
The Price of Opportunity Fairness in Matroid Allocation Problems

Rémi Castera*

Moroccan Center for Game Theory
University Mohammed VI Polytechnic
Rabat, Morocco

Felipe Garrido-Lucero

IRIT, Université Toulouse Capitole
Toulouse, France

Patrick Loiseau

Inria, Fairplay joint team
Palaiseau, France

Simon Mauras

Inria, Fairplay joint team
Palaiseau, France

Mathieu Molina

Inria, Fairplay joint team
Crest, ENSAE
Palaiseau, France

Vianney Perchet

ENSAE, Fairplay joint team
Criteo AI Lab
Palaiseau, France

Abstract

We consider matroid allocation problems under *opportunity fairness* constraints: resources need to be allocated to a set of agents under matroid constraints (which include classical problems such as bipartite matching). Agents are divided into C groups according to a sensitive attribute, and an allocation is opportunity-fair if each group receives the same share proportional to the maximum feasible allocation it could achieve in isolation. We study the Price of Fairness (PoF), i.e., the ratio between maximum size allocations and maximum size opportunity-fair allocations. We first provide a characterization of the PoF leveraging the underlying polymatroid structure of the allocation problem. Based on this characterization, we prove bounds on the PoF in various settings from fully adversarial (worst-case) to fully random. Notably, one of our main results considers an arbitrary matroid structure with agents randomly divided into groups. In this setting, we prove a PoF bound as a function of the (relative) size of the largest group. Our result implies that, as long as there is no dominant group (i.e., the largest group is not too large), opportunity fairness constraints do not induce any loss of social welfare (defined as the allocation size). Overall, our results give insights into which aspects of the problem’s structure affect the trade-off between opportunity fairness and social welfare.

1 Introduction

Allocating scarce resources among agents is a fundamental task in diverse fields such as online markets [20], online advertising [54], the labor market [22], university admissions [3, 38], refugee programs [1, 26, 36], or organ transplants [2]. Traditionally, central planners aim to efficiently compute allocations that maximize some metric of social welfare such as the total number of allocated resources. Unfortunately, optimal allocations that neglect equity considerations often result in disparate treatment and unfair outcomes for legally protected groups of individuals, as documented in domains such as job offerings [50, 63] and online advertising [5, 14], among others.

Matching markets, a prominent instance of resource allocation problems, are especially sensitive to these discrimination concerns. For example, initiatives such as the European Union’s proposed

*This research was partly conducted while at Université Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG.

job-matching platform for migrants [22], and the urgent demands arising from the global refugee crisis [34] are complex challenges where fairness must be accounted for: migrants can belong to different demographic groups defined by sensitive attributes such as age, ethnicity, gender, or wealth, and jobs or resettlement locations must be allocated in a fair manner.

Motivated by these challenges, we consider *matroid allocation problems*, i.e., resource allocation problems where the constraint has a matroid structure [61, Chapter 44]. An instance of our problem is defined by a pair (E, \mathcal{I}) where E is a finite set of agents and \mathcal{I} the set of feasible resource allocations, and such that E is partitioned based on sensitive attributes into $C \in \mathbb{N}$ distinct groups (for the formal definitions, please refer to Section 2). Each feasible solution selects a subset of agents who can simultaneously receive a resource, with each agent receiving at most one. Matroid allocation problems have a theoretical structure that gives tractability while being expressive enough to formulate many important problems. They include our main application of bipartite matching discussed above, but also other allocation problems where fairness is relevant. For instance, the allocation of research funding within a university, where the goal is to fairly distribute limited departmental budgets among professors while respecting diversity criteria such as seniority or gender balance; or the assignment of faculty positions across departments and teams, where office space and departmental capacities must be respected; can both be formulated as matroid allocation problems corresponding respectively to partition and laminar matroids. From a combinatorial optimization perspective, this constitutes a *selection problem* under matroid constraints. We use the term “matroid allocation” to highlight the fairness objective, as we focus on how resources are effectively allocated across demographic groups.

Prior works have tackled fairness challenges in (matroid) allocation problems. Chierichetti et al. [18] study fair matroid allocation problems, and Bandyapadhyay et al. [9] examine fair matching (see a more complete discussion in Appendix A). These works, however, focus on the efficient computation of fair allocations and provide approximation algorithms. In contrast, we focus on the *quality* of the optimal fair allocation. Indeed, imposing fairness constraints may reduce *social welfare* (the total number of resources allocated). To understand the trade-off between fairness and social welfare, we study the metric called *Price of Fairness*, introduced independently in [11] and [15] in the contexts of proportional fairness and equitability, respectively. This metric, formally defined by

$$\text{PoF}(\mathcal{I}) := \frac{\max\{|S| : S \in \mathcal{I}\}}{\max\{|S| : S \in \mathcal{I} \text{ is fair}\}},$$

where $|S|$ denotes the size of the allocation S , provides insights into the scenarios where fairness leads to a degradation of the social welfare.

We focus on a novel notion of fairness that we introduce: **opportunity fairness**. Opportunity fairness draws inspiration from the notion of Equality of Opportunity [42] in machine learning and from Kalai-Smorodinsky fairness [47] in fair division. This fairness notion is particularly adapted to the structure of the allocation problem as it accounts for the inherent capabilities of the agents groups. Formally, an allocation is called opportunity fair if for any two groups c, c' , it satisfies

$$\frac{\text{\# of resources allocated to } c}{\text{Total \# of resources that can be allocated to } c} = \frac{\text{\# of resources allocated to } c'}{\text{Total \# of resources that can be allocated to } c'}.$$

We consider a **large market setting** where the number of resources that can be allocated to each group is large, as is often the case, e.g., in job market platforms. In this situation, integral allocations can be well approximated by fractional—or randomized—allocations (as proved in Section 2.3). Hence, we focus on the PoF computation under *fractional allocations*.

Contributions. The price of opportunity fairness may depend on different features of the problem, in particular the set of feasible allocation and the agents’ group assignment. We prove tight PoF bounds in multiple settings from adversarial to fully random, in line with the beyond-worst-case paradigm [59], which seeks to capture more realistic and nuanced behaviors than worst-case analysis alone:

1. *Polymatroid representation:* We first show that for matroids constraints, the set of feasible per-group allocations can be represented as a *polymatroid*, a multi-set generalization of matroids. We then leverage the polymatroid representation to achieve a simpler characterization for the price of opportunity fairness, which is key for the subsequent analysis (Proposition 3.2).
2. *Adversarial analysis:* When both the group partition and the set of feasible allocations are chosen adversarially, we show that the worst-case PoF is $C - 1$, regardless of the number of agents (Theorem 4.1).

3. *Parametrized families of matroids*: By considering a parametrized family of matroids, we conduct a finer PoF analysis with bounds that interpolate the best case of no loss, $\text{PoF} = 1$, and the worst possible case, $\text{PoF} = C - 1$ (Propositions 4.4 and 4.6).
4. *Semi-random setting*: When agents are randomly partitioned into C groups according to some distribution $p = (p_1, \dots, p_C)$, we characterize the worst-case PoF as a function of $\max_{c \in [C]} p_c$ by reducing a joint infinite-dimensional combinatorial optimization problem to a one-dimensional optimization problem (Theorem 4.8). Remarkably, we show that as long as $\max_{c \in [C]} p_c \leq 1/(C - 1)$, no social welfare loss is incurred under opportunity fairness constraints, in particular suggesting that *the trade-off between fairness and social welfare is not entirely due to the presence of small groups, but rather due to the presence of a dominant group* (Corollary 4.9).
5. *Random graphs model*: Finally, we extend the no-social-welfare-loss result of the previous case to any groups partition distribution p whenever the set of feasible allocations \mathcal{I} is obtained from certain Erdős-Rényi random graphs (Propositions 4.10 and 4.11).

Overall, our results lead to a better understanding of how the structure of both the agents groups and feasible allocations can affect the price of fairness. The main qualitative takeaway is that *for realistic matroid and protected groups instances, opportunity fair allocations incur only a small social welfare loss*. While our main focus is on opportunity fairness, as it is the most relevant fairness notion for the setting we consider, we also provide additional results for other fairness notions in Appendix F.

2 Model

2.1 Matroids and Colored Matroids

Let E be a finite set of agents. We denote by $\mathcal{I} \subseteq 2^E$ a family of **feasible allocations**, where for any allocation $S \in \mathcal{I}$, $e \in S$ represents that agent e got a resource allocated. We assume that (E, \mathcal{I}) is a finite matroid:

Definition 2.1. The pair (E, \mathcal{I}) is a (finite) **matroid** if $\emptyset \in \mathcal{I}$ and the following properties are satisfied: (i) For any $S \in \mathcal{I}$ and $T \subseteq S$, $T \in \mathcal{I}$ (*hereditary property*); and (ii) For any $S, T \in \mathcal{I}$, such that $|S| < |T|$, there exists $e \in T \setminus S$ such that $S \cup \{e\} \in \mathcal{I}$ (*augmentation property*).

Matroids are particularly useful for combinatorial optimization. They are rich enough to describe many allocation problems often encountered in practice, e.g., the bipartite matching (traversal matroids), the allocation of research funding within a university (partition matroids), or the assignment of faculty positions across departments and teams (laminar matroids), mentioned in the introduction (see Appendix B for the definition of each sub-class of matroids). More importantly, due to the augmentation property, a maximal size allocation under matroid constraints can be computed in polynomial time via a greedy algorithm, provided there is a polynomial-time oracle to identify if a set is feasible.

To every matroid (E, \mathcal{I}) , we associate a **rank function** $r : 2^E \rightarrow \mathbb{R}_+$, which maps each $S \subseteq E$ to $r(S) := \max\{|T| \mid T \subseteq S, T \in \mathcal{I}\}$, that is, to the size of the maximum feasible allocation included in S . Basic results of matroid theory [61] show that the rank function is submodular², non-decreasing, and that $0 \leq r(S) \leq |S|$ for any $S \subseteq E$.

We consider matroids in which the set of agents is partitioned into C groups (based on sensitive attributes)—also termed colors [18]. Given a matroid (E, \mathcal{I}) , $C \in \mathbb{N}$, and $(E_c)_{c \in [C]}$ a partition of E into groups, the tuple $((E_c)_{c \in [C]}, \mathcal{I})$ is called a **C -colored matroid** (or simply colored matroid).

Given a colored matroid and a subset of groups $\Lambda \subseteq [C]$, we denote by $r(\Lambda) := r(\cup_{c \in \Lambda} E_c)$ the rank of the corresponding subset of agents. We call the function $r : 2^{[C]} \rightarrow \mathbb{R}_+$ the rank function of the colored matroid. The rank function of the colored matroid inherits all properties from the rank function of the original matroid. In addition, remark that $r([C])$ corresponds to the size of a maximum size allocation within \mathcal{I} , i.e., the maximum social welfare achievable in the corresponding resource allocation problem, while $r(c)$ ³ corresponds to the *opportunity level* of the color (the group) c , i.e., the maximum social welfare when considering only the agents within E_c .

²A set function f is submodular if for all finite sets S and T , $f(S) + f(T) \geq f(S \cap T) + f(S \cup T)$.

³We write $r(c)$ instead of $r(\{c\})$ for convenience, since there is no ambiguity.

2.2 Opportunity Fairness

Given a colored matroid $((E_c)_{c \in [C]}, \mathcal{I})$, we denote $I := \{x \in \mathbb{N}^C \mid \text{there exists } S \in \mathcal{I}, x_c = |S \cap E_c|, \text{ for any } c \in [C]\}$. The set I of **integer feasible group allocations** will be sufficient to find optimal fair allocations. We consider fractional allocations as feasible solutions (see Section 2.3 for the justification): we denote by $M := \text{convex hull}(I)$ the set of **fractional feasible group allocations**.

With this notation, we introduce opportunity fairness and the price of fairness for fractional allocations:

Definition 2.2. A fractional allocation $x \in M$ is **opportunity fair** if for any $c, c' \in [C]$,

$$\frac{x_c}{r(c)} = \frac{x_{c'}}{r(c')}.$$

We denote by F the **set of opportunity fair fractional allocations**. The idea behind this notion is that each group is entitled to a quantity of resources proportional to the maximal quantity this group can get across all feasible allocations.

Example. Consider a refugee resettlement task with n men and m women, each linked to cities in a bipartite graph. At most r_1 men can be matched, and r_2 women can be matched. Let the size of the largest allocation be $r_{1,2} \leq r_1 + r_2$. Under the optimal opportunity fair matching, each group's share is proportional to its isolated capacity, relocating $r_1 r_{1,2} / (r_1 + r_2)$ men and $r_2 r_{1,2} / (r_1 + r_2)$ women.

The social welfare of an allocation x , is simply the total number of resources allocated, that is to say $\sum_{c \in [C]} x_c$. Then the **Price of Opportunity Fairness (PoF)** is defined as

$$\text{PoF}(M) := \frac{\max_{x \in M} \sum_{c \in [C]} x_c}{\max_{x \in F} \sum_{c \in [C]} x_c}.$$

This quantity is always greater than 1 and grows as the fairness constraint becomes more restrictive. In the previous example, the Price of Opportunity Fairness equals 1, indicating no loss in total utility.

2.3 From Fractional Allocations to Approximately Fair Integral Allocations

While fractional allocations are not directly implementable, a natural and practical interpretation of a fractional allocation x is as a randomized integral allocation. Indeed, since M is a convex polytope, every point in M can be expressed as a convex combination of its vertices. Moreover, as M is a polymatroid defined by an integral submodular function, all its vertices are integral [30]. Consequently, the convex coefficient associated with each vertex can be interpreted as the probability mass assigned to the corresponding integral allocation.

We consider the relaxation to fractional allocations for two main reasons. First, when restricted to integer allocations, many resource allocation problems have no opportunity fair allocations besides the empty one. Indeed, even for two colors, whenever $r(1)$ and $r(2)$ are co-prime and $(r(1), r(2)) \notin I$, the only feasible fair integral allocation is to allocate 0 resources to each group (see Appendix C).

Second, by slightly relaxing the fairness constraints, any optimal fair fractional allocation can be implemented in an integral fashion through a specific randomized rounding technique, at a cost that becomes negligibly small in large markets. We first define the relaxed fairness notion:

Definition 2.3. For $\gamma \in [0, 1]$, an allocation $x \in M$ is said to be **γ -opportunity fair** if for any $c, c' \in [C]$, it holds that $x_c / r(c) \geq \gamma \cdot x_{c'} / r(c')$.

Setting $\gamma = 1$ recovers Definition 2.2, while $\gamma = 0$ corresponds to no fairness considerations.

Proposition 2.4. For any fractional opportunity fair maximum size allocation $x \in F$, there exists a random allocation X , such that, $\mathbb{E}[X] = x$ and for all realizations, X is feasible, integral, $\left(1 - \frac{2C}{\min_c r(c)}\right)$ -opportunity fair, and $\|X - x\|_1 \leq C$. Conversely, these bounds are tight up to constants: there exist instances (M, x) and (M', x') such that any rounding scheme X is at most $1 - O(\frac{C}{\min_c r(c)})$ -opportunity fair in the first case, or satisfies $\|X - x'\|_1 \geq \Omega(C)$ in the second case.

Proof Sketch. The argument relies on the polymatroid characterization of M (Proposition 3.2) that will be proven independently in Section 3. Then by Theorem 35 of [30], which guarantees that the extreme points of the intersection of two integral polymatroids remain integral, x can be written as a convex combination of nearby integral allocations. See Appendix E.1 for the proof. \square

We remark that the distribution of X can be computed in polynomial time by using an algorithmic version of Carathéodory's Theorem (see Theorem 6.5.11 of [40]). In a **large market**, that is, when $\min_{c \in [C]} r(c)$ is large, Proposition 2.4 shows that both the gap between $\text{PoF} = r([C])/\|x\|_1$ and $r([C])/\|X\|_1$, and the fairness degradation, become negligible. Hence, it provides the strongest fairness guarantees ex-ante (since $\mathbb{E}[X] = x$ and x is perfectly fair), jointly with an approximately fair integral allocation at low cost ex-post. Such a best-of-both-worlds approach that leverages fractional allocations to design lotteries satisfying ex-ante requirements (here fairness) and additional ex-post properties is widespread in market design [4, 12, 13, 46] and fair division [8, 35], with applications in school choice, housing, and kidney exchange, to name but a few.

Even when considering fractional allocations, one may want to impose approximate fairness constraints only. Interestingly, for any colored matroid, our bounds on PoF (with exact fairness constraint from Definition 2.2) can easily be transposed into bounds for the relaxed setting. Formally, denote by PoF_γ the price of γ -opportunity fairness (i.e., when fairness is defined by Definition 2.3).

Proposition 2.5. *Let M be a C -colored matroid. Then, either $\gamma < \max_{x \in F} \min_{c \in [C]} x_c/r(c)$ and $\text{PoF}_\gamma(M) = 1$, or*

$$\gamma \cdot \text{PoF}(M) \leq \text{PoF}_\gamma(M) \leq \frac{\gamma \cdot \text{PoF}(M)}{1 - (1 - \gamma)^{\frac{(C-1) - \text{PoF}(M)}{C-2}}}.$$

See proof in Appendix E.2. Proposition 2.5 allows us to seamlessly lift any PoF upper and lower bounds to PoF_γ , with $\text{PoF}_\gamma \approx \gamma \text{PoF}$ for small γ . As before, the same randomized rounding applies, at a small cost. Thus, **the rest of the paper will focus on fractional allocation and perfect fairness constraints**, as Propositions 2.4 and 2.5 can handle integrality and relaxed constraint considerations.

2.4 Comparison with proportional fairness and other fairness notions

Proportional fairness, a broadly studied fairness notion in the literature, aims to equalize the ratios $x_c/|E_c|$, i.e., the number of allocated resources to each group relative to their size. In machine learning, proportional fairness corresponds to *demographic parity* [10] while *opportunity fairness* is more closely related to *equality of opportunity*, as it accounts for the inherent quality of the groups.

A key limitation of proportional fairness in the matroid allocation setting is that it is *sensitive to the presence of irrelevant agents*, unlike opportunity fairness. Indeed, whenever the allocation problem possesses agents that cannot be allocated any resources, proportional fairness becomes too constraining to satisfy. For instance, adding irrelevant agents of color c increases $|E_c|$ without increasing the size of the feasible allocations to that group, thereby reducing the ratio $x_c/|E_c|$. To maintain proportional fairness, the allocation to other groups must be reduced, even though the underlying allocation problem remains unchanged.

This highlights a pathological behavior: even under fractional allocations, the *Price of Proportional Fairness* can be unbounded, as adding irrelevant agents drives all fair allocations for other groups to zero. In contrast, as we show throughout the paper, *opportunity fairness is robust to such manipulations* and remains bounded, since it accounts for the structure of the allocation problem. In Appendix F, we further discuss the relationships between different notions of fairness (including weighted and leximin fairness) and their respective prices of fairness, underscoring the relevance of opportunity fairness in our context.

3 Polymatroid Structure and PoF Characterization

This section is devoted to characterizing the price of opportunity fairness of a matroid as a simple combinatorial optimization problem. Our main technique will be the use of polymatroids.

Definition 3.1. The **polymatroid** associated to the submodular function $f : 2^C \rightarrow \mathbb{R}_+$ is the polytope

$$\{x \in \mathbb{R}_+^C \mid \sum_{c \in \Lambda} x_c \leq f(\Lambda), \forall \Lambda \subseteq [C]\}.$$

Polymatroids can be seen as a generalization of matroids, as there is a natural mapping from a matroid on a ground set E to a polymatroid included in $[0, 1]^E$, where a feasible allocation $S \in \mathcal{I}$ is associated to a vector $z \in [0, 1]^E$ with coordinates $z_e = 1$ if $e \in S$ and 0 otherwise. Polymatroids are strictly more general, as coordinates can be larger than 1. Proposition 3.2 shows that there is also a natural relation between colored matroids and polymatroids. Refer to Appendix E.3 for the proof.

Proposition 3.2. Let $((E_c)_{c \in [C]}, \mathcal{I})$ be a colored matroid with rank function r and set of feasible fractional allocations M . Then, M is the polymatroid associated to the function r , i.e.,

$$M = \{x \in \mathbb{R}_+^C \mid \sum_{c \in \Lambda} x_c \leq r(\Lambda), \forall \Lambda \subseteq [C]\}.$$

Note that while the usual natural mapping from matroids to polymatroids is not a surjection (the coordinates must remain bounded by 1, which is not the case of all polymatroids), the mapping $((E_c)_{c \in [C]}, \mathcal{I}) \mapsto M$ from the set of colored matroids to the set of all integral Polymatroids is a surjection. From now on, we will use interchangeably the names feasible fractional allocations, polymatroid, and colored matroid for M .

The set M inherits interesting properties from being a polymatroid. For instance, the Pareto frontier of the Multi-objective Optimization Problem $\max_{x \in M} (x_1, x_2, \dots, x_C)$ corresponds to the set of allocations maximizing social welfare $\sum_{c \in [C]} x_c$ [43]. In particular, the existence of an allocation of maximum size that is opportunity fair is reduced to verifying whether the intersection between the Pareto frontier and the line defined by the opportunity fairness condition is non-empty. Figure 1 illustrates M and P for a C -colored matroid with $C = 2$ and $C = 3$, respectively. Remark that P is simultaneously the Pareto frontier and the set of points which maximize $\sum_{c \in [C]} x_c$.




Figure 1: Examples of the set of fractional feasible allocations M (dark blue solid region) and the Pareto frontier P (light blue region) for $C = 2$ and $C = 3$.

The structure of M yields the following characterization of PoF. Refer to Appendix E.4 for the proof.

Corollary 3.3. The price of opportunity fairness of a polymatroid M is given by,

$$\text{PoF}(M) = \frac{r([C])}{\sum_{c \in [C]} r(c)} \cdot \max_{\Lambda \subseteq [C]} \frac{\sum_{c \in \Lambda} r(c)}{r(\Lambda)}. \quad (1)$$

Remark that the combinatorial optimization problem in (1) is exponential in C . Even though real-life applications typically involve a small number of sensitive attributes (often only 2), applications involving intersectional fairness between different sensitive features may introduce a larger number of colors [14, 49, 56]. The PoF can be computed in time $\text{poly}(C, |E|)$ whenever the underlying matroid possesses a **polynomial-time independence oracle** as is the case of transversal, graphic, partition, laminar and uniform matroids. Refer to Appendix D for the details.

In the following section we leverage Corollary 3.3 to tightly bound the PoF in various settings.

4 Bounding the Price of Fairness

4.1 Adversarial Price of Fairness

Our first result is to show that the price of opportunity fairness is always bounded, with a bound only depending linearly on C , the number of colors of the colored matroid, independent of the number of agents $|E|$ and the number of feasible allocations $|\mathcal{I}|$. Refer to Appendix E.5 for the proof.

Theorem 4.1. For any C -colored matroid M , we have $\text{PoF}(M) \leq C - 1$, and this bound is tight.

Theorem 4.1 implies the following remarkable result.

Corollary 4.2. For any 2-colored matroid M , $\text{PoF}(M) = 1$.

Corollary 4.2 shows that whenever agents are divided in two groups, no social welfare loss is incurred due to the opportunity fairness constraint. For an alternative geometrical proof of Corollary 4.2, please refer to Appendix E.6. As a direct application of Proposition 2.5 and Theorem 4.1, we obtain the following corollary, which follows immediately from the monotonicity of the upper bound in Proposition 2.5 with respect to $\text{PoF}(M)$.

Corollary 4.3. Let $\gamma \in [0, 1]$. For any C -colored matroid M , we have $\text{PoF}_\gamma(M) \leq \gamma \cdot (C - 1)$, and this bound is tight.

Theorem 4.1 raises the question of whether tighter bounds can be obtained by restricting the resource allocation problem to specific subclasses of matroids. However, the bound is in fact tight for any class of matroids that contains either graphic or partition matroids, which implies tightness for transversal and laminar matroids. Refer to Appendix E.5 for the details.

4.2 Parametric Price of Fairness

The worst-case bound derived in the previous section relies on the existence of a specific polymatroid, as outlined in the proof of Theorem 4.1. Essentially, this requires an underlying structure where one group E_c can have a rank that grows arbitrarily large while the ranks of other groups remain bounded. This raises the question of whether more favorable guarantees for the price of opportunity fairness can be attained when all groups exhibit similar ranks. Proposition 4.4 provides an upper bound on PoF based on the ranks of the groups. The proof is detailed in Appendix E.7.

Proposition 4.4. *For any C -colored matroid M , it holds,*

$$\text{PoF}(M) \leq \frac{1}{2} \cdot \frac{\max_{c \in [C]} r(c)}{\min_{c \in [C]} r(c)} + \frac{C}{4} \cdot \left(\frac{\max_{c \in [C]} r(c)}{\min_{c \in [C]} r(c)} \right)^2 + \frac{1}{4C} \cdot \mathbb{1}\{C \text{ odd}\}.$$

Moreover, whenever all groups have the same rank, the resulting bound is tight.

The bound in Proposition 4.4 takes into account the shape of the polytope M . When all colors have the same rank, PoF scales as $C/4$. While this upper bound is smaller than the one stated in Theorem 4.1, the price of fairness remains linear with respect to the number of colors.

To complement the analysis, we consider another geometrical parameter related to the shape of M that can interpolate PoF between 1 and $C/4$. Intuitively, PoF is expected to be low when either there is no competition between groups, or the competition is extremely fierce and no group can be unilaterally allocated resources without damaging the allocation of others. Similar behavior has been observed for the Price of Anarchy in congestion games [21], which approximates to 1 under both light and heavy traffic conditions. In the context of the price of opportunity fairness, the relevant problem complexity measure is the associated *independence index* that we define below.

Definition 4.5. We define the **independence index** of a polymatroid M as $\rho(M) := \frac{r([C])}{\sum_{c \in [C]} r(c)}$.

The independence index measures how close the maximal social welfare $r([C])$ is to the social welfare of the *utopian allocation*, the allocation where each group E_c receives $r(c)$ resources (which, in general, is not a feasible allocation). Note that ρ always falls within the interval $[1/C, 1]$, with $\rho = 1/C$ corresponding to complete competition between groups, and $\rho = 1$ corresponding to full independence between groups. These extreme values of the independence index impose a distinct shape on M , as illustrated in Figure 3.

Proposition 4.6. *Let M be a C -colored matroid. Suppose that for any c, c' in $[C]$, $r(c) = r(c')$. Whenever $\rho \in [1/C, 1/(C-1)]$, $\text{PoF}(M) = 1$. Otherwise,*

$$\text{PoF}(M) \leq \rho \max \left(\frac{C - \lfloor C\rho \rfloor + 1}{C\rho - \lfloor C\rho \rfloor + 1}, C - \lfloor C\rho \rfloor \right) \leq \rho((1 - \rho)C + 1).$$

In addition, the first upper bound is jointly tight in ρ and C .

The proof of Proposition 4.6 is provided in Appendix E.8, where Figure 11 illustrates a tight example of the first upper bound for a transversal matroid. Figure 2 illustrates both upper bounds from Proposition 4.6 for $C = 5$ groups with equal ranks. We observe that in both extremes, PoF tends towards 1. Notice that the second upper bound in Proposition 4.6, when maximized over ρ , aligns with the bound from Proposition 4.4 (when all groups have identical ranks). Therefore, the independence index interpolates the price of opportunity fairness between 1 and an order of $C/4$ when all groups have the same isolated social welfare.

Worst-case analysis results stand out by their robustness. However, the particular matroid examples attaining the upper bounds (even under the extra structural assumptions) are rarely observed in real-life. Due to this, the following sections will be dedicated to analyzing PoF in random settings.

4.3 Semi-Random Price of Fairness (Random Coloring)

Our first random setting considers an adversarial matroid choice with a random group agents partition. Formally, we denote by Δ^C the simplex of dimension C , that is, the set of all vectors $p \in [0, 1]^C$




Figure 2: PoF upper bounds stated in Proposition 4.6 for 5 groups with equal rank, with variable value of the independence index ρ .




Figure 3: Shape of M for extreme values of independence index. Left: $\rho = 1/C$, right: $\rho = 1$.

such that $\sum_{c \in C} p_c = 1$. Given a vector $p \in \Delta^C$ without null entries, we create a random partition of a matroid (E, I) (a coloring of the elements in E) by independently and identically assigning each element $e \in E$ to $c \in [C]$ with probability p_c . We denote by $M(p)$ the polymatroid obtained by the random coloring of the agents in E according to the vector p .

Let $(M_n(p))_{n \in \mathbb{N}}$ be a sequence of C -colored matroids over sets $(E^{(n)})$ such that $|E^{(n)}| = n$, randomly colored according to p , and $(r_n)_{n \in \mathbb{N}}$ the associated sequence of rank functions. For each $c \in [C]$, suppose the following limit exists,

$$R(p_c) := \lim_{n \rightarrow \infty} \frac{\mathbb{E}_{p_c}[r_n(c)]}{n}, \quad (2)$$

where $\mathbb{E}_{p_c}[r_n(c)]/n$ represents the rescaled expected social welfare of group c . Remark that assuming convergence is not particularly restrictive as $\mathbb{E}_p[r_n(c)]/n$ is bounded in $[0, 1]$ and thus, it always admits a converging subsequence. We extend the previous definition to any subset $\Lambda \subseteq [C]$ by,

$$R\left(\sum_{c \in \Lambda} p_c\right) := \lim_{n \rightarrow \infty} \frac{1}{n} \cdot \mathbb{E}_{(p_c)_{c \in \Lambda}} [r_n(\cup_{c \in \Lambda} E_c)].$$

Finally, we will assume that $r_n([C]) = \Omega(n)$, which ensures that the size of the optimal allocation grows with the size of the ground sets $E^{(n)}$. The semi-random model represents settings where agents' protected attributes are independent of their individual capabilities or opportunities, i.e., the matroid structure. Random coloring therefore isolates the impact of feasibility constraints and group size, relative to the protected attributes, in generating unfairness.

Remark 1. The use of a sequence of colored matroids $(M_n(p))_{n \in \mathbb{N}}$ is mainly a technical device to capture the large-scale limit. Conceptually, this can be viewed as selecting a large adversarial matroid (whose rank grows roughly proportionally to the size of its ground set) and then randomly coloring its elements according to p .

We first show that the price of opportunity fairness for large colored matroids is completely characterized by the function R .

Proposition 4.7. *If $\liminf_{n \rightarrow \infty} \frac{r_n([C])}{n} > 0$, then*

$$\text{PoF}(M_n(p)) \xrightarrow[n \rightarrow \infty]{P} \max_{\Lambda \subseteq [C]} \frac{R(1)}{\sum_{c \in [C]} R(p_c)} \cdot \frac{\sum_{c \in \Lambda} R(p_c)}{R(\sum_{c \in \Lambda} p_c)}, \quad (3)$$

where \xrightarrow{P} denotes convergence in probability. Moreover, the function R is such that $R(0) = 0$, R is concave, non-decreasing, and 1-Lipschitz. Conversely, any function with these three properties can be realized as the limit of a sequence of limit-matroid derived functions R'_m , i.e. $\|R - R'_m\|_\infty \xrightarrow{m \rightarrow \infty} 0$.

Proof sketch. We first show that, as R is the multi-linear extension of a submodular function, it satisfies the aforementioned properties. Using the concavity of R , we show that whenever $r([C])$ is large, with high probability, $r(c)$ is large as well. Then, by using McDiarmid's concentration inequality, the convergence in probability is concluded. The approximation result is proved by constructing a family of simple functions from specific sequences $(M_n)_n$ whose closed convex hull is equal to the desired set of functions. The full proof is included in Appendix E.9. \square

The first property of Proposition 4.7 shows that upper bounding the right-hand side of Equation (3) yields an upper bound on PoF. The second part shows an equivalence between sequences of C -colored matroids and the set of concave, non-decreasing, 1-Lipschitz functions, which we denote by \mathcal{R} . Therefore, we can shift the problem of bounding the price of opportunity fairness of C -colored matroids to bounding the right-hand side of Equation (3) over all functions in \mathcal{R} . We aim to find a bound that depends only on $\max_{c \in [C]} p_c$. The following theorem provides the exact solution:

Theorem 4.8. Fix $\pi \in [1/C, 1]$ and consider $\Delta_\pi^C := \{p \in \Delta^C \mid \max_{c \in [C]} p_c = \pi\}$ the set of probability distributions with maximum value of π . It follows that

$$\max_{p \in \Delta_\pi^C} \max_{R \in \mathcal{R}} \max_{\Lambda \subseteq [C]} \frac{R(1)}{\sum_{c \in [C]} R(p_c)} \frac{\sum_{c \in \Lambda} R(p_c)}{R(\sum_{c \in \Lambda} p_c)} = \max_{\lambda \in [C]} \psi_\lambda \left(\frac{1 - (C - \lambda)\pi}{C} \right) \leq \min(C - \frac{1}{\pi}, 1), \quad (4)$$

where $\psi_\lambda : [-\lambda, \frac{1}{C}] \rightarrow \mathbb{R}$, for each $\lambda \in [C]$, is given by

$$\psi_\lambda(q) = \begin{cases} \lambda & \text{if } q \in [-\lambda, 0], \\ \frac{\lambda}{(\lambda C q - 1)^2} \cdot (1 + q(1 - 2\lambda) + C(\lambda - 2 + \lambda q)q - 2\sqrt{(\lambda - 1)(C - 1)(1 - Cq)(1 - \lambda q)q}) & \text{if } q \in \left(0, \frac{(\lambda - 1)}{\lambda(C - 1)}\right], \\ 1 & \text{if } q \in \left(\frac{(\lambda - 1)}{\lambda(C - 1)}, \frac{1}{C}\right]. \end{cases}$$

Proof sketch. The left-hand side of Equation (4) defines an infinite-dimensional combinatorial optimization problem, for which classical techniques are difficult to apply. We tackle this by designing transformations that map a generic instance (p, Λ, R) to a new one (p', Λ', R') with a *larger Price of Fairness*. These include linearizing R over subintervals of $[0, 1]$, averaging probabilities in Λ via Karamata's inequality, and modifying Λ so that the maximum-probability coordinate $c^* \in [C]$ of p lies outside it, among others. Iteratively applying these transformations reduces the original problem, for fixed size $\lambda = |\Lambda|$, to a *single-variable* optimization problem solvable via first-order conditions, yielding ψ_λ , the worst-case value for that size. The final result is obtained by maximizing over all λ . See full proof in Appendix E.10. \square

We have reduced a complex optimization problem to a simple closed-form formula with a nice interpretation, the ψ_λ corresponding to the worst-case price of opportunity fairness for a fixed size $\lambda = |\Lambda|$. It may intuitively seem that the Price of Fairness should always be 1, as randomly coloring the ground set is equivalent to first drawing a random number of agents per group and then placing them uniformly on the ground set, seemingly ensuring group-independent opportunities. Nonetheless, a striking conclusion from Theorem 4.8 is that **the price of opportunity fairness can still exceed 1, even when all agents are treated identically**. In particular, whenever $\max_{c \in [C]} p_c$ is larger than $1/2$, the worst-case PoF is at least $C - 2$, as observed in Figure 4.

On the other hand, Theorem 4.8 provides meaningful PoF bounds whenever $\max_{c \in [C]} p_c \leq 1/2$. More importantly, Theorem 4.8 immediately implies the following corollary.

Corollary 4.9. Whenever $\max_{c \in [C]} p_c \leq 1/(C - 1)$, $\text{PoF}(M_n)$ converges in probability towards 1.

For matroid allocation problems in large markets, there is no loss of social welfare due to opportunity fairness as long as no group is overrepresented. This is quite striking as this may contradict the intuition that unfairness stems from the presence of small protected groups that must be catered to, sacrificing the welfare of larger groups. The above corollary shows that even with the presence of an arbitrarily small group, there might be no social welfare loss when being fair. Instead, it is the presence of a single overwhelming group which makes resources hard to fairly allocate, for the specific notion of opportunity fairness.




Figure 4: Theorem 4.8 for $C = 5$ groups: worst-case PoF, ψ_λ for $\lambda \in \{3, 4\}$, and the relaxed bound $C - 1/\pi$.

We have shown how to bound the price of opportunity fairness in the semi-random setting by reducing the combinatorial optimization problem to make it tractable. However, considering the underlying matroid (E, \mathcal{I}) to be fixed may still be a pessimistic assumption for some real-world applications. For this reason, we study next a setting where both the matroid and the colors are drawn randomly.

4.4 Random Graphs Price of Fairness

A first approach to generating random matroids is to uniformly sample one among the $2^{2^n - O(\log n)}$ possible matroids on a ground set of size n , but this would mix different resource allocation problems. Instead, we focus on well-studied subclasses of matroids with established random models: graph-based matroids derived from Erdős-Rényi random graphs.

We consider *graphic matroids* on general graphs, where the ground set consists of edges, the independent sets are acyclic subgraphs, and a maximal allocation corresponds to the largest forest; and *transversal matroids* on bipartite graphs, where the ground set consists of nodes, the independent sets are bipartite matchings, and a maximal allocation corresponds to the largest matching —see Appendix B for the formal matroid class definitions. We show that whenever the edge probability between any two nodes is constant (i.e., independent of n), the price of opportunity fairness converges to 1 with high probability in both cases. Furthermore, by leveraging literature results on the structure of random graphs [37], these findings can be extended to settings with non-constant edge probability.

Graphic matroid: Given $n \in \mathbb{N}$ and $q \in [0, 1]$, we consider the Erdős-Rényi random graph $G_{n,q} := ([n], A)$ with n nodes such that for any $i, j \in [n]$, the edge (i, j) belongs to A independently with probability q . Given a random graph $G_{n,q} = ([n], A)$ and $p \in \Delta^C$, we consider a p **randomly colored random graphic matroid**, denoted $G_{n,q}(p)$. Remark that the random coloring process and the random edges connections are done independently from one another.

Proposition 4.10. *Let $\omega = \omega(n)$ be a function such that $\omega(n) \rightarrow \infty$. Whenever $q \leq 1/(\omega n)$ or $q \geq \omega/n$, for any $p \in \Delta^C$, $\text{PoF}(G_{n,q}(p))$ converges to 1 with high probability as n grows.*

Thus, except near the critical threshold $q = \Theta(1/n)$ marking the transition from a forest to a giant component, almost all sufficiently large random graphic matroids incur no utility loss under opportunity fairness.

Transversal matroid (matching): Given $n \in \mathbb{N}$, $\beta \in (0, 1)$ such that $\beta n \in \mathbb{N}$, and $q \in [0, 1]$, we consider the random bipartite graph $B_{n,\beta,q} := ([n], [\beta n], A)$, where for any $i \in [n], j \in [\beta n]$, the edge $(i, j) \in A$ independently with probability q . Given a random Erdős-Rényi bipartite graph $B_{n,\beta,q}$ and $p \in \Delta^C$, we consider a p **randomly colored random transversal matroid** denoted $B_{n,\beta,q}(p)$. Recall, the coloring process and the edges are drawn independently between them.

Proposition 4.11. *Let $\omega = \omega(n)$ be a function such that $\omega(n) \rightarrow \infty$ arbitrarily slow as $n \rightarrow \infty$. Whenever $q \leq 1/(\omega n^{3/2})$ or $q \geq \omega \log(n)/n$, for any $p \in \Delta^C$, $\text{PoF}(B_{n,\beta,q}(p))$ converges to 1 with high probability as n grows.*

As with graphic matroids, imposing opportunity fairness on random bipartite matching problems yields no utility loss when the edge probability is sufficiently small or large. Unlike the graphic case, however, the behavior in the intermediate regime near the phase transition remains unclear.

5 Conclusions and Limitations

We address matroid allocation problems under a novel group-fairness notion—opportunity fairness—, and prove tight bounds for the Price of Fairness, i.e., the loss of social welfare due to the fairness restrictions, in multiple settings from adversarial to fully random. Our model has two main limitations. First, integral allocations are only well approximated by fractional ones in large markets, which corresponds well to our motivating examples (e.g., job market) but may not always hold. Second, we considered the allocation cardinality as our notion of social welfare, i.e., each allocation has the same weight. The extension to weighted matching is straightforward if weights are assigned at the level of groups as it simply skews the polymatroid. However, individual-level weights would break the anonymity property and the polymatroid nature of M , hence new techniques would be required.

Acknowledgments

This work has been supported by the French National Research Agency (ANR) through grant ANR-20-CE23-0007, ANR-23-CE23-0002, and PEPR IA FOUNDRY project (ANR-23-PEIA-0003). This work has also been funded by the European Union through the ERC OCEAN 101071601 grant. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible.

References

- [1] N. Ahani, P. Gözl, A. D. Procaccia, A. Teytelboym, and A. C. Trapp. Dynamic placement in refugee resettlement. In *ACM Conference on Economics and Computation*, 2021.
- [2] M. Akbarpour, J. Combe, Y. He, V. Hiller, R. Shimer, and O. Tercieux. Unpaired kidney exchange: Overcoming double coincidence of wants without money. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 465–466, 2020.
- [3] M. Akbarpour, A. Kapor, C. Neilson, W. Van Dijk, and S. Zimmerman. Centralized school choice with unequal outside options. *Journal of Public Economics*, 210:104644, 2022.
- [4] M. Akbarpour and A. Nikzad. Approximate random allocation mechanisms. *Microeconomics: Welfare Economics & Collective Decision-Making eJournal*, 2019.
- [5] M. Ali, P. Sapiezynski, M. Bogen, A. Korolova, A. Mislove, and A. Rieke. Discrimination through optimization: How facebook’s ad delivery can lead to biased outcomes. *Proceedings of the ACM on human-computer interaction*, 3(CSCW):1–30, 2019.
- [6] B. Alspach. The wonderful walecki construction. *Bull. Inst. Combin. Appl*, 52(52):7–20, 2008.
- [7] M. Aprile and S. Fiorini. Regular matroids have polynomial extension complexity. *Mathematics of Operations Research*, 47:540–559, 2022.
- [8] H. Aziz. Simultaneously achieving ex-ante and ex-post fairness. *ArXiv*, abs/2004.02554, 2020.
- [9] S. Bandyapadhyay, F. V. Fomin, T. Inamdar, and K. Simonov. Proportionally fair matching with multiple groups. In *Graph-Theoretic Concepts in Computer Science: 49th International Workshop*, page 1–15, 2023.
- [10] S. Barocas, M. Hardt, and A. Narayanan. *Fairness and Machine Learning: Limitations and Opportunities*. MIT Press, 2023.
- [11] D. Bertsimas, V. F. Farias, and N. Trichakis. The price of fairness. *Oper. Res.*, 59:17–31, 2011.
- [12] A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic theory*, 100:295–328, 2001.
- [13] E. Budish, Y.-K. Che, F. Kojima, and P. R. Milgrom. Designing random allocation mechanisms: Theory and applications. *The American Economic Review*, 103:585–623, 2013.
- [14] J. Buolamwini and T. Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, pages 77–91. PMLR, 23–24 Feb 2018.
- [15] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. The efficiency of fair division. *Theory of Computing Systems*, pages 589 – 610, 2012.
- [16] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum nash welfare. 2019.
- [17] R. Castera, P. Loiseau, and B. S. Pradelski. Statistical discrimination in stable matchings. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, EC ’22, page 373–374, New York, NY, USA, 2022. Association for Computing Machinery.
- [18] F. Chierichetti, R. Kumar, S. Lattanzi, and S. Vassilvitskii. Matroids, matchings, and fairness. In *International Conference on Artificial Intelligence and Statistics*, 2019.
- [19] Y. Chun. *Fair queueing*. Springer, 2016.
- [20] M. G. Coles and E. Smith. Marketplaces and matching. *International Economic Review*, 39:239–254, 1998.
- [21] R. Colini-Baldeschi, R. Cominetti, P. Mertikopoulos, and M. Scarsini. When is selfish routing bad? the price of anarchy in light and heavy traffic. *Operations Research*, 68(2):411–434, 2020.

- [22] E. Comission. Commission proposes new measures on skills and talent to help address critical labour shortages. 2023.
- [23] V. Conitzer, R. Freeman, N. Shah, and J. W. Vaughan. Group fairness for the allocation of indivisible goods. In *AAAI Conference on Artificial Intelligence*, 2019.
- [24] G. Călinescu, C. Chekuri, M. Pál, and J. Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM J. Comput.*, pages 1740–1766, 2011.
- [25] C. D’Aspremont and L. Gevers. Equity and the informational basis of collective choice. volume 44, pages 199–209, 1977.
- [26] D. Delacrétaz, S. D. Kominers, and A. Teytelboym. Matching mechanisms for refugee resettlement. *American Economic Review*, 2023.
- [27] R. Deschamps and L. Gevers. Leximin and utilitarian rules: A joint characterization. *Journal of Economic Theory*, 17:143–163, 1978.
- [28] S. Devic, D. Kempe, V. Sharan, and A. Korolova. Fairness in matching under uncertainty. In *ICML*, 2023.
- [29] J. P. Dickerson, A. D. Procaccia, and T. Sandholm. Price of fairness in kidney exchange. In *Adaptive Agents and Multi-Agent Systems*, 2014.
- [30] J. Edmonds. Submodular functions, matroids, and certain polyhedra. *Combinatorial Structures and Their Applications*, pages 69–87, 1970.
- [31] M. El Halabi, S. Mitrović, A. Norouzi-Fard, J. Tardos, and J. M. Tarnawski. Fairness in streaming submodular maximization: Algorithms and hardness. *Advances in Neural Information Processing Systems*, 33:13609–13622, 2020.
- [32] M. El Halabi, J. Tarnawski, A. Norouzi-Fard, and T.-D. Vuong. Fairness in submodular maximization over a matroid constraint. In *International Conference on Artificial Intelligence and Statistics*, pages 1027–1035. PMLR, 2024.
- [33] S.-A. Esmaeili, S. Duppala, V. Nanda, A. Srinivasan, and J. P. Dickerson. Rawlsian fairness in online bipartite matching: Two-sided, group, and individual. In *Adaptive Agents and Multi-Agent Systems*, 2022.
- [34] U. N. H. C. for Refugees. Resettlement fact sheet 2023, 2023.
- [35] R. Freeman, N. Shah, and R. Vaish. Best of both worlds: Ex-ante and ex-post fairness in resource allocation. *Proceedings of the 21st ACM Conference on Economics and Computation*, 2020.
- [36] D. N. Freund, T. Lykouris, E. Paulson, B. Sturt, and W.-Y. Weng. Group fairness in dynamic refugee assignment. In *ACM Conference on Economics and Computation*, 2023.
- [37] A. Frieze and M. Karoński. *Introduction to random graphs*. Cambridge University Press, 2016.
- [38] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 120:386 – 391, 2013.
- [39] T. L. Gouic, J.-M. Loubes, and P. Rigollet. Projection to fairness in statistical learning. *arXiv*, 2020.
- [40] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer-Verlag, 2nd edition, 1993.
- [41] C. Haas. The price of fairness - a framework to explore trade-offs in algorithmic fairness. In *International Conference on Interaction Sciences*, 2019.
- [42] M. Hardt, E. Price, E. Price, and N. Srebro. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*, volume 29, 2016.

- [43] J. Herzog and T. Hibi. Discrete polymatroids. *Journal of Algebraic Combinatorics*, 16:239–268, 2002.
- [44] J. Herzog and T. Hibi. *Monomial Ideals*. Springer, 2011.
- [45] H. Hosseini, Z. Huang, A. Igarashi, and N. Shah. Class fairness in online matching. In *AAAI*, 2023.
- [46] A. Hylland and R. J. Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political Economy*, 87:293 – 314, 1979.
- [47] E. Kalai and M. Smorodinsky. Other solutions to nash’s bargaining problem. *Econometrica: Journal of the Econometric Society*, pages 513–518, 1975.
- [48] Y. Kamada and F. Kojima. Fair matching under constraints: Theory and applications. *Review of Economic Studies*, 2023.
- [49] M. Kearns, S. Neel, A. Roth, and Z. S. Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pages 2564–2572. PMLR, 2018.
- [50] A. Lambrecht and C. E. Tucker. Algorithmic bias? an empirical study of apparent gender-based discrimination in the display of stem career ads. *Manag. Sci.*, 65:2966–2981, 2019.
- [51] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *ACM Conference on Electronic Commerce*, 2004.
- [52] O. Loukas and H.-R. Chung. Demographic parity: Mitigating biases in real-world data. *arXiv preprint arXiv:2309.17347*, 2023.
- [53] W. Ma, P. Xu, and Y. Xu. Fairness maximization among offline agents in online-matching markets. 2021.
- [54] A. Mehta. Online matching and ad allocation. *Found. Trends Theor. Comput. Sci.*, 8:265–368, 2013.
- [55] A. K. Menon and R. C. Williamson. The cost of fairness in binary classification. In *Conference on Fairness, Accountability, and Transparency*, 2021.
- [56] M. Molina and P. Loiseau. Bounding and approximating intersectional fairness through marginal fairness. In *Advances in Neural Information Processing Systems*, volume 35, pages 16796–16807, 2022.
- [57] G. Nicosia, A. Pacifici, and U. Pferschy. Price of fairness for allocating a bounded resource. *European Journal of Operational Research*, pages 933–943, 2017.
- [58] M. J. Osborne and A. Rubinstein. *A course in game theory*. MIT Press, 1995.
- [59] T. Roughgarden. *Beyond the Worst-Case Analysis of Algorithms*. Cambridge University Press, 2020.
- [60] G. S. Sankar, A. Louis, M. Nasre, and P. Nimbhorkar. Matchings with group fairness constraints: Online and offline algorithms. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 377–383, 2021.
- [61] A. Schrijver et al. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer, 2003.
- [62] A. Sen. *Collective Choice and Social Welfare*. Harvard University Press, 2017.
- [63] T. Speicher, M. Ali, G. Venkatadri, F. N. Ribeiro, G. Arvanitakis, F. Benevenuto, K. P. Gummadi, P. Loiseau, and A. Mislove. Potential for discrimination in online targeted advertising. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81, pages 5–19. PMLR, 2018.

- [64] H. Steinhaus. Sur la division pragmatique. *Econometrica: Journal of the Econometric Society*, pages 315–319, 1949.
- [65] S. Tang and J. Yuan. Beyond submodularity: a unified framework of randomized set selection with group fairness constraints. *Journal of Combinatorial Optimization*, 45:1–22, 2023.
- [66] H. R. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974.
- [67] D. Weller. Fair division of a measurable space. *Journal of Mathematical Economics*, 14:5–17, 1985.
- [68] J. Yuan and S. Tang. Group fairness in non-monotone submodular maximization. *Journal of Combinatorial Optimization*, 45:1–15, 2023.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: Each contribution presented Section 1 is clearly linked to one or multiple Theorems and Propositions.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: See in particular Section 5 as well as Section 2.4 and Section 2.3.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an impor-

tant role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The proof of all Theorems and Propositions are provided in the Appendix and are appropriately cross referenced.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.

- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.

- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: The main theme of the paper being algorithmic fairness, the societal impacts are a core part of the discussions (e.g. see Section 1).

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.

- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.

- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A Further Related Works

Fairness notions: The study of fair algorithmic decision-making cover a broad range of fields, with fair division [64] in economics that focuses on concepts such as envy-freeness [16, 51, 66, 67] and maximin fairness [62], and machine learning that emphasizes statistical fairness notions like group fairness [23, 36, 60], including demographic parity [52]. Our new notion of opportunity fairness is inspired from both Equality of Opportunity [42] and Kalai-Smorodinsky fairness [47, 57]. Equality of Opportunity aims for *true-positive* rates to be independent of sensitive attributes; for opportunity fairness, this translates to ensuring that the resources allocated to a group are proportional to its opportunity level, the maximum allocation it could receive if it were considered in isolation. On the other hand, Kalai-Smorodinsky fairness requires maximizing the ratio $|S \cap E_c| / \max_{S' \in \mathcal{I}} |S' \cap E_c|$. While any maximum-size opportunity fair allocation satisfies Kalai-Smorodinsky fairness, the reverse does not necessarily hold, making our fairness notion more restrictive.

The Price of Fairness: The Price of Fairness was concurrently introduced by [11], who focus on maximin fairness and proportional fairness, and [15] who prioritize equitability and envy-freeness. They provided bounds for each fairness notion depending on the number of agents. Subsequently, [57] studied the price of fairness under the Kalai-Smorodinsky fairness notion for the *subset sum problem* and [29] studied the price of fairness in kidney-exchange. The concept of price of fairness has been extended to others research domains such as supervised machine learning [41, 39, 55] where the cost of fairness is studied on different prediction tasks. In this article, we initiate the study of the Price of Fairness under opportunity fairness. Unlike equitability, which requires identical allocations across groups, opportunity fairness is a more robust notion in the context of price of fairness. When some groups are inherently unable to receive resources, enforcing equitability leads to significant welfare loss, whereas we show that the price of fairness under opportunity fairness remains bounded. More importantly, we go beyond the traditional adversarial worst-case analysis, and instead consider more structured inputs, in the vein of [59], allowing for an average case analysis that better reflect trade-offs in real-world instances.

Fair Matroid Allocation Problems: The main objective of the matroid and fairness literature, initiated by [18], is to efficiently approximate maximum size fair allocations. Subsequent works extend this framework to submodular function optimization under fairness and matroid constraints [31, 32, 65, 68], while [9] study the computational complexity of finding optimal proportionally fair matching for more than two groups. We remark that maximum size opportunity fair fractional allocations can be computed efficiently whenever the underlying matroid possesses a polynomial-time independence oracle ([61], Chapter 44), as is the case of bipartite matching and communication network problems.

Fair matching: Recent work in matching have increasingly examined the impact of different fairness notions on matching mechanisms and how to design fair algorithms. [17, 28] and [48] examine the relation between fairness and stability in matching. Additionally, fairness in online matching has been studied in various contexts, including waiting time, equality of opportunity, and fairness constraints on the offline side of the market [19, 33, 45, 53, 60]. Our work contributes to this growing literature by introducing a new fairness notion, opportunity fairness, that captures the structural constraints of allocation problems.

B Matroid classes

The three examples introduced in Section 1 can be modeled as matroid allocation problems. Indeed,

1. The bipartite matching problem corresponds to a **transversal matroid**: Let $G = (U, V, A)$ be a bipartite graph. For a matching $\mu \subseteq A$ we denote $\mu(U) := \{u \in U \mid \exists v \in V, (u, v) \in \mu\}$. The pair (U, \mathcal{I}) , with $\mathcal{I} := \{\mu(U), \forall \mu \subseteq A \text{ matching}\}$, is called a transversal matroid.
2. The research funding allocation problem corresponds to a **partition matroid**. Let E be a finite set partitioned into disjoint subsets E^1, E^2, \dots, E^k , each representing, for instance, a department or category of funding. Given nonnegative integers b_1, b_2, \dots, b_k , the partition matroid is the pair (E, \mathcal{I}) , where $\mathcal{I} := \{S \subseteq E \mid |S \cap E^i| \leq b_i \text{ for all } i = 1, \dots, k\}$. This formulation captures allocation constraints where each group (e.g., department) has its own capacity limit.
3. The faculty assignment problem corresponds to a **laminar matroid**. Let E be a finite set and \mathcal{L} a laminar family of subsets of E (that is, for any $L_1, L_2 \in \mathcal{L}$, either $L_1 \subseteq L_2$, $L_2 \subseteq L_1$, or $L_1 \cap L_2 = \emptyset$). Given a capacity function $b : \mathcal{L} \rightarrow \mathbb{N}$, the laminar matroid is the pair (E, \mathcal{I}) , where

$\mathcal{I} := \{S \subseteq E \mid |S \cap L| \leq b(L) \text{ for all } L \in \mathcal{L}\}$. This structure models hierarchical constraints, such as departmental and team-level capacity limits in faculty assignments, or in private companies.

We also give the definition of further matroid classes, which are mathematically relevant to the studied problem.

4. **Graphic matroid:** Let $G = (U, A)$ be a graph. The pair (A, \mathcal{I}) , with $\mathcal{I} := \{S \subseteq A \mid S \text{ is acyclic}\}$, is called a graphic matroid.
5. **Uniform matroid:** Let E be a finite set and $b \in \mathbb{N}$. The b -uniform matroid is the pair (E, \mathcal{I}) , such that, $\mathcal{I} := \{S \subseteq E \mid |S| \leq b\}$.

C Null integral opportunity fair allocations

Figure 5 illustrates an example on a graphic matroid⁴ for two groups. The figure on the right shows the integer feasible allocations (the blue dots) and the set of opportunity fair allocations (the orange line), whose only intersection is at the origin.




Figure 5: Graphic matroid example showing that integrality can lead to null opportunity fair allocations. (Left) Graph defining a colored graphic matroid with two groups. Group 1 is denoted by the green edges while Group 2 by the red edges. It follows that $r(1) = 5$, $r(2) = 3$, and $r(\{1, 2\}) = 7$. (Right) Set of integer feasible allocations (blue dots) and set of opportunity fair allocations (orange line), whose only intersection is at the origin.

D Efficient computation of Opportunity fair allocations of maximum size

Whenever a matroid (E, \mathcal{I}) possesses a poly-time independent oracle, that is, an oracle that for any subset S of E tells if S is an independent set or not in polynomial time, we can construct a poly-time separation oracle for the corresponding polymatroid in $\mathbb{R}^{|E|}$ [61][Theorem 40.4]. A separation oracle, given a vector $x \in \mathbb{R}^{|E|}$ and the polymatroid

$$Q := \left\{ x \in \mathbb{R}_+^{|E|} \mid \sum_{e \in S} x_e \leq \text{rank}(S), \forall S \subseteq E \right\},$$

either tells if x belongs to Q or outputs a separating hyperplane. Starting from the poly-time separation oracle of a matroid, we can obtain a poly-time separation oracle for set of feasible fractional allocation of the corresponding colored matroid (or polymatroid). Therefore, we can compute an opportunity fair allocation of maximum size by solving the following problem,

$$\begin{aligned} & \max \sum_{c \in [C]} x_c \\ & \text{s.t. } x_c = \sum_{e \in E: e \in c} y_e & \forall c \in [C] \\ & \frac{x_c}{\text{rank}(c)} = \frac{x_{c'}}{\text{rank}(c')}, & \forall c, c' \in [C] \\ & (y_e)_{e \in E} \text{ is a feasible fractional allocation} \\ & x \in [0, 1]^C, y \in [0, 1]^{|E|}. \end{aligned}$$

Indeed, the previous problem can be solved in polynomial time using the ellipsoid method provided with the poly-time separation oracle of the polymatroid and the (trivial) poly-time separation oracle of the polytope of fairness constraints.

⁴Allocations within a graphic matroid correspond to acyclic subgraphs of the given graph.

One possible downside of this approach, is that the ellipsoid method can be quite slow in practice. Nonetheless, for many common matroid classes, including transversal and graphic matroids, the fair optimization problem admits a polynomial-size linear programming formulation, which allows us to use more efficient algorithms such as Simplex or interior point methods.

Indeed, this relates to the concept of extension complexity, which basically refers to the minimal number of constraints needed to describe the independence polytope, possibly adding some variables. Notably, [7] show that regular matroids have polynomial size extension complexity.

E Missing proofs

E.1 Proof of Proposition 2.4

Proposition 2.4. For any fractional opportunity fair maximum size allocation $x \in F$, there always exists a random allocation X , such that, $\mathbb{E}[X] = x$ and for all realizations, X is feasible, integral, $\left(1 - \frac{2C}{\min_c r(c)}\right)$ -opportunity fair, and $\|X - x\|_1 \leq C$.

Proof. Let M be a set of fractional allocations and x be an optimal fair allocation. Consider the translated unit cube $K = \lfloor x \rfloor + [0, 1]^C$ where $\lfloor x \rfloor = (\lfloor x_1 \rfloor, \dots, \lfloor x_C \rfloor)$, which contains x and all the integer vertices closest to x . We aim to prove that $K \cap M$ yields integer extreme points.

Let $M' := \mathbb{R}_+^C \cap (M - \lfloor x \rfloor)$ to be the positive quadrant of M translated by $\lfloor x \rfloor$. It holds

$$\begin{aligned} M' &= \left\{ y \in \mathbb{R}_+^C \mid \sum_{c \in \Lambda} y_c + \lfloor x_c \rfloor \leq r(\Lambda), \forall \Lambda \subset [C] \right\} \\ &= \left\{ y \in \mathbb{R}_+^C \mid \sum_{c \in \Lambda} y_c \leq r(\Lambda) - \sum_{c \in \Lambda} \lfloor x_c \rfloor, \forall \Lambda \subset [C] \right\}. \end{aligned}$$

Consider the function $f(\Lambda) := r(\Lambda) - \sum_{c \in \Lambda} \lfloor x_c \rfloor$. Since $x \in M$, so is $\lfloor x \rfloor$, hence $f(\Lambda) \geq 0$, and f takes only integer values as both r and $\lfloor x \rfloor$ take integer values. Moreover, f is submodular as both r and the application $\Lambda \mapsto \sum_{c \in \Lambda} x_c$ are submodular. Therefore, M' is an integral polymatroid.

The unit cube being also an integral polymatroid, by Theorem 35 of [30], $[0, 1]^C \cap M'$ has integral vertices, and thus so it does,

$$\lfloor x \rfloor + ([0, 1]^C \cap M') = K \cap M.$$

It follows that, since $x \in K \cap M$, which is a convex set (as both K and M are convex sets), x can be described as the convex combination of the extreme points of $K \cap M$. Since all extreme points of $K \cap M$ are integral, they are, in particular, a subset of the extreme points of K . By Carathéodory's Theorem, there exists $X^{(1)}, \dots, X^{(k)}$ vertices of $K \cap M$ for $k \leq C+1$, and $(p_1, \dots, p_k) \in [0, 1]^{C+1}$ with $\sum_{i=1}^k p_i = 1$, such that $x = \sum_{i=1}^k p_i X^{(i)}$.

The randomized rounding X then is defined by drawing $X^{(i)}$ with probability p_i . Note that for any realization of $X^{(i)}$ of X , $X^{(i)}$ is feasible and integral by construction. Moreover,

$$\mathbb{E}[X] = \sum_{i=1}^k p_i X^{(i)} = x,$$

and since $X^{(i)}$ for any $i \in \{1, \dots, k\}$, and x belong to $[0, 1]^C$, it always follows that $\|X^{(i)} - x\|_1 \leq C$. Regarding the γ -opportunity fairness, recall that x is 1-opportunity fair, so there exists $\alpha \in [0, 1]$ such that $x = \alpha(r(1), \dots, r(C))$. Moreover, because x is optimal, $\alpha \geq 1/C$. Indeed, either $\text{PoF} = 1 = r([C]) / (\alpha \sum_c r(c))$ which yields $\alpha = r([C]) / \sum_c r(c) \geq r([C]) / (Cr([C])) = 1/C$, or $\text{PoF} > 1$ in which case $(r(1), \dots, r(C)) / (C - 1)$ is feasible (see proof of Theorem 4.1) which implies, by definition, that $\alpha \geq (C - 1)$. For any $i, j \in [C]$, it follows,

$$\frac{X_i/r(i)}{X_j/r(j)} \geq \frac{(x_i - 1)/r(i)}{(x_j + 1)/r(j)} \geq \frac{\alpha - 1/\min_c r(c)}{\alpha + 1/\min_c r(c)} = 1 - \frac{2}{\alpha \min_c r(c) + 1} \geq 1 - \frac{2C}{\min_c r(c)},$$

concluding the first part of the proof.

Worst-case L_1 deviation: Consider a submodular function f such that for $S \subset [C]$, $f(S) = \max(|S|, C \cdot \mathbf{1}[1 \in S])$. This submodular function induces an integral polymatroid M (hence a colored matroid by surjection). The Pareto front of this polymatroid is $P = \{x \in \mathbb{R}_+^C : x_i \in [0, 1] \text{ for } 2 \leq i \leq C \text{ and } x_1 = C - \sum_{i \geq 2} x_i\}$. The optimal fair allocation is $x = (C^2/(2C-1), C/(2C-1), \dots, C/(2C-1))$. We now show that x is far from all vertices. Let y be an integral point of M , and let p be the projection over the last $C-1$ coordinates. Because p is a projection, $\|x - y\|_1 \geq \|p(x) - p(y)\|_1$. Now $p(x) = C/(2C-1) \cdot \mathbf{1}_{C-1}$ lies in the unit cube $[0, 1]^{C-1}$, and the closest vertex is $y' = \mathbf{1}_{C-1}$. Overall,

$$\|x - y\|_1 \geq \|p(x) - p(y)\|_1 \geq \|C/(2C-1) \cdot \mathbf{1}_{C-1} - \mathbf{1}_{C-1}\|_1 = (C-1) \frac{C-1}{2C-1} \geq \frac{C}{2} - \frac{1}{4} = \Omega(C),$$

as $C \geq 2$. This is true for any vertex y , hence no rounding scheme X can do better.

Worst-case fairness degradation: Using the construction from fig. 6, suppose group 1 matches up to r_1 agents, and the remaining groups share r_2 agents, with $(r_1/(C-1), r_2/(C-2), \dots, r_2/(C-2))$ the optimal fair allocation. Let r_1 and r_2 be such that the nearest integers to $r_1/(C-1)$ and $r_2/(C-2)$ are respectively $r_1/(C-1) - 1/2 + \epsilon$ and $r_2/(C-2) + 1/2 - \epsilon$. Hence, for $(C-1) \leq r_2 \ll r_1$, we have

$$\begin{aligned} \frac{(r_1/(C-1) - 1/2)/r_1}{(r_2/(C-1) + 1/2)/r_2} &= \frac{1/(C-1) - 1/(2r_1)}{1/(C-1) + 1/(2r_2)} = 1 - \frac{1/(2r_1) + 1/(2r_2)}{1/(C-1) + 1/(2r_2)} \\ &\approx 1 - \frac{C-1}{C-1+2r_2} \leq 1 - \frac{C-1}{3r_2} = 1 - O\left(\frac{C}{\min_c r(c)}\right). \end{aligned}$$

So no integral allocation can have a better approximately fair guarantee than $1 - O(C/\min_c r(c))$. \square

E.2 Proof of Proposition 2.5

Proposition 2.5. Let M be a C -colored matroid. Then, either $\gamma < \max_{x \in F} \min_{c \in [C]} x_c/r(c)$ and $\text{PoF}_\gamma(M) = 1$, or

$$\gamma \cdot \text{PoF}(M) \leq \text{PoF}_\gamma(M) \leq \frac{\gamma \cdot \text{PoF}(M)}{1 - (1 - \gamma) \frac{((C-1) - \text{PoF}(M))}{C-2}}.$$

Proof. The lower bound arises by disregarding the feasibility constraints, and improving the original optimal fair solution by simply scaling non-tight coordinates by γ . The upper bound is obtained via a convex combination between the optimal fair and optimal unfair solution which is made to be exactly γ -fair, and which remains feasible by convexity.

Lower bound. Let $x = \arg\max_{y \in F} \|y\|_1 = \alpha(r(1), \dots, r(C))$, where $\alpha \in [0, 1]$. If x is maximal, $\text{PoF}_\gamma(M) = \text{PoF}(M) = 1$. Otherwise, there exists $\Lambda \subsetneq [C]$ such that $x \in \arg\max_{y \in M} \sum_{c \in \Lambda} y_c$. Define x' as $x'_i = x_i$ for all $i \in \Lambda$ and $x'_i = \frac{1}{\gamma} x_i$ for all $i \notin \Lambda$.

Let us denote by F_γ the γ -opportunity fair feasible set. First, we prove that $\|x'\|_1 \geq \max_{F_\gamma} \|x\|_1$. Suppose there exists $x'' \in F_\gamma$ a γ -opportunity realizable point such that $\|x''\|_1 > \|x'\|_1$. Let $\Gamma = \{i \in [C], x''_i > x'_i\}$. Since

$$\sum_{\Lambda} x''_i \leq r(\Lambda) = \sum_{\Lambda} x'_i,$$

it follows that, $\sum_{\Lambda^c} x''_i > \sum_{\Lambda^c} x'_i$, and thus $\Gamma \cap \Lambda^c \neq \emptyset$. Let $i \in \Gamma \cap \Lambda^c$, then $x''_i > x'_i = \frac{1}{\gamma} x_i$. Moreover, there must exist $j \in \Lambda$ such that $x''_j \leq x_j$ since $x \in \arg\max_M \sum_{c \in \Lambda} y_c$. Therefore,

$$\frac{x''_i}{r(i)} > \frac{1}{\gamma} \cdot \frac{x_i}{r(i)} = \frac{1}{\gamma} \cdot \frac{x_j}{r(j)} \geq \frac{1}{\gamma} \cdot \frac{x''_j}{r(j)},$$

which contradicts x'' being γ -opportunity fair. It follows,

$$\|x'\|_1 = \sum_{i \in \Lambda} x_i + \frac{1}{\gamma} \sum_{i \in \Lambda^c} x_i = \alpha \left(\sum_{i \in \Lambda} r(i) + \frac{1}{\gamma} \sum_{i \in \Lambda^c} r(i) \right),$$

and therefore,

$$\frac{\text{PoF}(M)}{\text{PoF}_\gamma(M)} = \frac{\max_{y \in F_\gamma} \|y\|_1}{\max_{y \in F} \|y\|_1} \leq \frac{\|x'\|_1}{\|x\|_1} = \frac{\sum_{i \in \Lambda} r(i) + \frac{1}{\gamma} \sum_{i \in \Lambda^c} r(i)}{\sum_{i \in [C]} r(i)} \leq \frac{1}{\gamma},$$

thus, $\text{PoF}_\gamma(M) \geq \gamma \text{PoF}(M)$.

Upper bound. As before, let $x = \arg\max_{y \in F} \|y\|_1 = \alpha(r(1), \dots, r(C))$. We suppose that x is not optimal, otherwise $\text{PoF}_\gamma = \text{PoF} = 1$. Let x' be a point in the Pareto frontier that Pareto dominates x . By the polymatroid characterization, it holds that x is minimal, i.e. $\|x\|_1 = r([C])$. For $\lambda \in \mathbb{R}_+$, consider the combination $z = \lambda x' + (1 - \lambda)x$. Because x' Pareto dominates x , we have $x'_i \geq x_i$ for all $i \in [C]$, hence $z_i \geq x_i$. By definition of $r(i)$ we also have that $x'_i \leq r(i)$, thus $z_i \leq \lambda r(i) + (1 - \lambda)x_i$. This implies that

$$\frac{z_i/r(i)}{z_j/r(j)} \geq \frac{x_i/r(i)}{\lambda + (1 - \lambda)(x_j/r(j))} = \frac{\alpha}{\lambda + (1 - \lambda)\alpha}.$$

In particular, z is γ -opportunity fair whenever $\alpha/(\lambda + (1 - \lambda)\alpha) \geq \gamma$, i.e., whenever $\lambda \leq \frac{\alpha(1 - \gamma)}{\gamma(1 - \alpha)}$. If $\alpha(1 - \gamma)/(\gamma(1 - \alpha)) > 1$, which is equivalent to $\gamma < \alpha$, then taking $\lambda = 1$ yields that x' is γ -fair, and it is already optimal so $\text{PoF}_\gamma = 1$. In the following we let $\gamma \geq \alpha$. Let us take $\lambda = \alpha(1 - \gamma)/(\gamma(1 - \alpha)) \leq 1$ such that z in M by convex combination. Because z is γ -opportunity fair and feasible, $z \in F_\gamma$. Hence

$$\begin{aligned} \text{PoF}_\gamma(M) &= \frac{r([C])}{\max_{y \in F_\gamma} \|y\|_1} \leq \frac{r([C])}{\|z\|_1} = \frac{r([C])}{\lambda \sum_{c \in [C]} x'_c + (1 - \lambda) \sum_{c \in [C]} x_c} \\ &= \frac{\text{PoF}(M)}{\lambda \text{PoF}(M) + (1 - \lambda) \cdot 1} \\ &= \frac{\text{PoF}(M)}{\frac{\alpha(1 - \gamma)}{\gamma(1 - \alpha)} \text{PoF}(M) + \frac{\gamma(1 - \alpha) - \alpha(1 - \gamma)}{\gamma(1 - \alpha)}} \\ &= \frac{\gamma(1 - \alpha) \text{PoF}(M)}{\alpha(1 - \gamma) \text{PoF}(M) + \gamma - \alpha} \\ &\leq \frac{\gamma(1 - \frac{1}{C-1}) \text{PoF}(M)}{\frac{1}{C-1} (1 - \gamma) \text{PoF}(M) + \gamma - \frac{1}{C-1}} \\ &= \frac{\gamma \text{PoF}(M)}{1 - (1 - \gamma) \frac{((C-1) - \text{PoF}(M))}{C-2}}, \end{aligned}$$

where we used that $\alpha \geq 1/(C - 1)$ as otherwise $\text{PoF}(M) = 1$. \square

E.3 Proof of Proposition 3.2

Proposition 3.2. Let $((E_c)_{c \in [C]}, \mathcal{I})$ be a colored matroid with rank function r and set of feasible fractional allocations M . Then, M is the polymatroid associated to the function r , i.e.,

$$M = \{x \in \mathbb{R}_+^C \mid \sum_{c \in \Lambda} x_c \leq r(\Lambda), \forall \Lambda \subseteq [C]\}.$$

Proof. Let $((E_c)_{c \in [C]}, \mathcal{I})$ be a C -colored matroid. It is sufficient to show that

$$\mathbf{I} := \{(|S \cap E_1|, \dots, |S \cap E_C|) \in \mathbb{N}^C \mid S \in \mathcal{I}\},$$

the set of integer feasible allocation, is a discrete polymatroid, and to conclude by taking the convex hull, as an integral polymatroid is equivalent to the convex hull of a discrete polymatroid [30]. We prove this by showing that \mathbf{I} satisfies equivalent conditions for a set to be a discrete polymatroid [44], which are:

1. For any $x \in \mathbb{N}^C$ and $y \in \mathbf{I}$ such that $x \leq y$ (component wise), $x \in \mathbf{I}$.
2. For any $x, y \in \mathbf{I}$, with $\|x\|_1 < \|y\|_1$, there exists $c \in [C]$ such that $x_c < y_c$ and $x + \vec{e}_c \in \mathbf{I}$, where \vec{e}_c is the canonical vector with value 1 at the c -th entry and 0 otherwise.

We will repeatedly leverage the properties of the underlying matroid to show, by exhaustion, that it is always possible to find an element in the allocation associated with y whose color is underrepresented in the allocation associated with x .

The first property is a direct consequence of the hereditary property of \mathcal{I} . Let $x, y \in \mathcal{I}$ such that $\|x\|_1 < \|y\|_1$. Let $A_x, A_y \in \mathcal{I}$ be two independent sets such that $(|A_x \cap E_1|, \dots, |A_x \cap E_C|) = x$ and $(|A_y \cap E_1|, \dots, |A_y \cap E_C|) = y$. In particular, $|A_x| < |A_y|$. By the augmentation property, there exists $e \in A_y \setminus A_x$ such that $A_x \cup \{e\} \in \mathcal{I}$. Let c be the group of e . If $x_c < y_c$, the proof is over. Suppose, otherwise, that $x_c \geq y_c$. Let $e' \in A_x$ such that $e' \in E_c$ as well, and define $A'_x = A_x \cup \{e\} \setminus \{e'\}$. By the hereditary property, $A'_x \in \mathcal{I}$ and, by construction, $(|A'_x \cap E_1|, \dots, |A'_x \cap E_C|) = x$. Applying again the augmentation property, there exists $e'' \in A_y \setminus A'_x$ such that $A'_x \cup \{e''\} \in \mathcal{I}$. Notice that $e'' \notin \{e, e'\}$. The proof is concluded by repeating the same argument until finding an element in some E_c such that $x_c < y_c$. The procedure stops after a finite amount of iteration as at every iteration the element in A_y obtained from the augmentation property must be different to all previous ones as well to all replaced elements in A_x . \square

E.4 Proof of Corollary 3.3

Corollary 3.3. The price of opportunity fairness of a polymatroid M is given by,

$$\text{PoF}(M) = \frac{r([C])}{\sum_{c \in [C]} r(c)} \cdot \max_{\Lambda \subseteq [C]} \frac{\sum_{c \in \Lambda} r(c)}{r(\Lambda)}.$$

Proof. Let x^* be a maximum size opportunity fair allocation. The opportunity fair requirement implies that x^* belongs to the line $t \cdot (r(1), \dots, r(C))$ for $t > 0$. Let $t^* > 0$ such that $x^* = t^* \cdot (r(1), \dots, r(C))$. Since x^* is a feasible fractional allocation, Proposition 3.2 implies that for any $\Lambda \subseteq [C]$,

$$t^* \sum_{c \in \Lambda} r(c) \leq r(\Lambda).$$

It follows that

$$t^* = \min_{\Lambda \subseteq [C]} \frac{r(\Lambda)}{\sum_{c \in \Lambda} r(c)},$$

and, in particular, that,

$$\text{PoF}(M) = \frac{r([C])}{t^* \sum_{c \in [C]} r(c)} = \frac{r([C])}{\sum_{c \in [C]} r(c)} \cdot \max_{\Lambda \subseteq [C]} \frac{\sum_{c \in \Lambda} r(c)}{r(\Lambda)},$$

concluding the proof. \square

E.5 Proof of Theorem 4.1

Theorem 4.1. For any C -colored matroid M , we have $\text{PoF}(M) \leq C - 1$, and this bound is tight.

Proof. Using the monotonicity and the sub-additivity of the colored-matroid rank function r , we will upper bound the alternate formulation of the PoF in Equation (1).

Let M be a C -dimensional polymatroid. Let Λ^* be the maximizer of Equation (1). Whenever $\Lambda^* = [C]$ it holds $\text{PoF}(M) = 1 \leq C - 1$. Suppose then that $\Lambda^* \subsetneq [C]$. It follows,

$$\text{PoF}(M) = \frac{r([C])}{r(\Lambda^*)} \cdot \frac{\sum_{c \in \Lambda^*} r(c)}{\sum_{c \in [C]} r(c)}.$$

Since $r(\Lambda^*) \geq \max_{c \in \Lambda^*} r(c) \geq \frac{1}{|\Lambda^*|} \sum_{c \in \Lambda^*} r(c)$, it follows that,

$$\text{PoF}(M) \leq |\Lambda^*| \cdot \frac{r([C])}{\sum_{c \in \Lambda^*} r(c)} \cdot \frac{\sum_{c \in \Lambda^*} r(c)}{\sum_{c \in [C]} r(c)} \leq |\Lambda^*| \leq C - 1,$$

where we have used that $r([C]) \leq \sum_{c \in [C]} r(c)$ because r is submodular and thus sub-additive, and $|\Lambda^*| \leq C - 1$ as $\Lambda^* \neq [C]$.

To show that this bound is tight, we exhibit a sequence of C -dimensional polymatroids for which the bound is tight in the limit. Consider a bipartite graph as in Figure 6. Let E_1 be independently

connected to r_1 nodes, and E_2, \dots, E_C be completely connected to the same r_2 nodes. The last $C - 1$ groups are in competition for resources while the first group suffers no competition. It holds $r(1) = r_1$, while for any $\Lambda \subseteq [C] \setminus \{1\}$, $r(\Lambda) = r_2$, and $r(\Lambda \cup \{1\}) = r_1 + r_2$. In particular, it follows that Λ^* , the maximizer of Equation (1), is given by $\Lambda^* = [C] \setminus \{1\}$, and,

$$\text{PoF}(M) = \frac{r_1 + r_2}{r_1 + (C - 1)r_2} \cdot \frac{(C - 1)r_2}{r_2} = \frac{r_1 + r_2}{r_1 + (C - 1)r_2} \cdot (C - 1) \xrightarrow{r_1 \rightarrow \infty} C - 1. \quad \square$$




Figure 6: Transversal polymatroid that makes the PoF bound of $C - 1$ tight (Theorem 4.1). Group 1 is totally independent of the rest of the groups, while groups $\{2, 3, \dots, C\}$ compete for the same resources.

The proof of Theorem 4.1 shows that for transversal matroid, i.e. bipartite matching, our main application, the bound is tight. Regarding **graphic matroids** (see Appendix B for a definition), the Walecki construction [6] which states that any clique of $2C - 1$ vertices has a decomposition into $C - 1$ disjoint Hamiltonian cycles, allows to design a tight example for the PoF upper bound. Indeed, associate each Hamiltonian cycle to a color $c \in [C - 1]$ and add one extra group of edges with color C as illustrated in Figure 7. It is not hard to see that this construction achieves a PoF equal to $C - 1$.




Figure 7: Graphic polymatroid tight example for worst-case PoF equal to $C - 1$ (Theorem 4.1), based on the Walecki construction, for 5 groups. Groups are represented by colored edges. The first 4 figures show the edges of each group, while the clique corresponds to their union. The final figure considers a 5-th group that is totally independent on the rest, leading to a similar construction as in Figure 6.

The same transversal matroid example can be used for **partition matroid**. This also implies the tightness of this bound for larger sub-classes of matroid which include either graphic or partition matroids, such as **linear** or **laminar matroids**. For **uniform matroids**, in exchange, it is immediate to see that $\text{PoF}(M) = 1$: if the opportunity fairness constraint is violated, it is always possible to take the excessive (potentially fractional) resources from over-represented colors and give them to the under-represented ones. It remains open to prove if intermediate cases (not 1 nor $C - 1$) exist for some family of matroids.

E.6 Alternative proof of Corollary 4.2

Corollary 4.2 states that no loss of social welfare is incurred by imposing opportunity fairness when agents are divided into two groups. The same conclusion can be easily proven from a geometric point of view, as the line directed by (r_1, r_2) necessarily intersects with the Pareto frontier, which corresponds to the set of social welfare maximizing allocations and therefore, it is directed by $(1, -1)$. For $C > 2$, the property does not hold anymore, as proven by the tight $C - 1$ bound, since the line directed by $(r(1), \dots, r(C))$ does not necessarily intersect with the Pareto frontier. Figure 8 illustrates these situations for $C = 2$ and $C = 3$.




Figure 8: Relation between the Pareto frontier P (light blue) and the set of opportunity fair allocations F (orange) for two and three groups. For $C = 2$ they always intersect, however it is not always the case for $C > 2$ as illustrated on the right (the cross marks the largest feasible fair matching).

E.7 Proof of Proposition 4.4

Proposition 4.4. For any C -colored matroid M , it holds,

$$\text{PoF}(M) \leq \frac{1}{2} \cdot \frac{\max_{c \in [C]} r(c)}{\min_{c \in [C]} r(c)} + \frac{C}{4} \cdot \left(\frac{\max_{c \in [C]} r(c)}{\min_{c \in [C]} r(c)} \right)^2 + \frac{1}{4C} \cdot \mathbb{1}\{C \text{ odd}\}.$$

Moreover, whenever all groups have the same rank, the resulting bound is tight.

Proof. From Corollary 3.3 we know that $\text{PoF}(M) = \frac{r([C])}{\sum_{c \in [C]} r(c)} \cdot \max_{\Lambda \subseteq [C]} \frac{\sum_{c \in \Lambda} r(c)}{r(\Lambda)}$. Let Λ^* be the argmax in the equation, and let us reorder the groups so $\Lambda^* = [\ell]$ for some $\ell \in [C]$. Denote $\gamma := \frac{\max_{c \in [C]} r(c)}{\min_{c \in [C]} r(c)}$. First, remark that the sub-additivity of the r function (consequence of non-negativity and submodularity) implies the following inequality,

$$r([C]) - r([\ell]) \leq r([C] \setminus [\ell]) \leq \sum_{c=\ell+1}^C r(c) \leq (C - \ell) \max\{r(c), c \in [C]\}.$$

It follows that

$$\text{PoF}(M) = \frac{r([C])}{r([\ell])} \cdot \frac{\sum_{c \in [\ell]} r(c)}{\sum_{c \in [C]} r(c)} \leq \left(1 + \frac{r([C]) - r([\ell])}{r([\ell])} \right) \frac{\ell}{C} \gamma \leq (1 + (C - \ell)\gamma) \frac{\ell}{C} \gamma.$$

The right-hand side of the previous inequality is maximized (subject to $\ell \in \mathbb{N}$) at $\ell = C/2 + \mathbb{1}\{C \text{ odd}\}/2\gamma$, which leads to the stated upper bound. Concerning the general result of the tightness of the bound, consider the constructions illustrated in Figures 9 and 10 where $\lfloor \frac{C}{2} \rfloor$ groups are isolated and $\lceil \frac{C}{2} \rceil$ compete for the same resources, respectively for transversal and graphic matroids, for (left) $C = 4$ and $r(c) = 3$ for all $c \in [C]$, and (right) $C = 5$ with $r(c) = 5$ for all $c \in [C]$. Suppose, moreover, that all groups have rank r , for some $r \in \mathbb{N}$. An opportunity fair allocation must allocate at most $r/\lceil \frac{C}{2} \rceil$ resources to each group. It follows that

$$\text{PoF} = \frac{Cr/\lceil \frac{C}{2} \rceil}{r + r\lfloor \frac{C}{2} \rfloor} = \frac{(1 + \lfloor C/2 \rfloor)\lceil C/2 \rceil}{C} = \frac{1}{2} + \frac{C}{4} + \frac{\mathbb{1}\{C \text{ odd}\}}{4C}. \quad \square$$

E.8 Proof of Proposition 4.6

Proposition 4.6. Let M be a C -colored matroid. Suppose that for any c, c' in $[C]$, $r(c) = r(c')$. Whenever $\rho \in [1/C, 1/(C - 1)]$, $\text{PoF}(M) = 1$. Otherwise,

$$\text{PoF}(M) \leq \rho \max \left(\frac{C - \lfloor C\rho \rfloor + 1}{C\rho - \lfloor C\rho \rfloor + 1}, C - \lfloor C\rho \rfloor \right) \leq \rho((1 - \rho)C + 1).$$

In addition, the first upper bound is jointly tight in ρ and C .

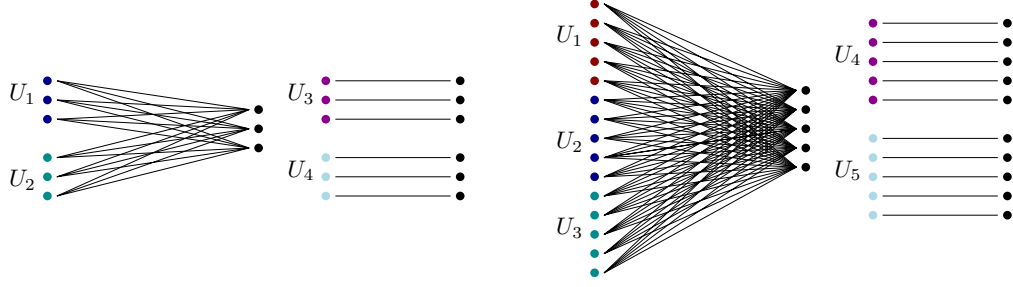


Figure 9: Tight bound example for PoF as stated in Proposition 4.4 for transversal matroids with (left) four groups and (right) five groups.




Figure 10: Tight bound example for PoF as stated in Proposition 4.4 for graphic matroids with (left) four groups and (right) five groups.

Proof. Let M be a polymatroid such that all groups $c \in [C]$ have the same rank r . Suppose $\rho \in [1/C, 1/(C-1)]$ and $\text{PoF}(M) > 1$. Bounding the PoF as in Theorem 4.1 plus using the fact that $\rho \leq 1/(C-1)$ leads to,

$$\text{PoF}(M) \leq \rho(C-1) \leq 1,$$

which is a contradiction.

Suppose $\rho \in [1/(C-1), 1]$. Let $\alpha^* = \max\{\alpha \in [0, 1] : \alpha(r, \dots, r) \in M\}$, the price of fairness is written as,

$$\text{PoF} = \frac{r([C])}{\alpha^* \sum_{c \in [C]} r} = \frac{\rho \sum_{c \in [C]} r}{\alpha^* \sum_{c \in [C]} r} = \frac{\rho}{\alpha^*}.$$

Therefore, the stated upper bound comes from proving that

$$\frac{1}{\alpha^*} \leq \max \left\{ \frac{C - \lfloor C\rho \rfloor + 1}{C\rho - \lfloor C\rho \rfloor + 1}, C - \lfloor C\rho \rfloor \right\}. \quad (5)$$

Let $\sigma \in \Sigma([C])$ be a permutation, $c \in [C]$, and denote, for $\ell \in [C]$,

$$\alpha_\ell(\sigma) := \frac{r(\sigma([\ell]))}{\sum_{t \in [\ell]} r(\sigma(t))},$$

where $r(\sigma([\ell])) = r(\{\sigma(1), \dots, \sigma(\ell)\})$ corresponds to the size of a maximum size allocation in the submatroid obtained by restricting to the groups in the first ℓ entries of σ . With this in mind, it follows,

$$\alpha^* = \min_{\sigma \in \Sigma([C])} \min_{\ell \in [C]} \alpha_\ell(\sigma).$$

Therefore, in order to prove Equation (5) it is enough to prove that for any permutation $\sigma \in \Sigma([C])$ and any $\ell \in [C]$, Equation (5) holds for $\alpha_\ell(\sigma)$. Let us prove the property for $\sigma = I_C$, the identity permutation. Notice this is done without loss of generality as the same argument will work for any other permutation σ . It follows,

$$\begin{aligned} \alpha_\ell &= \frac{r([\ell]) + \sum_{t \in [\ell]} r(t) - \sum_{t \in [\ell]} r(t)}{\sum_{t \in [\ell]} r(t)} \\ &= 1 - \frac{\sum_{t \in [\ell]} r(t) - r([\ell]) }{\sum_{t \in [\ell]} r(t)} \end{aligned}$$

$$\begin{aligned}
&= 1 - \frac{\sum_{t \in [L]} r(t) - \sum_{t=\ell+1}^C r(t) - r([\ell])}{\sum_{t \in [\ell]} r(t)} \\
&= 1 - \frac{Cr - \sum_{t=\ell+1}^C r(t) - r([\ell])}{\ell r} \\
&= 1 - \frac{Cr - r([C]) - \sum_{t=\ell+1}^C r(t) - r([\ell]) + r([C])}{\ell r} \\
&= 1 - \frac{Cr - \rho Cr}{\ell r} + \frac{\sum_{t=\ell+1}^C r(t) + r([\ell]) - r([C])}{\ell r} \\
&= 1 - \frac{C(1-\rho)}{\ell} + \frac{\sum_{t=\ell+1}^C r(t) + \sum_{t=\ell+1}^C r([t-1]) - r([t])}{\ell r} \\
&= 1 - \frac{C(1-\rho)}{\ell} + \frac{\sum_{t=\ell+1}^C [r(t) - r([t]) + r([t-1])]}{\ell r}.
\end{aligned}$$

The numerator of the third term satisfies,

$$\sum_{t=\ell+1}^C [r(t) - r([t]) + r([t-1])] \geq \max\{0, C(1-\rho) - (\ell-1)\}.$$

Indeed, the term is always non-negative as the rank function is submodular and non-negative, therefore, $r([t]) = r([t-1] \cup \{t\}) \leq r([t-1]) + r(t)$. The second lower bound comes from,

$$\begin{aligned}
\sum_{t=\ell+1}^C [r(t) - r([t]) + r([t-1])] &= \sum_{t \in [C]} r(t) - r([C]) + r([\ell]) - \sum_{t \in [\ell]} r(t) \\
&= Cr(1-\rho) + r([\ell]) - \ell r \\
&\geq Cr(1-\rho) - (\ell-1)r,
\end{aligned}$$

where we have used that $r([\ell]) \geq r(\ell) = r$. It follows,

$$\alpha_\ell \geq 1 - \frac{C(1-\rho)}{\ell} + \max\left\{0, \frac{C(1-\rho) - (\ell-1)}{\ell}\right\} = \max\left\{\frac{C(\rho-1) + \ell}{\ell}, \frac{1}{\ell}\right\}.$$

In particular, as the lower bound over α_ℓ does not depend on the chosen permutation,

$$\alpha^* \geq \min_{\ell \in [C]} \max\left\{\frac{C(\rho-1) + \ell}{\ell}, \frac{1}{\ell}\right\},$$

whose minimum is attained at $\ell^* = (1-\rho)C + 1$. Remark the second upper bound is obtained by replacing ℓ^* in the previous inequality. Regarding the first upper bound, as ℓ must be an integer, the minimum is either reached at $\lfloor \ell^* \rfloor$ or $\lceil \ell^* \rceil$. It follows,

$$\alpha^* \geq \min\left\{\frac{C(\rho-1) + \lceil (1-\rho)C + 1 \rceil}{\lceil (1-\rho)C + 1 \rceil}, \frac{1}{\lfloor (1-\rho)C + 1 \rfloor}\right\} = \min\left\{\frac{C\rho - \lfloor C\rho \rfloor + 1}{C - \lfloor C\rho \rfloor + 1}, \frac{1}{C - \lfloor C\rho \rfloor}\right\},$$

which concludes the proof of the stated upper bound.

Regarding the tightness of the bound, we provide the example for transversal matroids. A similar construction can be done for graphic matroids by using the Hamiltonian cycle decomposition. Let $\rho \in [1/C, 1]$, $\rho \in \mathbb{Q}$, $r \gg 1$, and denote $\ell^* := \lfloor C\rho \rfloor$ and $r_1 = (C\rho - \ell^*)r$. We take r such that $r_1 \in \mathbb{N}$. Consider the following bipartite graph,




Figure 11: Tight bound example for PoF as stated in Proposition 4.6 for transversal matroids.

where each group E_ℓ has r elements. All groups $\ell \in \{1, \dots, \ell^* - 1\}$ are independent and have $r(\ell) = r$. All groups $\ell \in \{\ell^* + 1, \dots, C\}$ share resources and have $r(\ell) = r$. Finally, group E_{ℓ^*} is a *semi-independent* group, where r_1 agents are connected to r_1 resources and $r - r_1$ agents belong to the clique. We obtain $r(\ell^*) = r$ as well. Finally, remark $r([C]) = (\ell^* - 1)r + r_1 + r = C\rho r$.

We focus next on the maximum size opportunity fair allocation. Notice that, as all groups have the same rank, an allocation x is opportunity fair if $x_\ell = x_k$ for all $\ell, k \in [C]$. Since all groups $\ell \in \{\ell^* + 1, \dots, C\}$ share all their resources, the highest share than can be fairly allocated to them is,

$$x_\ell = \frac{r}{C - \ell^*}.$$

This allocation is feasible if and only if the remaining available resources to be allocated to ℓ^* are enough to fulfill its demand, i.e, if and only if,

$$r_1 \geq \frac{r}{C - \ell^*} \quad (6)$$

Moreover, remark that depending on whether Equation (6) holds or not, the maximum on the stated upper bound gets a different value,

$$\begin{aligned} \frac{r}{C - \lfloor C\rho \rfloor} &\leq (L\rho - \lfloor C\rho \rfloor)r \\ \Leftrightarrow \frac{1}{C - \lfloor C\rho \rfloor} + 1 &\leq C\rho - \lfloor C\rho \rfloor + 1 \\ \Leftrightarrow \frac{C - \lfloor C\rho \rfloor + 1}{C - \lfloor C\rho \rfloor} &\leq C\rho - \lfloor C\rho \rfloor + 1 \\ \Leftrightarrow \frac{C - \lfloor C\rho \rfloor + 1}{C\rho - \lfloor C\rho \rfloor + 1} &\leq C - \lfloor C\rho \rfloor. \end{aligned}$$

Suppose Equation (6) holds. It follows the allocation is feasible and,

$$\|x\|_1 = \frac{Cr}{C - \ell^*} = \frac{Cr}{C - \lfloor C\rho \rfloor},$$

and the Price of Fairness is equal to,

$$\text{PoF} = \frac{C\rho r}{\frac{Cr}{C - \lfloor C\rho \rfloor}} = \rho(C - \lfloor C\rho \rfloor),$$

which is indeed equal to the upper bound. Suppose Equation (6) does not hold. In particular, the opportunity fair allocation must allocate some share of the r resources on the clique to E_{ℓ^*} . Let $s \in (0, 1)$ denote the share. We obtain the following system,

$$sr + r_1 = \frac{(1 - s)r}{C - \ell^*},$$

whose solution is given by

$$s^* = \frac{1 - (C\rho - \ell^*)(C - \ell^*)}{C - \ell^* + 1}.$$

It follows the opportunity fair allocation x has size,

$$\|x\|_1 = C \cdot \frac{(1 - s^*)r}{C - \ell^*} = \frac{C(C\rho - \ell^* + 1)r}{C - \ell^* + 1},$$

which yields,

$$\text{PoF} = \frac{C\rho r}{\frac{C(C\rho - \ell^* + 1)r}{C - \ell^* + 1}} = \rho \cdot \frac{C - \ell^* + 1}{C\rho - \ell^* + 1} = \rho \cdot \frac{C - \lfloor C\rho \rfloor + 1}{C\rho - \lfloor C\rho \rfloor + 1},$$

which corresponds to the stated upper bound when Equation (6) does not hold. \square

E.9 Proof of Proposition 4.7

We recall that $(M_n(p))_{n \in \mathbb{N}}$ is a sequence of C -colored matroids over sets $(E^{(n)})$ such that $|E^{(n)}| = n$, randomly colored according to $p = (p_1, p_2, \dots, p_C)$, and $(r_n)_{n \in \mathbb{N}}$ the associated sequence of rank functions. For each $c \in [C]$, suppose the following limit exists,

$$R(p_c) := \lim_{n \rightarrow \infty} \frac{\mathbb{E}_{p_c}[r_n(c)]}{n}, \quad (7)$$

and recall its natural extension to any subset $\Lambda \subseteq [C]$,

$$R\left(\sum_{c \in \Lambda} p_c\right) := \lim_{n \rightarrow \infty} \frac{1}{n} \cdot \mathbb{E}_{(p_c)_{c \in \Lambda}}[r_n(\bigcup_{c \in \Lambda} E_c)].$$

Proposition 4.7. If $R(1) > 0$, it follows

$$\text{PoF}(M_n(p)) \xrightarrow[n \rightarrow \infty]{P} \max_{\Lambda \subseteq [C]} \frac{R(1)}{\sum_{c \in [C]} R(p_c)} \cdot \frac{\sum_{c \in \Lambda} R(p_c)}{R(\sum_{c \in \Lambda} p_c)}, \quad (8)$$

where \xrightarrow{P} denotes convergence in probability. Moreover, the function R is such that $R(0) = 0$, R is concave, non-decreasing, and 1-Lipschitz. Finally, for any such function R , there exists a double sequence of C -colored matroids $(M'_{n,m})$ such that, for R'_m defined similarly to R for the sequence $(M'_{n,m})_{n \in \mathbb{N}}$,

$$\|R - R'_m\|_\infty \xrightarrow{m \rightarrow \infty} 0.$$

Proof. Let $\mathcal{R} := \{f : [0, 1] \rightarrow \mathbb{R} \mid f \text{ is a concave, non-decreasing, 1-Lipschitz function, and } f(0) = 0\}$. Remark that \mathcal{R} is closed and convex. We divide the proof in several steps. First, we prove the function R defined in Equation (7) belongs to \mathcal{R} . Second, we prove the PoF converges in probability to the stated limit. Third, we construct a double sequence of colored matroids $(M'_{n,m})_{n,m \in \mathbb{N}}$ such that R is well approximated by R_m , where each R_m is defined as in Equation (7) for $(M'_{n,m})_{n \in \mathbb{N}}$.

1. Function R belongs to \mathcal{R} . Since \mathcal{R} is closed, it is enough to prove that for each $n \in \mathbb{N}$, the mapping

$$Q_n : [0, 1] \rightarrow \mathbb{R}_+ \\ p_c \mapsto \frac{\mathbb{E}_{p_c}[r_n(c)]}{n}$$

belongs to \mathcal{R} . Clearly, $Q_n(0) = 0$. Regarding concavity and monotonicity, remark

$$\mathbb{E}_{p_c}[r_n(c)] = \sum_{S \subseteq E} r_n(S) \mathbb{P}(E_c = S) = \sum_{S \subseteq E} r_n(S) p_c^{|S|} (1 - p_c)^{|E| - |S|},$$

which can be seen as the multi-linear extension of the rank function r_n of M_n evaluated at (p_c, \dots, p_c) . It follows that $\mathbb{E}_{p_c}[r_n(c)]$ is a concave and non-decreasing function as r_n is submodular [24]. Moreover, Q_n is also concave and non-decreasing. Finally, remark,

$$\mathbb{E}_{p_c}[r_n(c)] \leq \mathbb{E}_{p_c}[|E_c|] = np_c,$$

where n is the total number of agents in E . In particular, Q_n is 1-Lipschitz⁵.

2. Convergence of PoF. To prove the convergence of the PoF, remark first that $r_n(c)$ concentrates around its mean $\mathbb{E}_{p_c}[r_n(c)]$. Indeed, $r_n(c)$ is a function on the indicator variables $\mathbb{1}[e \in E_c]$ for $e \in E$, which are i.i.d. according to $\text{Ber}(p_c)$. In particular, $r_n(c)$ has a bounded difference of 1, as for any of the indicator variables that changes of value, the rank modifies at most in 1. The McDiarmid concentration inequality implies that,

$$\mathbb{P}\left(|r_n(c) - \mathbb{E}_{p_c}[r_n(c)]| \geq \sqrt{n \log(n)}\right) \leq \exp\left(\frac{-2n \log(n)}{n}\right) = \frac{1}{n^2}.$$

⁵Remark the function is defined over the interval $[0, 1]$. Therefore, Q_n is concave, increasing, $Q(0) = 0$, and $Q(x) \leq x$ for any $x \in [0, 1]$, if and only if the function is 1-Lipschitz.

Added to the union bound, we obtain that,

$$\left| \sum_{c \in [C]} \frac{r_n(c)}{n} - \sum_{c \in [C]} \frac{\mathbb{E}_{p_c}[r_n(c)]}{n} \right| \xrightarrow[n \rightarrow \infty]{P} 0,$$

in other words,

$$\lim_{n \rightarrow \infty} \sum_{c \in [C]} \frac{r_n(c)}{n} = \sum_{c \in [C]} R(p_c).$$

For any $c \in [C]$, notice that,

$$R(1) = \lim_{n \rightarrow \infty} \frac{\mathbb{E}_{p_c=1}[r_n(c)]}{n} = \lim_{n \rightarrow \infty} \frac{\mathbb{E}[r_n(E)]}{n} = \lim_{n \rightarrow \infty} \frac{r_n(E)}{n}.$$

Next, since R is concave,


$$\lim_{n \rightarrow \infty} \frac{\mathbb{E}_{p_c}[r_n(c)]}{n} = R(p_c) = R(p_c \cdot 1 + (1 - p_c) \cdot 0) \geq p_c R(1) + (1 - p_c) R(0) = p_c \lim_{n \rightarrow \infty} \frac{r_n(E)}{n}.$$

Since $R(1) > 0$, we obtain that both $r_n(E) = \Omega(n)$ and $\mathbb{E}_{p_c}[r_n(c)] = \Omega(n)$. Putting all together, we conclude the following,

$$\text{PoF}(M_n) = \max_{\Lambda \subseteq [C]} \frac{r_n([C])}{\sum_{c \in [C]} r_n(c)} \cdot \frac{\sum_{c \in \Lambda} r_n(c)}{r_n(\Lambda)} \xrightarrow[n \rightarrow \infty]{} \max_{\Lambda \subseteq [C]} \frac{R(1)}{\sum_{c \in [C]} R(p_c)} \cdot \frac{\sum_{c \in \Lambda} R(p_c)}{R(\sum_{c \in \Lambda} p_c)},$$

where we used the assumption that R exists and that $R(p_c) > 0$ for all c .

3. Approximation result. To approximate the functions in \mathcal{R} , we will construct a family of matroids able to produce a family of piece-wise functions $f \in \mathcal{R}$ whose convex hull is dense on the set \mathcal{R} . Let $0 \leq b \leq a \leq 1$ be two real values and $n \in \mathbb{N}$, such that an , bn , and $(1 - a)n$ are integer values. Consider the following graph containing a complete bipartite graph with sides of sizes an and bn , respectively, and $(1 - a)n$ isolated vertices, as in the figure below,



Let M_n be the associated transversal matroid. Given a random coloring according to a vector $p = (p_1, \dots, p_C)$, notice that, for n large enough,

$$\mathbb{E}_{p_c}[r_n(c)] = \min\{ap_cn, bn\}.$$

For this sequence of matroids, it follows that,

$$R(p_c) = \lim_{n \rightarrow \infty} \frac{\mathbb{E}_{p_c}[r_n(c)]}{n} = \min\{ap_c, b\}.$$

We denote

$$\mathcal{T} := \{f : [0, 1] \rightarrow \mathbb{R} \mid \exists a, b \in \mathbb{R}_+, f(t) = \min\{at, b\}, \forall t \in [0, 1]\}.$$

In particular, all functions in \mathcal{T} can be obtained by the previous construction. Consider next the set,

$$\mathcal{H} := \{f : [0, 1] \rightarrow \mathbb{R} \mid f \text{ is piece-wise linear, concave, non-decreasing, 1-Lipschitz, and } f(0) = 0\}.$$

We claim that any function in \mathcal{H} can be obtained as convex combinations of functions within \mathcal{T} . Indeed, for $f \in \mathcal{H}$ consisting in two pieces of value a and then $b \leq a$, i.e., such that there exists $t^* \in [0, 1]$,

$$f(t) = \begin{cases} at & t \leq t^*, \\ b(t - t^*) + at^* & t \geq t^*, \end{cases}$$

$$f_1 : [0, 1] \rightarrow \mathbb{R}, \quad f_2 : [0, 1] \rightarrow \mathbb{R}$$

$$t \mapsto at \quad \quad \quad t \mapsto \max(at, at^*)$$

it is enough to take

as $f \equiv \frac{b}{a}f_1 + (1 - \frac{b}{a})f_2$. For the rest of functions within \mathcal{H} , the construction is done inductively. Consider $f \in \mathcal{H}$ to be $(m+1)$ -piece-wise linear, for $m \geq 2$. Let $0 \leq c \leq b \leq a \leq 1$ be the last three linear slopes of f with respective changes at $t_1 \leq t_2$, as illustrated in Figure 12.




Figure 12: Example piece-wise function

Consider next,

$$f_1(t) := \begin{cases} f(t) & t \leq t_1 \\ f(t_1) + at & t \in [t_1, t_2] \\ f(t_1) + a(t_2 - t_1) + ct & t \geq t_2 \end{cases} \quad \text{and} \quad f_2(t) := \begin{cases} f(t) & t \leq t_1 \\ f(t_1) + ct & t \geq t_1 \end{cases}$$

Remark both f_1 and f_2 are m -piece-wise linear. It is not hard to check that

$$f \equiv \left(\frac{b-c}{a-c} \right) f_1 + \left(1 - \left(\frac{b-c}{a-c} \right) \right) f_2.$$

We conclude the proof by showing that \mathcal{H} is dense in \mathcal{R} . Let $R \in \mathcal{R}$ be fixed. For $m \in \mathbb{N}$, divide the interval $[0, 1]$ in m pieces $\{0, \frac{1}{m}, \frac{2}{m}, \dots, \frac{m-1}{m}, 1\}$, and define the m -piece-wise linear function that interpolates R , as it follows,

$$f(t) = R(i) + t \left(R\left(\frac{i+1}{m}\right) - R\left(\frac{i}{m}\right) \right), \text{ for } t \in \left[\frac{i}{m}, \frac{i+1}{m} \right], \text{ for } i \in \{0, 1, \dots, m\}.$$

For $t \in [\frac{i}{m}, \frac{i+1}{m}]$, by monotonicity of R and f , it follows,

$$\begin{aligned} |R(t) - f(t)| &\leq \max \left\{ R\left(\frac{i+1}{m}\right) - f\left(\frac{i}{m}\right); f\left(\frac{i+1}{m}\right) - R\left(\frac{i}{m}\right) \right\} \\ &= R\left(\frac{i+1}{m}\right) - R\left(\frac{i}{m}\right) \\ &\leq \frac{1}{m} \xrightarrow{m \rightarrow \infty} 0, \end{aligned}$$

where the last inequality comes from the fact that R is 1-Lipschitz. In particular, $\|R - f\|_\infty \rightarrow 0$. \square

E.10 Proof of Theorem 4.8

Theorem 4.8. Let $\pi \in [0, 1]$ be fixed. Consider the sets,

$$\mathcal{R} := \{f : [0, 1] \rightarrow \mathbb{R} \mid f \text{ is a concave, non-decreasing, 1-Lipschitz function, and } f(0) = 0\},$$

$$\Delta_\pi^C := \{p \in \Delta^C \mid \max_{c \in [C]} p_c = \pi\}.$$

It follows,

$$\max_{p \in \Delta_\pi^C} \max_{R \in \mathcal{R}} \max_{\Lambda \subseteq [C]} \frac{R(1)}{\sum_{c \in [C]} R(p_c)} \cdot \frac{\sum_{c \in \Lambda} R(p_c)}{R(\sum_{c \in \Lambda} p_c)} = \max_{\lambda \in [C]} \psi_\lambda \left(\frac{1 - (C - \lambda)\pi}{C} \right) \leq C - \frac{1}{\pi}, \quad (9)$$

where $\psi_\lambda : [-\lambda, \frac{1}{C}] \rightarrow \mathbb{R}$, for each $\lambda \in [C]$, is given by,

$$\psi_\lambda(q) = \begin{cases} \lambda & q \in [-\lambda, 0], \\ \frac{\lambda}{(\lambda C q - 1)^2} \cdot (1 + q(1 - 2\lambda) + C(\lambda - 2 + \lambda q)q - 2\sqrt{(\lambda - 1)(C - 1)(1 - Cq)(1 - \lambda q)q}) & q \in \left(0, \frac{(\lambda - 1)}{\lambda(C - 1)}\right], \\ 1 & q \in \left(\frac{(\lambda - 1)}{\lambda(C - 1)}, \frac{1}{C}\right]. \end{cases} \quad (10)$$

The proof of Theorem 4.8 consists of constructing an optimal solution of the triple optimization problem in Equation (9) by starting from an instance (p_0, Λ_0, R_0) and iteratively modifying p , Λ , and R . Before giving the formal proof, we show some useful technical lemmas. We define

$$F : \Delta^C \times \mathcal{R} \times 2^{[C]} \rightarrow [1, \infty)$$

$$(p, R, \Lambda) \rightarrow F(p, R, \Lambda) := \frac{R(1)}{\sum_{c \in [C]} R(p_c)} \cdot \frac{\sum_{c \in \Lambda} R(p_c)}{R(\sum_{c \in \Lambda} p_c)}.$$

Remark that whenever $|\Lambda| \in \{1, C\}$, $F(p, R, \Lambda) = 1$, for any $p, R \in \Delta^C \times \mathcal{R}$. Indeed,

$$F(p, R, \{\bar{c}\}) = \frac{R(1)}{\sum_{c \in [C]} R(p_c)} \leq 1,$$

where the inequality comes from the concavity of R and the fact that $\sum_{c \in [C]} p_c = 1$, so $\sum_{c \in [C]} R(p_c) \leq R(\sum_{c \in [C]} p_c)$. Similarly,

$$F(p, R, [C]) = \frac{R(1)}{R(\sum_{c \in [C]} p_c)} = 1.$$

Therefore, from now on, **we suppose** $1 < |\Lambda| < C$. The function F is invariant to scaling R by non-null constants, i.e., $F(p, R, \Lambda) = F(p, \alpha R, \Lambda)$ for any $\alpha \neq 0$. In addition, F evaluates R at $C + 2$ points: $(p_c)_{c \in [C]}$, $\sum_{c \in \Lambda} p_c$, and 1. Since,

$$F(p, R, \Lambda) = \frac{R(1)}{R(\sum_{c \in \Lambda} p_c)} \left(1 - \frac{\sum_{c \in [C] \setminus \Lambda} R(p_c)}{\sum_{c \in [C]} R(p_c)} \right),$$

F is decreasing on $R(\sum_{c \in \Lambda} p_c)$ and $(R(p_c))_{c \in [C] \setminus \Lambda}$ and increasing on $R(1)$ and $(R(p_c))_{c \in \Lambda}$. Figure 13 illustrates a function $R \in \mathcal{R}$ for $C = 5$ and $\Lambda = \{1, 2, 4\}$, with the red dots indicating the values where F is decreasing, and the blue dots those where F is increasing.




Figure 13: Increasing (blue) and decreasing (red) points for function F

The construction in the proof of Theorem 4.8 will be done by playing with both: the position (over the horizontal axis) of the red and blue dots and their values.

Lemma E.1. Given $(p, R, \Lambda) \in \Delta^C \times \mathcal{R} \times 2^{[C]}$, we can always construct $R' \in \mathcal{R}$ such that either $F(p, R', \Lambda) > F(p, R, \Lambda)$ or $R' = R$.

Proof. Given $(p, R, \Lambda) \in \Delta^C \times \mathcal{R} \times 2^{[C]}$, it is enough with picking $R' \in \mathcal{R}$ satisfying,

$$\begin{aligned} R' \left(\sum_{c \in \Lambda} p_c \right) &\leq R \left(\sum_{c \in \Lambda} p_c \right) \\ R'(p_c) &\leq R(p_c), \forall c \in [C] \setminus \Lambda \\ R(1) &\leq R'(1) \\ R(p_c) &\leq R'(p_c), \forall c \in \Lambda. \end{aligned}$$

For example, suppose $C = 3$, $0 < p_1 < p_2 < p_3 < p_2 + p_3 < 1$, and $\Lambda = \{2, 3\}$. Starting from R we can take $R' \in \mathcal{R}$ such that

$$R'(x) = \begin{cases} x \cdot \frac{R(p_2)}{p_2} & x \in [0, p_2] \\ R(x) & x \in [p_2, p_3] \\ R(p_3) + (x - p_3) \cdot \frac{R(1) - R(p_3)}{1 - p_3} & x \in [p_3, 1] \end{cases}$$

as illustrated in Figure 14,




Figure 14: Function R'

Remark that

$$\begin{aligned} R'(p_2 + p_3) &= R(p_3) + p_2 \cdot \frac{R(1) - R(p_3)}{1 - p_3} \\ &= \frac{p_2}{1 - p_3} \cdot R(1) + \left(1 - \frac{p_2}{1 - p_3}\right) \cdot R(p_3) \\ &\leq R \left(\frac{p_2}{1 - p_3} + \left(1 - \frac{p_2}{1 - p_3}\right) \cdot p_3 \right) \\ &= R \left(\frac{1}{1 - p_3} \cdot p_2(1 - p_3) + p_3 \right) = R(p_2 + p_3), \end{aligned}$$

where the inequality comes from R 's concavity. Similarly,

$$\begin{aligned} R'(p_1) &= \frac{p_1}{p_2} \cdot R(p_2) \\ &= \frac{p_1}{p_2} \cdot R(p_2) + \left(1 - \frac{p_1}{p_2}\right) R(0) \\ &\leq R \left(\frac{p_1}{p_2} \cdot p_2 + \left(1 - \frac{p_1}{p_2}\right) \cdot 0 \right) = R(p_1), \end{aligned}$$

where we have used R 's concavity and that $R(0) = 0$. Finally, remark $R(1) = R'(1)$ and $R'(p_c) = R(p_c)$ for $c \in \Lambda$. \square

Lemma E.2. Let $\pi \in [0, 1]$ be fixed, $(p, R, \Lambda) \in \Delta_\pi^C \times \mathcal{R} \times 2^{[C]}$, and $c^* = \operatorname{argmax}_{c \in \Lambda} p_c$ (if several c^* exists, pick one at random). Consider $p', p'' \in \Delta_\pi^C$ given by,

$$\forall c \in [C], p'_c = \begin{cases} p_c & c \in [C] \setminus \Lambda \text{ or } c = c^* \\ \frac{1}{|\Lambda|-1} \sum_{c \in \Lambda \setminus \{c^*\}} p_c & c \in \Lambda \setminus \{c^*\}, \end{cases}$$

$$\forall c \in [C], p''_c = \begin{cases} p_c & c \in [C] \setminus \Lambda \\ \frac{1}{|\Lambda|} \sum_{c \in \Lambda} p_c & c \in \Lambda. \end{cases}$$

It follows $F(p, R, \Lambda) \leq F(p', R, \Lambda)$ and $F(p, R, \Lambda) \leq F(p'', R, \Lambda)$.

To prove Lemma E.2 we introduce the following definition.

Definition E.3. For $x \in \mathbb{R}_+^C$ a vector, we denote $x_{(c)}$ to its c -th highest entry. Given $x, y \in \mathbb{R}_+^C$, we say that x **majorizes** y if

$$\sum_{c=1}^{\lambda} x_{(c)} \geq \sum_{c=1}^{\lambda} y_{(c)}, \text{ for all } \lambda \in [C], \text{ and } \sum_{c=1}^C x_c = \sum_{c=1}^C y_c.$$

In addition, we state **Kamarata's inequality**: Let $x, y \in \mathbb{R}^C$ be two vectors such that x majorizes y . For any concave function f , it follows,

$$\sum_{c \in [C]} f(x_c) \leq \sum_{c \in [C]} f(y_c).$$

Proof of Lemma E.2. We prove the stated result for p' . For p'' the argument is analogous. Recall that F is increasing on $\sum_{c \in \Lambda} R(p_c)$. We prove that $(p_c)_{c \in \Lambda}$ majorizes $(p'_c)_{c \in \Lambda}$ and conclude by using Karamata's inequality over R . First, remark

$$\sum_{c \in \Lambda} p'_c = p_{c^*} + \sum_{c \in \Lambda \setminus \{c^*\}} \frac{1}{|\Lambda|-1} \sum_{c \in \Lambda \setminus \{c^*\}} p_c = p_{c^*} + \sum_{c \in \Lambda \setminus \{c^*\}} p_c = \sum_{c \in \Lambda} p_c.$$

Regarding the inequality, assume without loss of generality that $\Lambda = \{1, \dots, m\}$ and $p_1 \geq p_2 \geq \dots \geq p_m$. In particular, notice $p'_1 \geq p_1$ (as they are equal). For any $\lambda \in \{2, \dots, m-1\}$, it follows,

$$\begin{aligned} \sum_{c=1}^{\lambda} p'_c &= p_1 + \sum_{c=2}^{\lambda} \frac{1}{m-1} \sum_{c=2}^m p_c \\ &= p_1 + \frac{\lambda-1}{m-1} \sum_{c=2}^m p_c \\ &= p_1 + \frac{\lambda-1}{m-1} \sum_{c=2}^{\lambda} p_c + \frac{\lambda-1}{m-1} \sum_{c=\lambda+1}^m p_c \\ &= p_1 + \sum_{c=2}^{\lambda} p_c - \frac{m-\lambda}{m-1} \sum_{c=2}^{\lambda} p_c + \frac{\lambda-1}{m-1} \sum_{c=\lambda+1}^m p_c \\ &\leq p_1 + \sum_{c=2}^{\lambda} p_c = \sum_{c=1}^{\lambda} p_c, \end{aligned}$$

where the last inequality comes from the fact that,


$$(m-\lambda)(\lambda-1)p_{\lambda} \leq (m-\lambda) \cdot \sum_{c=2}^{\lambda} p_c \text{ and } (\lambda-1) \cdot \sum_{c=\lambda+1}^m p_c \leq (\lambda-1)(m-\lambda)p_{\lambda+1},$$

and therefore,

$$\frac{\lambda-1}{m-1} \sum_{c=\lambda+1}^m p_c - \frac{m-\lambda}{m-1} \sum_{c=2}^{\lambda} p_c \leq \frac{(m-\lambda)(\lambda-1)}{m-1} \cdot (p_{\lambda+1} - p_{\lambda}) \leq 0,$$

for any $\lambda \in \{2, \dots, m-1\}$. □

Lemma E.2 allows to replace all elements $p_c \in \Lambda$ by one single value equal to their mean. In particular, Figure 13 becomes



The main issue with the uniformization of the probabilities within Λ is the fact that any transformation done to the vector $p \in \Delta_\pi^C$ must produce a vector within Δ_π^C , i.e., the maximum-value must remain unchanged (although it could eventually change of index). The following Lemma shows that for any instance $(p, R, \Lambda) \in \Delta_\pi^C \times \mathcal{R} \times 2^{[C]}$, we can always modify R and Λ such that the maximum-value entry of p stays outside of Λ , with a transformation that does not decrease the value of $F(p, R, \Lambda)$.

Lemma E.4. *Let $\pi \in [0, 1]$ be fixed, $(p, R, \Lambda) \in \Delta_\pi^C \times \mathcal{R} \times 2^{[C]}$, and $c^* = \operatorname{argmax}_{c \in [C]} p_c$ (if several c^* exists, pick one at random), i.e., $p_{c^*} = \pi$. Then, we can always construct $(R', \Lambda') \in \mathcal{R} \times 2^{[C]}$ such that $c^* \notin \Lambda'$ and $F(p, R, \Lambda) \leq F(p, R', \Lambda')$.*

Proof. Suppose that $c^* \in \Lambda$. Apply the partial uniformization technique to p of Lemma E.2 leaving p_{c^*} unchanged. Denote

$$q := \frac{1}{|\Lambda| - 1} \sum_{c \in \Lambda \setminus \{c^*\}} p_c.$$

Since $p_{c^*} \in \Lambda$, remark F is increasing at $R(1)$, $R(q)$, and $R(\pi)$. Notice that $0 < q < \pi < \sum_{c \in \Lambda} p_c < 1$. Apply Lemma E.1 and replace R by

$$R'(x) = \begin{cases} R(x) & x \in [0, q] \\ R(q) + (x - q) \frac{R(\pi) - R(q)}{\pi - q} & x \in [q, \pi] \\ R(\pi) + (x - \pi) \frac{R(1) - R(\pi)}{1 - \pi} & x \in [\pi, 1], \end{cases}$$

as illustrated in Figure 15, for $C = 5$, $\Lambda = \{1, 2, 4\}$, and $p_5 = \pi$. Remark that for $x \in [q, \pi]$ no value $R(x)$ is considered on F , which in particular allows to replace R by the linear segment between the points $(q, R(q))$ and $(\pi, R(\pi))$.




Figure 15: Function R' Lemma E.4

Next, we show we can find $\Lambda' \subseteq [C] \setminus \{c^*\}$ and $R'' \in \mathcal{R}$ starting from R' such that $F(p, R', \Lambda) \leq F(p, R'', \Lambda')$. Given $\varepsilon > 0$, consider

$$R'_\varepsilon(x) := \begin{cases} R'(x) & x \in [0, \pi] \\ R'(\pi) + (x - \pi) \left(\frac{R(1) - R(\pi)}{1 - \pi} + \varepsilon \right) & x \in [\pi, 1]. \end{cases}$$

Let

$$\varepsilon^* = \operatorname{argmax} \{ \varepsilon : R'_\varepsilon \in \mathcal{R} \text{ and } F(p, R'_\varepsilon, \Lambda) \geq F(p, R', \Lambda) \},$$

and set $R'' = R'_{\varepsilon^*}$. We claim that

$$\frac{R(1) - R(\pi)}{1 - \pi} + \varepsilon^* = \frac{R(\pi) - R(q)}{\pi - q},$$

i.e., the segment between $(q, R''(q))$ and $(\pi, R''(\pi))$ has the same slope as the one between $(\pi, R''(\pi))$ and $(1, R''(1))$, as illustrated in Figure 16, for $C = 5$, $\Lambda = \{1, 2, 4\}$, and $p_5 = \pi$. Clearly, R'_{ε^*} belongs to \mathcal{R} . Regarding the increase on the value of the function F , we show that the




Figure 16: Function R''

for any $\varepsilon > 0$,

$$R'_\varepsilon\left(\sum_{c \in \Lambda} p_c\right) - R'\left(\sum_{c \in \Lambda} p_c\right) \leq R'_\varepsilon(1) - R'(1),$$

i.e., that the increase of the blue dot in Figure 16 is larger than the increase of the red dot. It follows,

$$\begin{aligned} \frac{d}{d\varepsilon} \left[\frac{R'_\varepsilon(1)}{R'_\varepsilon(\sum_{c \in \Lambda} p_c)} \right] &= \frac{d}{d\varepsilon} \left[\frac{R'(\pi) + (1 - \pi) \left(\frac{R(1) - R(\pi)}{1 - \pi} + \varepsilon \right)}{R'(\pi) + (\sum_{c \in \Lambda} p_c - \pi) \left(\frac{R(1) - R(\pi)}{1 - \pi} + \varepsilon \right)} \right] \\ &= \frac{R'(\pi)(1 - \sum_{c \in \Lambda} p_c)}{\left(R'(\pi) + (\sum_{c \in \Lambda} p_c - \pi) \left(\frac{R(1) - R(\pi)}{1 - \pi} + \varepsilon \right) \right)^2}, \end{aligned}$$

which is always non-negative. In particular, R'' is rewritten as

$$R''(x) := \begin{cases} R(x) & x \in [0, q] \\ R(q) + (x - q) \left(\frac{R(\pi) - R(q)}{\pi - q} \right) & x \in [q, 1]. \end{cases} \quad (11)$$

To ease the notation, we drop the $''$ from R'' and denote $\alpha = (R(\pi) - R(q))/(\pi - q)$. Finally, we construct $\Lambda' \subseteq [C] \setminus \{c^*\}$ such that $F(p, R, \Lambda) \leq F(p, R, \Lambda')$. The analysis is split depending on whether a value $p_{\bar{c}}$ with $\bar{c} \in [C] \setminus \Lambda$ (a red dot) lies between q and π or not. Suppose it does. We claim that considering $\Lambda' := \Lambda \setminus \{c^*\} \cup \{\bar{c}\}$ we obtain the stated result. Indeed, although swapping the elements should decrease the value of the function F (as we obtain a higher-value red dot and a lower-value blue dot), the effect is compensated by the fact that $\sum_{c \in \Lambda'} p_c < \sum_{c \in \Lambda} p_c$. Figure 17 illustrates the swapping.




Figure 17: Swapping $p_{\bar{c}}$ and π within Λ

To prove that $F(p, R, \Lambda) \leq F(p, R, \Lambda')$, we check that

$$\frac{\sum_{c \in \Lambda} R(p_c)}{R(\sum_{c \in \Lambda} p_c)} \leq \frac{\sum_{c \in \Lambda'} R(p_c)}{R(\sum_{c \in \Lambda'} p_c)}. \quad (12)$$

For $z \in [p_{\bar{c}}, \pi]$, consider,

$$Q(z) := \frac{(|\Lambda| - 1)R(q) + R(\pi - p_{\bar{c}} + z)}{R((|\Lambda| - 1)q + \pi - p_{\bar{c}} + z)} = \frac{|\Lambda|R(q) + \alpha(\pi - p_{\bar{c}} + z - q)}{R(q) + \alpha((|\Lambda| - 2)q + \pi - p_{\bar{c}} + z)}$$

where the last equality comes from using R 's definition (11). In particular Equation (12) holds if and only if $Q(\pi) \leq Q(p_{\bar{c}})$. Notice,

$$\begin{aligned} \frac{d}{dz} Q(z) &= \frac{(R(q) + \alpha((|\Lambda| - 2)q + \pi - p_{\bar{c}} + z))\alpha - (|\Lambda|R(q) + \alpha(\pi - p_{\bar{c}} + z - q))\alpha}{[R(q) + \alpha((|\Lambda| - 2)q + \pi - p_{\bar{c}} + z)]^2} \\ &= \frac{\alpha(|\Lambda| - 1)(\alpha q - R(q))}{[R(q) + \alpha((|\Lambda| - 2)q + \pi - p_{\bar{c}} + z)]^2} \leq 0, \end{aligned}$$

where we have used that $R(q) \geq \alpha q$, which holds as $\alpha \leq 1$ is the slope of the last piece-wise part of the function R , which extended up to the origin remains positive, in particular implying that the image of 0 (given by $R(q) - \alpha q$) is at least 0. We conclude $Q(z)$ is decreasing over $[p_{\bar{c}}, \pi]$, concluding that Equation (12) holds.

To finish the proof, suppose that such as $p_{\bar{c}}$ did not exist, as in Figure 16. Keep increasing the slope of the last pice-wise linear function until achieving the slop between q and the closest red dot placed at the left of q , namely $p_{\underline{c}}$, as in Figure 18 and set $\Lambda' := \Lambda \setminus \{\bar{c}\} \cup \{\underline{c}\}$.




Figure 18: Final construction Λ'

As in the previous cases, it can be proved that $F(p, R, \Lambda) \leq F(p, R, \Lambda')$. We omit the proof. \square

Lemma E.5. Let $\pi \in [0, 1]$ be fixed and $(p, R, \Lambda) \in \Delta_\pi^C \times \mathcal{R} \times 2^{[C]}$. Define

$$\Gamma := \left\{ k \in [C] \setminus \Lambda : p_k \leq \frac{1}{|\Lambda|} \sum_{c \in \Lambda} p_c \right\}.$$

There exists $(p', R') \in \Delta_\pi^C \times \mathcal{R}$ such that $F(p, R, \Lambda) \leq F(p', R', \Lambda \cup \Gamma)$.

Proof. Apply Lemmas E.2 and E.4 so the entry of p of value π is not included in Λ and for any $c \in \Lambda$, $p_c = q := \frac{1}{|\Lambda|} \sum_{c \in \Lambda} p_c$. In particular, the only values where F is increasing are $R(q)$ and $R(1)$. Define $\Gamma := p_c < q$. It follows that F is decreasing at $(R(p_c))_{c \in \Gamma}$. Replace R by

$$R(x) = \begin{cases} x \frac{R(q)}{q} & x \in [0, q] \\ R(x) & x \in [q, \pi]. \end{cases}$$

Moreover, since $F(p, \alpha R, \Lambda) = F(p, R, \Lambda)$ for any $\alpha \neq 0$, redefine $R \equiv \frac{q}{R(q)} R$. Finally, since for any $q < p_c < 1$ the function F is decreasing on $R(p_c)$, replace R by,

$$R(x) = \begin{cases} x & x \in [0, q] \\ q + (x - q) \cdot \frac{R(1) - q}{1 - q} & x \in [q, \pi], \end{cases}$$

where we have used that $R(x) = x$ for any $x \leq q$ because of the previous scaling. The resulting function is illustrated in Figure 19 for $\Gamma = \{1, 2, 3\}$.




Figure 19: Function R Lemma E.5

Finally, we prove that $F(p, R, \Lambda) \leq F(p, R, \Lambda \cup \Gamma)$. For $z \in [0, q]$, consider

$$Q(z) := \frac{|\Lambda|R(q) + R(z)}{R(|\Lambda|q + z)} = \frac{|\Lambda|q + z}{q + (|\Lambda|q + z - q)\alpha},$$

where $\alpha = \frac{R(1)-q}{1-q}$. Remark the previous claim holds if and only if $Q(0) \leq Q(z)$, i.e., adding elements from Γ to Λ increases the part of F that depends on Λ . We obtain,

$$\begin{aligned} \frac{d}{dz}Q(z) &= \frac{d}{dz} \left[\frac{|\Lambda|q + z}{q + (|\Lambda|q + z - q)\alpha} \right] \\ &= \frac{(q + (|\Lambda|q + z - q)\alpha) - (|\Lambda|q + z)\alpha}{(q + (|\Lambda|q + z - q)\alpha)^2} \\ &= \frac{q(1 - \alpha)}{(q + (|\Lambda|q + z - q)\alpha)^2}, \end{aligned}$$

which is always positive. We conclude the proof. \square

Lemma E.6. Let $\pi \in [0, 1]$ be fixed and $(p, R, \Lambda) \in \Delta_\pi^C \times \mathcal{R} \times 2^{[C]}$. Then, there always exists $(p', R', \Lambda') \in \Delta_\pi^C \times \mathcal{R} \times 2^{[C]}$ such that,

$$F(p', R', \Lambda') = \frac{\lambda}{\alpha(\lambda - 1) + 1} \cdot \frac{\alpha + (1 - \alpha)q}{\alpha + (1 - \alpha)Cq} =: \hat{F}(q),$$

where $\lambda = |\Lambda'|$, $q = \frac{1}{\lambda} \sum_{c \in \Lambda'} p_c$, $\alpha = \frac{R(1)-q}{(1-q)}$, and $F(p, R, \Lambda) \leq F(p', R', \Lambda')$. In particular,

$$\operatorname{argmax}_{q \in [0, 1]} \hat{F}(q) = \begin{cases} 0 & \frac{1 - (C - \lambda)\pi}{\lambda} \end{cases}$$

Proof. Let $\pi \in [0, 1]$ be fixed and $(p, R, \Lambda) \in \Delta_\pi^C \times \mathcal{R} \times 2^{[C]}$. Apply Lemmas E.1, E.2, E.4 and E.5 to construct $(p', R', \Lambda') \in \Delta_\pi^C \times \mathcal{R} \times 2^{[C]}$ such that,

$$F(p, R, \Lambda) \leq F(p', R', \Lambda')$$

$$\text{for any } c \in \Lambda', p'_c = q := \frac{1}{\lambda} \sum_{c \in \Lambda'} p_c$$

$$\text{for any } c \in [C] \setminus \Lambda', p'_c \geq q,$$

$$R'(x) = \begin{cases} x & x \in [0, q] \\ q + \alpha(x - q) & x \in [q, 1] \end{cases}$$

Moreover, $c^* \in [C]$ such that $p_{c^*} = \pi$, is not included in Λ' . It is not hard to see that,

$$F(p', R', \Lambda') = \hat{F}(q) = \frac{\lambda}{\alpha(\lambda - 1) + 1} \cdot \frac{\alpha + (1 - \alpha)q}{\alpha + (1 - \alpha)Cq}.$$

Remark that $F(p', R', \Lambda')$ is decreasing on q . In particular, making $q \rightarrow q - \varepsilon$ increases the value of F . However, remark the value of q defines the kink of the function R' . In addition, any decreasing on q implies to decrease the value on the entries of p within Λ' . Since p is a probability distribution, the decrease of mass must be re-injected on all other entries whose values are below π , as we cannot modify the value of the highest-value entry of p . In conclusion, whenever solving

$$\max_{q \in [0,1]} \hat{F}(q),$$

we obtain the solution

$$q = \begin{cases} 0 \\ \frac{1-(C-\lambda)\pi}{\lambda} \end{cases}$$

where the second case comes from attaining the constraint of maximizing all entries $c \in [C] \setminus \{\Lambda \cup \{c^*\}\}$ up to π . Remark that when the previous optimization problem we do not consider anymore the space of functions \mathcal{R} , in particular allowing for $q = 0$ to be a possible solution. \square

We are ready to prove Theorem 4.8.

Proof of Theorem 4.8. For $\lambda \in [C]$ and $\alpha \in [0, 1]$, consider the function

$$\psi_\lambda(\alpha, q) = \frac{\lambda}{\alpha(\lambda - 1) + 1} \cdot \frac{\alpha + (1 - \alpha)q}{\alpha + (1 - \alpha)Cq}.$$

From Lemma E.6, it follows,

$$\max_{p \in \Delta_\pi^C} \max_{R \in \mathcal{R}} \max_{\Lambda \subseteq [C]} \frac{R(1)}{\sum_{c \in [C]} R(p_c)} \cdot \frac{\sum_{c \in \Lambda} R(p_c)}{R(\sum_{c \in \Lambda} p_c)} = \max_{\lambda \in [C]} \max_{\alpha \in [0,1]} \max_{q \in [0,1]} \psi_\lambda(\alpha, q). \quad (13)$$

We know that for $(\lambda, \alpha) \in [C] \times [0, 1]$,

$$\operatorname{argmax}_{q \in [0,1]} \psi_\lambda(\alpha, q) = \begin{cases} 0 \\ \frac{1-(C-\lambda)\pi}{\lambda} \end{cases}$$

Suppose $1 - (C - \lambda)\pi \leq 0$. It follows,

$$\max_{p \in \Delta_\pi^C} \max_{R \in \mathcal{R}} \max_{\Lambda \subseteq [C]} \frac{R(1)}{\sum_{c \in [C]} R(p_c)} \cdot \frac{\sum_{c \in \Lambda} R(p_c)}{R(\sum_{c \in \Lambda} p_c)} \leq \max_{\lambda \in [C]} \max_{\alpha \in [0,1]} \frac{\lambda}{\alpha(\lambda - 1) + 1} = \max_{\lambda \in [C]} \lambda.$$

Suppose $1 - (C - \lambda)\pi \geq 0$, i.e., $q \in [0, 1/C]$. We study the first order conditions of $\psi_\lambda(\alpha, q)$ over α . It follows,

$$\begin{aligned} \frac{d}{d\alpha} \psi_\lambda(\alpha, q) = & -\frac{\lambda}{(\alpha\lambda - \alpha + 1)^2 (\alpha Cq - \alpha - Cq)^2} [\alpha^2 (C\lambda q^2 - C\lambda q - Cq^2 + Cq - \lambda q + \lambda + q - 1) \\ & + \alpha(-2C\lambda q^2 + 2Cq^2 + 2\lambda q - 2q) + C\lambda q^2 - Cq^2 - Cq + q]. \end{aligned}$$

Imposing $\frac{d}{d\alpha} \psi_\lambda(\alpha, q) = 0$ we obtain the solutions,

$$\begin{aligned} \alpha_1 &= \frac{2(-1 + \lambda)q(-1 + Cq) - \sqrt{4(-1 + C)(-1 + \lambda)q(-1 + Cq)(-1 + \lambda)q}}{2(q + Cq(-1 + (-1 + \lambda)q))}, \\ \alpha_2 &= \frac{2(-1 + \lambda)q(-1 + Cq) + \sqrt{4(-1 + C)(-1 + \lambda)q(-1 + Cq)(-1 + \lambda)q}}{2(q + Cq(-1 + (-1 + \lambda)q))}. \end{aligned}$$

Since $q \leq 1/C$, it follows that $2(-1 + \lambda)q(-1 + Cq) \leq 0$ and, therefore, $\alpha_1 < 0$. The only possible solution being α_2 , we show that either

1. $\alpha_2 \in [0, 1]$ and then the stated value of $\psi_\lambda(q)$ for $q \in (0, \frac{\lambda-1}{\lambda(C-1)})$ comes from plugging α_2 into Equation (13) or,
2. $\alpha_2 \geq 1$ and then Equation (13) is upper bounded by 1 for any $q \in (\frac{\lambda-1}{\lambda(C-1)}, \frac{1}{C}]$.

The first point is direct. For the second point, notice that $\alpha_2 \geq 1$ if and only if

$$\sqrt{4(C-1)(\lambda-1)q(Cq-1)(\lambda q-1)} \geq 2(\lambda-1)(p-1)(Cq-1) + 2(\lambda-1)q(1-Cq),$$

and, as $2(\lambda-1)(p-1)(Cq-1) + 2(\lambda-1)q(1-Cq) \geq 0$, this is equivalent to

$$\begin{aligned} 4(C-1)(\lambda-1)q(Cq-1)(\lambda q-1) &\geq (2(\lambda-1)(p-1)(Cq-1) + 2(\lambda-1)q(1-Cq))^2 \\ \iff 4(\lambda-1)(1-q)(1-Cq)(\lambda((C-1)q-1)+1) &\geq 0 \\ \iff q \in \left(\frac{\lambda-1}{\lambda(C-1)}, \frac{1}{C} \right]. \end{aligned}$$

Since $\alpha_2 \geq 1$, the optimal value of α is 1, yielding

$$\max_{\alpha \in [0,1]} \max_{q \in [0,1]} \psi_\lambda(\alpha, q) = \max_{q \in \left(\frac{\lambda-1}{\lambda(C-1)}, \frac{1}{C} \right]} \psi_\lambda(1, q) = 1.$$

To conclude the proof, set $\psi_\lambda(q) := \max_{\alpha \in [0,1]} \psi_\lambda(\alpha, q)$. The relaxed upper bound

$$\max_{\lambda \in [C]} \psi_\lambda \left(\frac{1 - (C-\lambda)\pi}{C} \right) \leq C - \frac{1}{\pi},$$

is obtained through symbolic computation in Mathematica (with the Reduce function). Indeed, it can be verified that the inequality system

$$\frac{\lambda(Cq-1)}{\lambda q-1} - \frac{\lambda \left(Cq(\lambda q + \lambda - 2) - 2\sqrt{(C-1)(\lambda-1)q(Cq-1)(\lambda q-1)} - 2\lambda q + q + 1 \right)}{(C\lambda q-1)^2} \geq 0$$

for $q \in \left[0, \frac{\lambda-1}{C\lambda-\lambda} \right]$ and $\lambda \in [2, C-1]$,

is always feasible. It follows that $\lambda(Cq-1)/(\lambda q-1) \geq \psi_\lambda(q)$ for $0 \leq q \leq (\lambda-1)/(C\lambda-\lambda)$. In particular, for $q = (1 - (C-\lambda)\pi)/\lambda$, it follows that $C - 1/\pi \geq \psi_\lambda(q)$ over $[1/(C-1), 1/(C-\lambda)]$. Similarly, since $C - 1/\pi$ is greater than $\lambda = \psi_\lambda(q)$ whenever $\pi \geq 1/(C-\lambda)$, we conclude $C - 1/\pi \geq \psi_\lambda(q)$ for any $\lambda \in [C]$ and $q \in [0, 1]$. \square

E.11 A technical lemma

The following technical lemma gives a sufficient condition for a polymatroid to have a price of opportunity fairness equal to 1. In particular, several of the posterior results in the stochastic setting will use it.

Lemma E.7. *Let M be a polymatroid. Given a permutation $\sigma \in \Sigma([C])$, consider the sequence $\mathbf{r}(\sigma) = (\mathbf{r}_c(\sigma))_{c \in [C]}$ such that,*

$$\text{for any } c \in [C], \mathbf{r}_c(\sigma) := \frac{\mathbf{r}(\sigma(1, \dots, c)) - \mathbf{r}(\sigma(1, \dots, c-1))}{\mathbf{r}(\sigma(c))},$$

where, $\mathbf{r}(\sigma(1, \dots, c))$ corresponds to the size of a maximum size allocation in the submatroid obtained by the groups in the first c entries of $\sigma([C])$. Whenever the sequences $\mathbf{r}(\sigma)$ for any $\sigma \in \Sigma([C])$, are all decreasing, it holds $\text{PoF}(M) = 1$.

Proof. Let $\Lambda^* = \arg\max_{\Lambda \subseteq [C]} \frac{\sum_{c \in \Lambda} \mathbf{r}(c)}{\mathbf{r}(\Lambda)}$. We aim at proving that the monotonicity of the sequences $\{\mathbf{r}(\sigma), \sigma \in \Sigma([C])\}$ implies $\Lambda^* = [C]$, which yields $\text{PoF}(M) = 1$. Without loss of generality, take $\sigma = \text{Id}_C$ to be the identity permutation (the same argument works for any other permutation). Denote

$$\rho_t := \frac{\mathbf{r}([t])}{\sum_{\ell \in [t]} \mathbf{r}(\ell)},$$

the competition index of the submatroid obtained by the the first t groups. Denoting $\mathbf{r}(0) = 0$, it follows,

$$\rho_{t+1} - \rho_t = \frac{\mathbf{r}([t+1])}{\sum_{\ell \in [t+1]} \mathbf{r}(\ell)} - \frac{\mathbf{r}([t])}{\sum_{\ell \in [t]} \mathbf{r}(\ell)}$$

$$\begin{aligned}
&= \frac{\sum_{\ell \in [t+1]} r(\ell) - r(\ell-1)}{\sum_{\ell \in [t+1]} r(\ell)} - \frac{\sum_{\ell \in [t]} r(\ell) - r(\ell-1)}{\sum_{\ell \in [t]} r(\ell)} \\
&= \frac{\sum_{\ell \in [t]} [r(t+1) - r(t)] r(\ell) - r(t+1) [r(\ell) - r(\ell-1)]}{(\sum_{\ell \in [t+1]} r(\ell)) (\sum_{\ell \in [t]} r(\ell))}.
\end{aligned}$$

Since $r(\sigma)$ is decreasing, for any $s < t+1$ it follows,

$$\frac{r(s) - r(s-1)}{r(s)} \geq \frac{r(t+1) - r(t)}{r(t+1)}.$$

In particular, $r(t+1)[r(s) - r(s-1)] \geq [r(t+1) - r(t)]r(s)$ and therefore, $\rho_t \geq \rho_{t+1}$. It follows that the optimal solution corresponds to $\Lambda^* = [C]$. \square

E.12 Proof of Propositions 4.10 and 4.11

Proposition 4.10. Let $\omega = \omega(n)$ be a function such that $\omega(n) \rightarrow \infty$. Whenever $q \leq 1/(\omega n)$ or $q \geq \omega/n$, for any $p \in \Delta^C$, $\text{PoF}(G_{n,q}(p))$ converges to 1 with high probability as n grows.

Proof. We show this proposition by leveraging results from random graph theory. Suppose $q \leq 1/(\omega n)$. By Theorem 2.1 [37], $G_{n,q}$ is a forest w.h.p.. It follows that, independent of the label realization, the maximal allocation contain all edges in the graph. Hence $\sum_{c \in [C]} r(c) = r([C])$, which implies the matroid has PoF equal to 1 as the independence index is equal to 1.

Suppose $q \geq \omega/n$. For any $c \in [C]$ the subgraph induced by considering only the subset E_c over $G_{n,q}$ is distributed according to $G_{n,p_c q}$. Since $p_c q \geq p_c \omega/n$, with $p_c \omega \rightarrow \infty$ arbitrarily slow, Theorem 2.14 [37] states that w.h.p. $G_{n,p_c q}$ has a giant component of size $(1 - \frac{x}{p_c \omega})n$, for a fixed $x \in [0, 1]$. In particular, $r(c) = (1 - \frac{x}{p_c \omega})n$ as connected components contain spanning trees. Intersecting the events over all $c \in [C]$ we obtain, w.h.p. as n goes to infinity,

$$\rho(G_{n,q}(p)) = \frac{r([C])}{\sum_{c \in [C]} r(c)} = \frac{n}{\sum_{c \in [C]} n} + o(1) = \frac{1}{C} + o(1),$$

which from Proposition 4.6 shows that PoF is also equal to 1 w.h.p.. \square

Proposition 4.11. Let $\omega = \omega(n)$ be a function such that $\omega(n) \rightarrow \infty$ arbitrarily slow as $n \rightarrow \infty$. Whenever $q \leq 1/(\omega n^{3/2})$ or $q \geq \omega \log(n)/n$, for any $p \in \Delta^C$, $\text{PoF}(B_{n,\beta,q}(p))$ converges to 1 with high probability as n grows.

Proof. Suppose $q \geq \omega \log(n)/n$. Let $\Lambda \subseteq [C]$, we have that $\sum_{c \in [C]} |E_c|$ is a sum of independent bernoulli random variables (the E_c are disjoint), hence it has an expected value of $n \sum_{c \in \Lambda} p_c$ and Hoeffding's concentration inequality show that

$$\mathbb{P}\left(\left|\sum_{c \in \Lambda} |E_c| - n \sum_{c \in \Lambda} p_c\right| > \sqrt{n \log(n)}\right) \leq 2 \exp\left(-2 \frac{n \log(n)}{n}\right) = \frac{2}{n^2} \xrightarrow{n \rightarrow \infty} 0.$$

Since $q \geq \omega \log(n)/n$, Theorem 6.1 [37] states that w.h.p. for any $\Lambda \subseteq [C]$, the subgraph considering only the vertices in Λ on the left-hand side has a matching of size $\min\{\beta n, \sum_{c \in \Lambda} p_c n\}$, therefore, $r(\Lambda) = \min\{\beta n, \sum_{c \in \Lambda} p_c n\}$. We will conclude by applying Lemma E.7. As usual, w.l.o.g. consider $\sigma = I_L$. For any $c \in [C-1]$, it follows,

$$r_{c+1}(\sigma) = \frac{\min\{\beta n, \sum_{c' \in [c+1]} p_{c'} n\} - \min\{\beta n, \sum_{c' \in [c]} p_{c'} n\}}{\min\{\beta n, p_c n\}}.$$

In particular, as $\sum_{c' \in [c]} p_{c'}$ is increasing in c , the sequence $r_{c+1}(\sigma)$ initially consists on only 1 (given by all times that $r_{c+1}(\sigma) \leq \beta$), eventually some value between 0 and 1 (given by the first time that $r_{c+1}(\sigma) \geq \beta \geq r_c(\sigma)$), and finally a sequence of only zeros (given by all times when $r_{c+1}(\sigma) > \beta$). In particular, the sequence is decreasing, concluding the proof. \square

Remark 2. In both graphic and transversal random matroids, taking the same $q \in [0, 1]$ for all colors is done without loss of generality. Indeed, a coupling argument based on stochastic dominance allows us to consider edge probabilities q_c per color c and to obtain the same results.

F Price of Fairness under Other Fairness Notions

F.1 Weighted Fairness

The main fairness definition that we have used is opportunity fairness. We now discuss how this specific fairness notion relates with other fairness concepts, in particular with maximin fairness and proportionality in [11] and equitability in [15]. We can think more generally about group fairness in terms of what amount of social welfare protected group of agents are entitled to. Should each group be entitled to the same amount as others, or proportionally to their size? We introduce weighted fairness, where the weights correspond to group entitlement:

Definition F.1. Let $(w_c)_{c \in [C]} \in \mathbb{R}_+^C$ be a fixed weights vector. An allocation $x \in \mathbb{R}_+^C$ is w -fair if for any $i, j \in [C]$, $x_i/w_i = x_j/w_j$.

As an example, Figure 20 illustrates for a 2-colored matroid three fairness notions mentioned in the paper that are now framed as specific instances of weighted fairness.




Figure 20: Weighted Fairness for matroid with two groups

Compared to other weighted fairness notions, opportunity fairness remains bounded because the weights depend on the structure of the polymatroid M , while the weights of proportional fairness and equitability are independent of M and arbitrarily bad examples can easily be constructed.

Proposition F.2. *The price of proportional fairness and the price of equitability are unbounded in the worst-case, even when allowing for fractional allocations.*

Proof. Take a ground set with n agents of color 1 and n agents of color 2, and consider the feasible allocations where at most $n/2$ resources can be allocated to individuals of color 2 but only one resource may be allocated to individuals of color 1. This is a partition matroid, where the partition corresponds exactly to the color partition. Now, the optimal proportionally fair allocation, as well as the optimal equitable allocation, is $x_1 = 1$ and $x_2 = 1$, for a total of 2 resources allocated. Because the optimal allocation is of size $n + 1$, the price of fairness is $(n + 1)/2$, which goes to ∞ as $n \rightarrow \infty$. The price of fairness in both cases is unbounded. \square

Another common concept of fairness to divide resources, used in transferable utility cooperative game theory, is that of Shapley value [58]: it is the unique utility transfer that satisfies axioms of symmetry, additivity, nullity and efficiency. It can be shown that for $\Sigma([C])$ the set of permutations over $[C]$, the Shapley value of group c is

$$\varphi_c := \frac{1}{C!} \sum_{\sigma \in \Sigma([C])} r(\{i \in [C] \mid \sigma(i) < \sigma(c)\} \cup \{c\}) - r(\{i \in [C] \mid \sigma(i) < \sigma(c)\}),$$

that is to say φ_c is the expected marginal contribution of group c when groups are prioritized according to σ a uniformly drawn random permutation. When $w_c = \varphi_c$, we say that an allocation is Shapley fair.

The allocation problem we study can be seen as a type of non transferrable utility game, and as such there is no reason in general for the allocation $(\varphi_1, \dots, \varphi_C)$ to be realizable. Nonetheless, from the polymatroid characterization of M do have this property:

Proposition F.3. *The allocation $(\varphi_1, \dots, \varphi_C)$ is always feasible.*

Proof. For a given permutation σ , the marginal contribution allocation x^σ where $x_c^\sigma = r(\{i \in [C] \mid \sigma(i) < \sigma(c)\} \cup \{c\}) - r(\{i \in [C] \mid \sigma(i) < \sigma(c)\})$ is always feasible. Hence, the Shapley allocation $(\varphi_1, \dots, \varphi_C)$ the barycenter of all the x^σ , which belong to the Pareto front by definition. Moreover, by the polymatroid characterization, the Pareto front is convex [43]. Hence the barycenter, being a convex combination, is also feasible. We note that the x^σ are the extreme points of the Pareto front. \square

From the efficiency of the Shapley allocation, it is immediate that the Shapley Price of Fairness is always 1 for colored matroids.

In the semi-random model of Theorem 4.8, we can easily show the following property:

Proposition F.4. *For any distribution $p \in \Delta^C$ of the agents colors, in the large market setting with $\liminf r_n([C]) = \Omega(n)$, we have that the price of proportional fairness converges to 1 with high probability.*

Proof. Let S_n be any maximal allocation, taken independently of the random coloring. We have that $|S_n| = \Omega(n)$ and $|S_n \cap E_c|$ concentrates towards $p_c |S_n|$ by Hoeffding's inequality. We also have that $|E_c|$ concentrates around $p_c n$. Hence $|S_n \cap E_c|/|E_c|$ concentrates around $|S_n|/n$, which is independent of the colors, and therefore is a proportionally fair allocation. Moreover S_n is maximal by definition, so the price of fairness is equal to 1. \square

This shows that the price of proportional fairness goes from $\text{PoF} = +\infty$ in the adversarial setting to $\text{PoF} = 1$ in the semi-random setting.

Finally let us mention another fairness definition.

F.2 Leximin Fairness

Requiring that a fair allocation satisfies exactly $x_i/w_i = x_j/w_j$ can be considered wasteful, as it is possible to improve the total social welfare without making any group worst off. Maximin fairness [11], also called egalitarian rule or Rawlsian fairness, corresponds to ensuring that the worst off group has the best allocation possible. In other words, an allocation is maximin fair if it maximizes $\min_{x \in M} x_c$, or with entitlement w , maximizes $\min_{x \in M} x_c/w_c$. Most of the time there are multiple maximin fair feasible allocations, and thus one may seek to maximize the second minimum, and so forth. This is called the leximin rule, and has also been studied in the social choice literature [25, 27].

For a vector $x = (x_1, \dots, x_C)$, we denote the ordered coordinates by $x_{(1)} \geq x_{(2)} \geq \dots \geq x_{(C)}$. We say that a vector $x = (x_1, \dots, x_C)$ is leximin larger than $y = (y_1, \dots, y_C)$ if $x_{(C)} \geq y_{(C)}$, or $x_{(C)} = y_{(C)}$ and $x_{(C-1)} \geq y_{(C-1)}$, or $x_{(C)} = y_{(C)}$ and $x_{(C-1)} = y_{(C-1)}$ and $x_{(C-2)} \geq y_{(C-2)}$ and so on. The leximin order is a total preorder. Leveraging the more general notion of weighted fairness, we have the following definition:

Definition F.5. For a weight vector $w \in \mathbb{R}_+^C$, an allocation (x_1, \dots, x_C) is said to be w -lexmaxmin fair if $(\frac{x_1}{w_1}, \dots, \frac{x_C}{w_C})$ is maximal according to the leximin order for $x \in M$.

Clearly, the w -lexmaxmin fair allocation is w -maxmin fair. It is also Pareto efficient, and therefore by the polymatroid characterization Proposition 3.2 achieves maximal social welfare: the price of w -lexmaxmin fairness is always 1 in C -colored matroids.