, we have

$$B_p'(\varepsilon) \;=\; -\,\frac{2D_p(D_p+1)^2}{(1+\varepsilon\,D_p)^3} \;<\; 0.$$

Hence, $B_p(\varepsilon)$ is strictly decreasing in ε and $B_p(0) = (D_p+1)^2 - 1$, $B_p(1) = 0$, as claimed.

The strict monotonicity of $B_p(\varepsilon)$ in ε implies that, as ε increases, the upper bound on $\mathrm{CV}(\boldsymbol{x})$ decreases, thereby decreasing the dispersion of $\boldsymbol{x}$. This relationship was previously established for the case $p = 2$ in [27].

In the forthcoming sections, we present the main theoretical results that concern the equivalence or ordered relation between sets $\mathcal{Y}(p,\varepsilon)$ and correspondingly of $\mathcal{X}(p,\varepsilon)$ for varying values of p.

3.2 Equivalence Condition

In this section, we give conditions under which requiring $\boldsymbol{x}$ to be at least (ε,p)-fair is equivalent for every integer $p \geqslant 2$. We start with a key result that directly follows from the definition of (ε,p)-fairness.

Proposition 3.2 *Suppose that $\boldsymbol{e}_i$ is a unit vector with the i^{th} coordinate set to 1 and $\boldsymbol{e}$ is a vector of 1s, then for any integer $p \geqslant 2$, the following statements hold:*

$$\textit{For any } i \in \{1,\ldots,n\},\quad \boldsymbol{e}_i \in \mathcal{Y}(\varepsilon,p) \implies \varepsilon = 0 \tag{12}$$

$$\textit{and } \varepsilon = 1 \iff \mathcal{Y}(\varepsilon,p) = \left\{\tfrac{1}{n}\boldsymbol{e}\right\}. \tag{13}$$

Proof The proof follows from the observation in Section 2 that the term $(1+\varepsilon D_p)$ is the convex combination of the two scalars 1 and (D_p+1) with ε as the factor.

Theorem 3.1 *For any integer $p \geqslant 2$, and $\varepsilon \in \{0,1\}$, the sets of at least (ε, p)-fair vectors are equivalent, i.e.*

$$\mathcal{Y}(\varepsilon, p) = \mathcal{Y}(\varepsilon, 2) \quad \forall\ p > 2.$$

Proof We treat the two cases $\varepsilon = 0$ and $\varepsilon = 1$ separately. When $\varepsilon = 0$, the inequality in (9) is trivially satisfied for any $\boldsymbol{x} \in \Delta_n$ and any $p \geqslant 2$. Hence,

$$\mathcal{Y}(0, p) = \Delta_n = \mathcal{Y}(0, 2) \quad \forall, p > 2.$$

When $\varepsilon = 1$, Proposition 3.2 indicates that $\mathcal{Y}(1, p) = \left\{\frac{1}{n}\boldsymbol{e}\right\}$ for every $p \geqslant 2$.

Theorem 3.1 shows that at the two extreme fairness levels, the feasible set of vectors $\boldsymbol{x}$ is the same regardless of the choice of p.

3.3 Strict Inclusion Property

Here, we prove a result that shows a strict inclusion property between $\mathcal{Y}(\varepsilon, p_1)$ and $\mathcal{Y}(\varepsilon, p_2)$ for any pair of integers (p_1, p_2) such that $2 \leqslant p_1 < p_2$, and ε in the open interval $(0, 1)$. We begin by presenting a proposition that will be useful to prove the main theorem.

Proposition 3.3 *For any $\boldsymbol{x} \in \Delta_n \setminus \left\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n, \frac{1}{n}\boldsymbol{e}\right\}$,*

$$F(p) \triangleq \frac{1 - \|\boldsymbol{x}\|_p}{D_p \|\boldsymbol{x}\|_p} \text{ is strictly decreasing in } p \geqslant 2. \tag{14}$$

Proof See Appendix A.

Theorem 3.2 *For any pair of integers (p_1, p_2) such that $2 \leqslant p_1 < p_2$ and $\varepsilon \in (0, 1)$, $\mathcal{Y}(\varepsilon, p_2) \subset \mathcal{Y}(\varepsilon, p_1)$.*

Proof For each $\boldsymbol{x} \in \Delta_n$ and integer $p \geqslant 2$, define

$$\varepsilon_p(\boldsymbol{x}) \ \triangleq\ \frac{1 - \|\boldsymbol{x}\|_p}{D_p \, \|\boldsymbol{x}\|_p}. \tag{15}$$

Since

$$(1 + \varepsilon\, D_p)\, \|\boldsymbol{x}\|_p \leqslant 1 \iff \varepsilon \leqslant \varepsilon_p(\boldsymbol{x})$$

we have

$$\boldsymbol{x} \in \mathcal{Y}(\varepsilon, p) \iff 0 < \varepsilon \leqslant \varepsilon_p(\boldsymbol{x}).$$

Therefore $\mathcal{Y}(\varepsilon, p_2) \subset \mathcal{Y}(\varepsilon, p_1)$ will follow once we show

$$\varepsilon_{p_1}(\boldsymbol{x}) > \varepsilon_{p_2}(\boldsymbol{x}) \quad \forall \quad \boldsymbol{x} \in \Delta_n \setminus \left\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n, \tfrac{1}{n}\boldsymbol{e}\right\}. \tag{16}$$

The excluded cases correspond to the boundary situation $\varepsilon \in \{0, 1\}$, already treated in Proposition 3.2 and Theorem 3.1. For any such $\boldsymbol{x} \in \Delta_n \setminus \left\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n, \frac{1}{n}\boldsymbol{e}\right\}$, the proof is complete if we show that the map

$$p \mapsto \varepsilon_p(\boldsymbol{x}) = \frac{1 - \|\boldsymbol{x}\|_p}{D_p \|\boldsymbol{x}\|_p}$$

is *strictly decreasing* on $[2, \infty]$ which is exactly the statement of Proposition 3.3. Figure 1 illustrates the idea graphically.

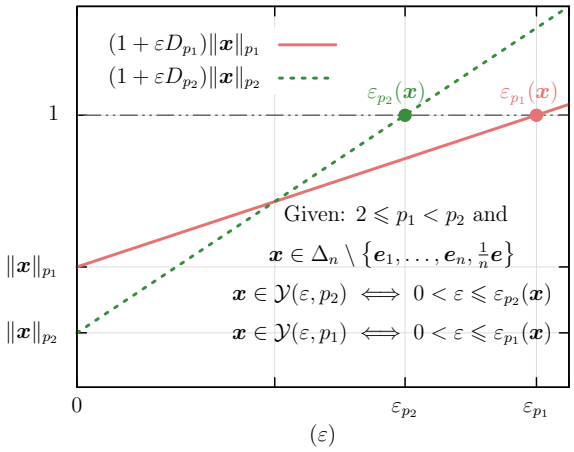

Fig. 1: Graphical illustration for $\varepsilon_{p_1}(\boldsymbol{x}) > \varepsilon_{p_2}(\boldsymbol{x})$.

Theorem 3.2 shows that the size of the feasible sets of requiring $\boldsymbol{x}$ to be at least (ε, p)-fair decreases with p when $\varepsilon \in (0, 1)$ and Theorem 3.1 shows that the feasible sets are equal for any value of integer $p \geqslant 2$ when $\varepsilon \in \{0, 1\}$. Next, we consider an example application to illustrate the use of the models presented thus far in this note.

4 Multi-Agent Assignment Problem

In this section, we formulate a basic version of a multi-agent assignment problem (MAAP) and incorporate the fairness models introduced in Section 2 to

control the dispersion of the total time required to complete the tasks across different agents. We start by introducing the relevant notation. Let $\mathcal{A}$ denote the set of agents and $\mathcal{J}$ the set of tasks. We assume that $|\mathcal{J}| > |\mathcal{A}|$. Each task $j \in \mathcal{J}$ must be completed by exactly one agent. For any agent $a \in \mathcal{A}$ and task $j \in \mathcal{J}$, let $t_{a,j}$ denote the time taken by agent a to complete task j. The multi-agent assignment problem seeks to assign each task $j \in \mathcal{J}$ to one of the agents in $\mathcal{A}$ to minimize the total completion time. Furthermore, we introduce binary variables

$$x_{a,j} = \begin{cases} 1, & \text{if task } j \text{ is assigned to agent } a, \\ 0, & \text{otherwise.} \end{cases}$$

Using these binary decision variables, a Mixed-Integer Linear Program (MILP) for the MAAP is as follows:

$$\min: \quad \sum_{a \in \mathcal{A}} t_a \tag{17a}$$

$$\text{subject to:} \quad t_a = \sum_{j \in \mathcal{J}} t_{a,j} \cdot x_{a,j} \quad \forall a \in \mathcal{A}, \tag{17b}$$

$$\sum_{a \in \mathcal{A}} x_{a,j} = 1 \quad \forall j \in \mathcal{J}, \tag{17c}$$

$$x_{a,j} \in \{0,1\} \quad \forall a \in \mathcal{A},\ j \in \mathcal{J}. \tag{17d}$$

Although the formulation in (17) is presented as an MILP, its constraint matrix is the node–arc incidence matrix of a bipartite graph, which is totally unimodular. Consequently, the linear programming relaxation has integral extreme points, and solving it via any linear programming solver yields the optimal integer solution [25]. Equivalently, one can apply the Hungarian algorithm directly to obtain the optimum in $O(|\mathcal{J}|^3)$ time [17]. In contrast, a natural generalization in which each task carries a weight and each agent has a capacity constraint becomes NP-hard (the generalized assignment problem) [6]. In this work, we restrict attention to the uncapacitated variant in (17), focusing on illustrating how (ε, p)-fairness models control the dispersion of the processing times t_a across agents in the simplest assignment problem, rather than a general variant of the MAAP.

4.1 Fairness-Constrained Multi-Agent Assignment Problem

We let $\boldsymbol{t}$ denote the vector of t_a values over all agents $a \in \mathcal{A}$. The fairness-constrained MAAP seeks to control the dispersion of the vector $\boldsymbol{t}$. This is enforced by requiring $\boldsymbol{t}$ to be at least (ε, p)-fair for any $\varepsilon \in (0, 1]$ and $p \geqslant 2$, as in (5). When $\varepsilon = 0$, constraint (5) becomes trivial and the problem reduces to (17). The fairness constraint for fixed $p \geqslant 2$ and $\varepsilon \in (0, 1]$ is

$$(1 + \varepsilon D_p) \, \|\boldsymbol{t}\|_p \;\leqslant\; \sum_{a \in \mathcal{A}} t_a. \tag{18}$$

Although the problem in (17) is solvable in polynomial time, adding (18) makes the resulting problem NP-complete. Indeed, for two agents with $\varepsilon = 1$, it reduces to the classical partition problem [16] that is known to be strongly NP-complete. As noted in Section 2, when $p < \infty$, (18) is a convex constraint—specifically, an ℓ_p-norm conic constraint (for $p = 2$, an ℓ_2-norm cone is referred to as a second-order cone). Hence, the fairness-constrained MAAP is a mixed-integer convex program for $p < \infty$. When $p = \infty$, applying the reduction in (7) to (18) yields a mixed-integer linear program that can be solved directly using off-the-shelf commercial or open-source mixed-integer linear programming solvers. For mixed-integer convex optimization problems ($2 \leqslant p < \infty$), constraint (18) can be lazily outer approximated by linear inequalities [7], enabling MILP solvers to obtain the solution. In the next section, we briefly present the outer approximation procedure for a general ℓ_p-norm cone. We remark that the outer approximation algorithm is well-understood in the literature, and the presentation in the next section is only intended to make the article self-contained for the readers.

4.2 Outer Approximation of ℓ_p-Norm Cone

We start with the definition of an ℓ_p-norm cone [5] of dimension $(n + 1)$ as

$$P = \left\{ (z, \boldsymbol{y}) \in \mathbb{R}^{n+1} : z \geqslant \|\boldsymbol{y}\|_p \right\}. \tag{19}$$

Note that the fairness constraint in (18) for any $p \geqslant 2$ can be represented in the form (19) by introducing an auxiliary variable. It is known in the literature that when convex cones are outer approximated by linear inequalities, it is both theoretically stronger and computationally more efficient to decompose a high-dimensional cone into several lower-dimensional cones and approximate each of those individually, rather than approximating the high-dimensional

cone directly [19]. Converting a high-dimensional conic constraint into many low-dimensional cones is called "cone disaggregation" [19]. To that end, we present the steps to disaggregate the $(n+1)$-dimensional ℓ_p-norm cone in (19) into n three-dimensional cones plus a single linear constraint:

$$\begin{aligned} z^p \;\geqslant\; \sum_{i=1}^{n} y_i^p \quad &\equiv \quad z \;\geqslant\; \sum_{i=1}^{n} \frac{y_i^p}{z^{p-1}}, \\ \equiv \quad &\left\{ z = \sum_{i=1}^{n} Y_i \text{ and } Y_i\, z^{p-1} \geqslant y_i^p \;\; \forall\, i \right\} \\ \equiv \quad &\left\{ z = \sum_{i=1}^{n} Y_i \text{ and } Y_i^{\frac{1}{p}}\, z^{1-\frac{1}{p}} \geqslant y_i \;\; \forall\, i \right\}. \end{aligned}$$

In the last equivalence, each constraint $Y_i^{\frac{1}{p}}\, z^{1-\frac{1}{p}} \geqslant y_i$ is convex for $2 \leqslant p < \infty$ and in fact defines a three-dimensional power cone [5]:

$$\mathcal{P}_3^{\alpha,1-\alpha} \triangleq \left\{ (v_1, v_2, v_3) \in \mathbb{R}^3_{\geqslant 0} : v_1^{\alpha} v_2^{1-\alpha} \geqslant v_3 \right\}. \tag{21}$$

Hence, the cone in (19) can be rewritten as

$$\left\{ z = \sum_{i=1}^{n} Y_i \text{ and } (Y_i, z, y_i) \in \mathcal{P}_3^{1/p,1-1/p} \;\; \forall\, i \right\}, \tag{22}$$

which comprises n three-dimensional power cones and one linear equality. A linear outer approximation of each power cone is obtained by the first-order Taylor expansion at a point (Y_i^0, z^0, y_i^0) on its surface (i.e., satisfying the conic constraint at equality). Specifically, the linear outer approximation of $(Y_i, z, y_i) \in \mathcal{P}_3^{1/p,1-1/p}$ at (Y_i^0, z^0, y_i^0) is

$$\left(\frac{1}{p}\right) y_i^0\, z^0\, Y_i \;+\; \left(1 - \frac{1}{p}\right) y_i^0\, Y_i^0\, z \;\geqslant\; Y_i^0\, z^0\, y_i. \tag{23}$$

In the algorithm to solve the mixed-integer convex optimization problem, we relax the n three-dimensional power cones and iteratively add these linear outer approximations whenever the solution to the relaxed problem violates one of them. The required point on the surface of the power cone is computed by projecting the infeasible solution onto the corresponding power cone.

5 Computational Results

This section provides empirical validation of the theoretical results in Section 3 using the formulations from Section 4. All experiments were run on 50 randomly generated instances of the MAAP, each with 5 agents and 250 tasks; the processing times $t_{a,j}$ were drawn uniformly at random. Both variants–with and without fairness constraints–were implemented in Kotlin and solved using CPLEX 22.1. The outer-approximation algorithm of Section 4.2 was realized via CPLEX's lazy-constraint callback. Source code for instance generation and both solution methods is available at `https://github.com/abhay-singh11/FairTaskAllocation`.

For each instance we solved the base assignment problem (17) augmented by the (ε, p)–fairness constraint (18) for

$$p \in \{2, 3, 5, 10, \infty\}, \quad \varepsilon \in \{0, 0.1, \ldots, 0.9, 0.99, 1.0\}.$$

We found that all 50 instances were infeasible at $\varepsilon = 1.0$, indicating that that perfect fairness cannot be enforced for the instances; we attribute this behaviour to the fact that $\{t_{a,j}\}$ are generated randomly. For this reason, we include $\varepsilon = 0.99$ in our study: at this level, all instances remain feasible for every chosen p.

We let $\boldsymbol{t}^*(\varepsilon, p)$ denote the the optimal solution vector for the MAAP under fixed (ε, p), and $\boldsymbol{t}^*(0)$ for the solution of the unconstrained problem (17). This notation is chosen to explicitly indicate that when $\varepsilon = 0$ in (18), it is equivalent to imposing no fairness constraints on the MAAP.

5.1 Dispersion Control

Here, we present the results for one representative instance out of the 50 randomly generated ones, since the other 49 exhibit the same qualitative behaviour. Figure 2 reports the coefficient of variation $\mathrm{CV}(\boldsymbol{t}^*(\varepsilon, p))$, computed as in Proposition 3.1. First, we observe that $\mathrm{CV}(\boldsymbol{t}^*(\varepsilon, p))$–the measure of dispersion–decreases monotonically with ε. Proposition 3.1 establishes that the upper bound on the coefficient of variation decreases in ε; Figure 2 confirms that the actual $\mathrm{CV}(\boldsymbol{t}^*(\varepsilon, p))$ likewise decreases. Next, we note the trend of diminishing returns: beyond $\varepsilon \approx 0.70$, the marginal reduction in dispersion is negligible, indicating that least fairness gains occur at moderate values of ε.

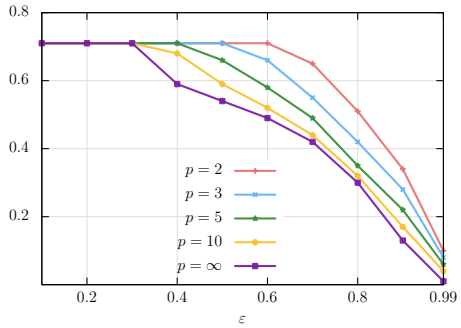

Fig. 2: Coefficient of variation $\mathrm{CV}(\boldsymbol{t}^*(\varepsilon, p))$ versus ε for different p for one instance.

5.2 Price of Fairness

The term "Price of Fairness" (PoF) is defined as the percentage increase in the objective value of the fairness-constrained MAAP of Section 4.1 relative to that of the unconstrained problem (17). This definition is standard in the literature [3,27]. Formally,

$$\mathrm{PoF}(\varepsilon, p) \triangleq \left(\frac{\boldsymbol{t}^*(\varepsilon, p)^\intercal \boldsymbol{e}}{\boldsymbol{t}^*(0)^\intercal \boldsymbol{e}} - 1 \right) \times 100\%.$$

Figure 3 illustrates this trend for one representative instance; the other

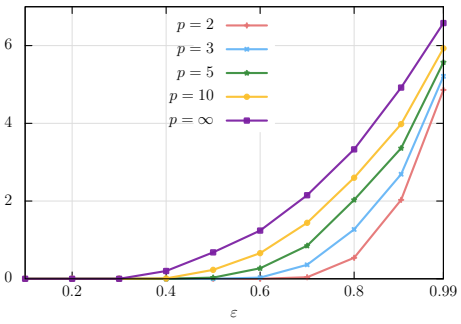

Fig. 3: $\mathrm{PoF}(\varepsilon, p)$ versus ε for various p.

49 instances show the same qualitative behavior. The increasing ordering $\mathrm{PoF}(\varepsilon, p_1) \geqslant \mathrm{PoF}(\varepsilon, p_2)$ for $p_1 > p_2$ reflects the inclusion property of Theorem 3.2.

5.3 Computation Times

The box-plots of computation times are shown in Figure 4. Two trends stand out: (i) The computation time increases with ε. As ε grows, the feasible region shrinks, making it harder for the solver to find and prove feasibility. (ii) For $p = \infty$, solve times are generally lower than for finite p. In that case, the fairness constraint reduces to a set of linear inequalities, whereas for $p < \infty$ the conic constraints incur additional overhead. Furthermore, when $p < \infty$, computation times are observed to increase with p.

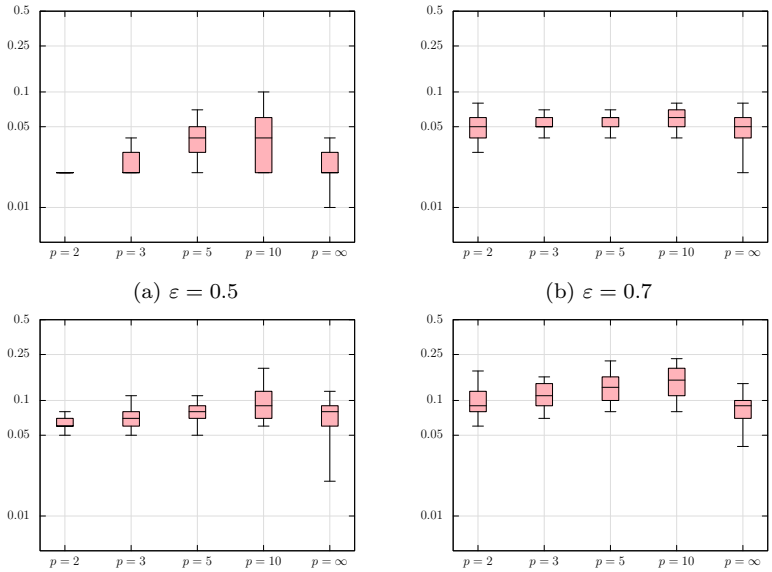

(a) $\varepsilon = 0.5$ (b) $\varepsilon = 0.7$

(c) $\varepsilon = 0.9$ (d) $\varepsilon = 0.99$

Fig. 4: Box-plot of computation times in seconds for all the 50 instances for $\varepsilon \in \{0.5, 0.7, 0.9, 0.99\}$.

6 Conclusion

The note presents a countably infinite family of convex models–parameterized by $\varepsilon \in [0, 1]$–to control the dispersion of a subset of decision variables in an optimization problem. The parameter ε bounds dispersion: when $\varepsilon = 0$, dispersion is unconstrained; when $\varepsilon = 1$, dispersion is forced to zero. This family is indexed by integers $p \geqslant 2$, and all the models in the family are

derived from first principles using the fundamental result of equivalence of finite-dimensional norms. For $p = 2$, the model is a second-order cone; for $2 < p < \infty$, it is an ℓ_p-norm cone; and for $p = \infty$, it reduces to a set of linear inequalities. The note also establishes several theoretical properties that align with intuitive notions of fair solutions. The $p = \infty$ variant is appealing computationally since the model is linear. We then evaluate the trade-offs of these models on the MAAP. In particular, the $p = \infty$ variant is appealing both computationally—since the model is linear—and in terms of dispersion control, as increasing ε yields a more gradual reduction in dispersion. Finally, these models can be integrated into optimization models in many application domains to control dispersion or enforce the fairness of a vector of decision variables with minimal modifications to existing solution algorithms.

Appendix A Proof of Proposition 3.3

We start with introducing additional notation and a few lemmas that will make the presentation of the proof of the main result cleaner. For any vector $\boldsymbol{x} \in \Delta_n \setminus \left\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n, \frac{1}{n}\boldsymbol{e}\right\}$, we let

$$S_p = \sum_{i=1}^{n} x_i^p \;\Rightarrow\; S_p^{\frac{1}{p}} = \|\boldsymbol{x}\|_p \tag{24}$$

and

$$w_i = \frac{x_i^p}{S_p} \;\Rightarrow\; \sum_{i=1}^{n} w_i = 1 \text{ and } \sum_{i=1}^{n} w_i^{\frac{1}{p}} = \frac{1}{\|\boldsymbol{x}\|_p}. \tag{25}$$

Also, applying $\mathcal{NE}(1, p)$ to $\boldsymbol{x}$ yields

$$\|\boldsymbol{x}\|_p < 1 < (D_p + 1)\,\|\boldsymbol{x}\|_p \;\Longrightarrow\; \|\boldsymbol{x}\|_p \in \left((D_p + 1)^{-1}, 1\right) \tag{26}$$

Observe that $\boldsymbol{x} = \boldsymbol{e}_i$ for any $i \in \{1, \ldots, n\}$ if and only if $\|\boldsymbol{x}\|_p = 1$ and $\boldsymbol{x} = \frac{1}{n}\boldsymbol{e}$ if and only if $\|\boldsymbol{x}\|_p = (D_p + 1)^{-1}$. Since these values are excluded from the possible values $\boldsymbol{x}$ can take, the inequality in (26) is strict.

Lemma A.1 *For any $\boldsymbol{x} \in \Delta_n \setminus \left\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n, \frac{1}{n}\boldsymbol{e}\right\}$ and integer $p \geqslant 2$, the following inequality holds*

$$\frac{p}{p-1}\left[\frac{-\ln \|\boldsymbol{x}\|_p}{1 - \|\boldsymbol{x}\|_p}\right] - \ln n \left[1 + \frac{1}{D_p}\right] < 0 \tag{27}$$

Proof We start the proof by observing that the function

$$f(\|\boldsymbol{x}\|_p) = \frac{-\ln \|\boldsymbol{x}\|_p}{1 - \|\boldsymbol{x}\|_p} \quad \text{strictly decreases for } \|\boldsymbol{x}\|_p \in (0, 1).$$

The above result is easy to see by evaluating the derivative and checking that the sign of the derivative is strictly negative. Now, for any fixed $p \geqslant 2$, we define $t_0 = (D_p + 1)^{-1}$. Since $\|\boldsymbol{x}\|_p \mapsto f(\|\boldsymbol{x}\|_p)$ is a strictly decreasing function, for every value of $\|\boldsymbol{x}\|_p > t_0$ we have $f(\|\boldsymbol{x}\|_p) < f(t_0)$, i.e.,

$$\begin{aligned}
-\frac{\ln \|\boldsymbol{x}\|_p}{1 - \|\boldsymbol{x}\|_p} &< -\frac{\ln(D_p + 1)^{-1}}{1 - (D_p + 1)^{-1}} \\
&= \frac{\ln(D_r + 1)}{1 - (D_r + 1)^{-1}} \\
&= \frac{\ln n^{1-\frac{1}{r}}}{1 - (D_r + 1)^{-1}} \\
&= \frac{r-1}{r} \ln n \left[\frac{D_r + 1}{D_r}\right] \\
&= \frac{p-1}{p} \ln n \left[1 + \frac{1}{D_p}\right] \\
\Rightarrow \quad \frac{p}{p-1}\left[\frac{-\ln \|\boldsymbol{x}\|_p}{1 - \|\boldsymbol{x}\|_p}\right] &< \ln n \left[1 + \frac{1}{D_p}\right] \\
\Rightarrow \quad \frac{p}{p-1}\left[\frac{-\ln \|\boldsymbol{x}\|_p}{1 - \|\boldsymbol{x}\|_p}\right] - \ln n \left[1 + \frac{1}{D_p}\right] &< 0.
\end{aligned}$$

We remark that when $\boldsymbol{x} = \frac{1}{n}\boldsymbol{e}$, $\|\boldsymbol{x}\|_p$ evaluates to t_0 and the above inequality is satisfied at equality, and the claim of the lemma does not hold. Hence, $\boldsymbol{x} \neq \frac{1}{n}\boldsymbol{e}$ is assumed in the statement of the lemma.

Lemma A.2 *Suppose* $\boldsymbol{w} = [w_1, \ldots, w_n]^\intercal$, *then for any* $p \geqslant 2$, *the function* $H(\boldsymbol{w}) = -\sum_i w_i \ln w_i$ *satisfies*

$$0 \leqslant H(\boldsymbol{w}) \leqslant -\frac{p}{p-1} \ln \|\boldsymbol{x}\|_p. \tag{28}$$

Proof The definitions in (25) tell us that $\boldsymbol{w} = [w_1, \ldots, w_n]^\intercal$ is a probability vector. Hence, $H(\boldsymbol{w})$ is the Shannon-entropy of $\boldsymbol{w}$, which we know is non-negative [9]. To derive the upper bound, we use the definition of Rényi entropy of order α [24] given below

$$H_\alpha(\boldsymbol{w}) = \frac{1}{1-\alpha} \ln \sum_i w_i^\alpha. \tag{29}$$

The following two facts about $H_\alpha(\cdot)$ are well known: $H_1(\boldsymbol{w}) = H(\boldsymbol{w})$ and $H_\alpha(\cdot)$ is a decreasing function of α. We now let $\alpha = \frac{1}{p}$. Then since $\alpha < 1$ (because $p \geqslant 2$), we have

$$H(\boldsymbol{w}) \leqslant H_{\frac{1}{p}}(\boldsymbol{w}) = \frac{p}{p-1} \ln \sum_{i=1}^{n} w_i^{\frac{1}{p}} = -\frac{p}{p-1} \ln \|\boldsymbol{x}\|_p. \tag{30}$$

The final equality follows from (25).

Lemma A.3 *For any integer $p \geqslant 2$ and $\boldsymbol{x} \in \Delta_n$, given the definitions in* (24), (25) *and Lemma A.2, the following relationship holds*

$$\ln S_p - p\frac{S_p'}{S_p} = H(\boldsymbol{w})$$

where S_p' is the derivative of S_p with respect to p.

Proof We start by computing the derivative of S_p with respect to p as

$$\begin{aligned}
S_p' &= \sum_{i=1}^{n} x_i^p \ln x_i = S_p \sum_{i=1}^{n} w_i \ln x_i \\
\iff p\frac{S_p'}{S_p} &= \sum_{i=1}^{n} w_i \ln x_i^p = \sum_{i=1}^{n} w_i \ln(S_p w_i) \\
\iff p\frac{S_p'}{S_p} &= \ln S_p \left(\sum_{i=1}^{n} w_i \right) + \sum_{i=1}^{n} w_i \ln w_i \\
&= \ln S_p - H(\boldsymbol{w}) \\
\iff H(\boldsymbol{w}) &= \ln S_p - r\frac{S_p'}{S_p}.
\end{aligned}$$

We now present the proof of Proposition 3.3 by performing direct sign analysis of the derivative of $F(p)$. We start by taking logarithm of $F(p)$ as

$$\ln F(p) \triangleq g(p) = \ln(1 - \|\boldsymbol{x}\|_p) - \ln \|\boldsymbol{x}\|_p - \ln D_p. \tag{31}$$

We now let $F'(p)$, $g'(p)$, $\|\boldsymbol{x}\|_p'$ and D_p' denote the derivative with respect to p of $F(p)$, $g(p)$, $\|\boldsymbol{x}\|_p$ and D_p, respectively. Then, we can differentiate (31), we get

$$\frac{F'(p)}{F(p)} = g'(p) = -\frac{\|\boldsymbol{x}\|_p'}{1 - \|\boldsymbol{x}\|_p} - \frac{\|\boldsymbol{x}\|_p'}{\|\boldsymbol{x}\|_p} - \frac{D_p'}{D_p}$$

Simplifying the above equation further, we obtain

$$\begin{aligned}
g'(p) &= -\frac{\|\boldsymbol{x}\|_p'}{\|\boldsymbol{x}\|_p(1-\|\boldsymbol{x}\|_p)} - \frac{D_p'}{D_p} \\
&= \frac{-1}{(1-\|\boldsymbol{x}\|_p)}\left[\frac{1}{p}\frac{S_p'}{S_p} - \frac{1}{p^2}\ln S_p\right] - \frac{\ln n}{p^2}\left[1+\frac{1}{D_p}\right] \\
&= \frac{1}{r^2}\left[\frac{\ln S_r - r\frac{S_r'}{S_r}}{1-\|\boldsymbol{x}\|_r}\right] - \frac{\ln n}{r^2}\left[1+\frac{1}{D_r}\right] \\
&= \frac{H(\boldsymbol{w})}{p^2(1-\|\boldsymbol{x}\|_p)} - \frac{\ln n}{p^2}\left[1+\frac{1}{D_p}\right] \quad \text{(using Lemma A.3)} \\
&\leqslant \frac{1}{p^2}\left(\frac{p}{p-1}\left[\frac{-\ln\|\boldsymbol{x}\|_p}{1-\|\boldsymbol{x}\|_p}\right] - \ln n\left[1+\frac{1}{D_p}\right]\right) \\
&< 0 \quad \text{(using Lemma A.1).}
\end{aligned}$$

The penultimate inequality follows from Lemma A.2. We have proved that $g'(p) < 0$. But $F'(p) = F(p) \cdot g'(p)$ and since $F(p) > 0$ for any $\boldsymbol{x}$ that satisfies the assumptions of the proposition, we obtain $F'(p) < 0$. □

Acknowledgements The authors thank Dirk Lauinger of the MIT Energy Initiative for insightful discussions on the potential for expanding on a family of fairness models.

Conflict of interest

The authors declare that they have no conflict of interest.

References

1. Atkinson, A.B., et al.: On the measurement of inequality. Journal of economic theory **2**(3), 244–263 (1970)
2. Bektaş, T., Letchford, A.N.: Using ℓp-norms for fairness in combinatorial optimisation. Computers & Operations Research **120**, 104975 (2020)
3. Bertsimas, D., Farias, V.F., Trichakis, N.: The price of fairness. Operations research **59**(1), 17–31 (2011)
4. Bhadoriya, A.S., Deka, D., Sundar, K.: Equitable routing–rethinking the multiple traveling salesman problem. arXiv preprint arXiv:2404.08157 (2024)
5. Boyd, S.P., Vandenberghe, L.: Convex optimization. Cambridge university press (2004)
6. Cattrysse, D.G., Van Wassenhove, L.N.: A survey of algorithms for the generalized assignment problem. European journal of operational research **60**(3), 260–272 (1992)
7. Duran, M.A., Grossmann, I.E.: An outer-approximation algorithm for a class of mixed-integer nonlinear programs. Mathematical programming **36**(3), 307–339 (1986)

8. Everitt, B.S., Skrondal, A.: The Cambridge dictionary of statistics. Cambridge university press Cambridge, UK (2010)
9. Gallager, R.G.: Information theory and reliable communication, vol. 588. Springer (1968)
10. Ghomi, E.J., Rahmani, A.M., Qader, N.N.: Load-balancing algorithms in cloud computing: A survey. Journal of Network and Computer Applications **88**, 50–71 (2017)
11. Gini, C.: On the measure of concentration with special reference to income and statistics, colorado college publication. General series **208**(1) (1936)
12. Halmos, P.R.: Finite-dimensional vector spaces. Courier Dover Publications (2017)
13. Jain, R.K., Chiu, D.M.W., Hawe, W.R., et al.: A quantitative measure of fairness and discrimination. Eastern Research Laboratory, Digital Equipment Corporation, Hudson, MA **21**(1), 2022–2023 (1984)
14. Janiar, S.B., Pourahmadi, V.: Deep-reinforcement learning for fair distributed dynamic spectrum access in wireless networks. In: 2021 IEEE 18th Annual Consumer Communications & Networking Conference (CCNC), pp. 1–4. IEEE (2021)
15. Kalai, E.: Proportional solutions to bargaining situations: interpersonal utility comparisons. Econometrica: Journal of the Econometric Society pp. 1623–1630 (1977)
16. Korf, R.E.: A complete anytime algorithm for number partitioning. Artificial Intelligence **106**(2), 181–203 (1998)
17. Kuhn, H.W.: The hungarian method for the assignment problem. Naval research logistics quarterly **2**(1-2), 83–97 (1955)
18. Long, L., Nucci, A.: The hoover index of population concentration: A correction and update. The Professional Geographer **49**(4), 431–440 (1997)
19. Lubin, M., Yamangil, E., Bent, R., Vielma, J.P.: Polyhedral approximation in mixed-integer convex optimization. Mathematical Programming **172**(1), 139–168 (2018)
20. Matyjas, J.D., Kumar, S., Hu, F.: Spectrum sharing in wireless networks: Fairness, efficiency, and security. CRC Press (2016)
21. Mo, J., Walrand, J.: Fair end-to-end window-based congestion control. IEEE/ACM Transactions on networking **8**(5), 556–567 (2002)
22. Nash, J.F., et al.: The bargaining problem. Econometrica **18**(2), 155–162 (1950)
23. Rawls, J.: A theory of justice. In: Applied ethics, pp. 21–29. Routledge (2017)
24. Rényi, A.: On measures of entropy and information. In: Proceedings of the fourth Berkeley symposium on mathematical statistics and probability, volume 1: contributions to the theory of statistics, vol. 4, pp. 547–562. University of California Press (1961)
25. Schrijver, A., et al.: Combinatorial optimization: polyhedra and efficiency, vol. 24. Springer (2003)
26. Shakkottai, S., Srikant, R., et al.: Network optimization and control. Foundations and Trends® in Networking **2**(3), 271–379 (2008)
27. Sundar, K., Deka, D., Bent, R.: A parametric, second-order cone representable model of fairness for decision-making problems. Optimization and Engineering pp. 1–16 (2025)
28. Xinying Chen, V., Hooker, J.N.: A guide to formulating fairness in an optimization model. Annals of Operations Research **326**(1), 581–619 (2023)