Asymptotic Existence of Class Envy-free Matchings

Tomohiko Yokoyama
The University of Tokyo
Tokyo, Japan
tomohiko_yokoyama@mist.i.u-tokyo.ac.jp

Ayumi Igarashi
The University of Tokyo
Tokyo, Japan
igarashi@mist.i.u-tokyo.ac.jp

ABSTRACT

We consider a one-sided matching problem where agents who are partitioned into disjoint classes and each class must receive fair treatment in a desired matching. This model, proposed by Benabbou et al. [9], aims to address various real-life scenarios, such as the allocation of public housing and medical resources across different ethnic, age, and other demographic groups. Our focus is on achieving class envy-free matchings, where each class receives a total utility at least as large as the maximum value of a matching they would achieve from the items matched to another class. While class envy-freeness for worst-case utilities is unattainable without leaving some valuable items unmatched, such extreme cases may rarely occur in practice. To analyze the existence of a class envy-free matching in practice, we study a distributional model where agents' utilities for items are drawn from a probability distribution. Our main result establishes the asymptotic existence of a desired matching, showing that a round-robin algorithm produces a class envy-free matching as the number of agents approaches infinity.

KEYWORDS

Matching; Fairness; Envy-freeness; Asymptotics

ACM Reference Format:
Tomohiko Yokoyama and Ayumi Igarashi. 2025. Asymptotic Existence of Class Envy-free Matchings. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 9 pages.

1 INTRODUCTION

One-sided matching is a fundamental problem that forms the economic foundation of numerous practical applications, spanning from kidney exchange [33], assigning drivers to customers [8], to house allocation [1, 21, 41]. An instance of the one-sided matching problem consists of a set of agents, a set of indivisible items, and preferences of the agents over the items. The goal is to find an assignment of items to agents while ensuring desirable normative properties. In particular, guaranteeing fairness is of paramount importance in scenarios such as allocating tasks to workers or providing social housing to residents.

A substantial body of the literature is dedicated to ensuring fairness among individuals, focusing on concepts such as *envy-freeness* [11]. Envy-freeness ensures that no agent prefers another agent's allocated item to their own. On the other hand, when addressing various real-life resource allocation problems, it becomes increasingly crucial to consider fairness requirements among different classes of agents, particularly when implementing solutions in large-scale systems. For instance, consider the allocation of scarce medical resources to different regions in a country. Another motivating example involves student placement in public schools among different ethnic groups and social housing allocation among different income groups. To gain public acceptance, the social planner must ensure an equitable allocation of resources across various regional groups or classes.

The scenarios we consider are effectively captured by a model proposed by Benabbou et al. [9]. This model consists of m items and n agents, divided into k disjoint classes, with each agent allocated at most one item. A key feature of this framework is the use of *assignment valuations* to evaluate *envy* among classes. Assignment valuations quantify the potential value a class could derive from the set of items allocated to another class, determined by the optimal matching between the items and the members of the class. By adapting fairness notions from fair division literature into the one-sided matching problem, Benabbou et al. [9] introduced the notion of *class envy-freeness* [9, 20],[1] which requires that no class prefers the set of items assigned to any other class over its own bundle, in the sense of assignment valuations.

Unfortunately, a class envy-free matching is not guaranteed to exist without wasting any items. Consider a simple example with one item and two single-agent classes, both valuing the item. In any nonempty matching, one class receives nothing whereas the other class receives one item, violating class envy-freeness. However, this strong conflict between fairness and efficiency does not preclude the existence of practical solutions that strike a balance between the two. The question then becomes: can we achieve a fair and efficient matching on average cases?

In the context of fair division, related works have examined the existence of envy-free allocations in probabilistic settings, where agents' utilities for items are modeled using probability distributions [4, 6, 7, 10, 15, 25, 28–30, 35]. These approaches of asymptotic analysis are not limited to the fair division but have been widely studied in broader economic models; examples include one-to-one house allocation [19, 30], two-sided matchings [5, 23, 24], and voting theory [32, 39]. These studies provide insights into the behavior of mechanisms in large-scale scenarios, and offer both theoretical foundations and practical implications for real-world applications.

Our contributions. In this paper, to investigate the asymptotic existence of a matching that is efficient and fair among classes on average, we introduce a distributional model where each agent's utilities are drawn from a probability distribution.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA.

[1] In Benabbou et al. [9], the concept of class envy-freeness was referred to as typewise envy-freeness. We adopt the terminology used in Hosseini et al. [20] who considered class envy-freeness in the context of online matching.

We first prove that when m increases quadratically with respect to n, a welfare-maximizing mechanism, which returns a maximum-weight matching between agents and items, is asymptotically class envy-free. Importantly, this matching is also efficient since it maximizes the total utility of the agents. While this result provides asymptotic fairness and efficiency guarantees, it may not be satisfactory in all scenarios. Specifically, the mechanism may fail to treat classes fairly for worst-case inputs; for example, the welfare-maximizing mechanism could allocate all items to one class if the members of that class assign large values to each item.

To address this limitation, we integrate the *round-robin algorithm* from fair division into our framework. Specifically, in each round of the algorithm, each class selects an item with the highest marginal utility among the remaining items, creating a maximum-weight matching. Our main result establishes that this algorithm asymptotically yields a class envy-free matching as the number of agents approaches infinity, contingent upon milder assumptions regarding the number and sizes of classes. Furthermore, based on the recent result of Amanatidis et al. [3], the round-robin algorithm is known to produce a matching that satisfies an approximation of *class envy-freeness up to one item* (CEF1) for every input. We also prove that it achieves *non-wastefulness*, satisfying the property that no item can be reallocated to increase one class's valuation without decreasing another's.

While our algorithm shares similarities with the round-robin algorithm designed for fair division instances with additive agents, the non-additivity of valuations introduces notable distinctions in our analysis compared to the additive setting [30], which leads to unique technical challenges. Specifically, analyzing the behavior of the round-robin algorithm becomes complex because the set of items available at each round depends on the selections made in prior rounds.

In the case of additive valuations, Manurangsi and Suksompong [30] utilized a key property: envy between any pair of agents, based on the output of the round-robin algorithm, can be decomposed into the maximum value of an item obtained by another agent and the differences in value between the items each agent received and those received by another agent in consecutive rounds. Exploiting this fact, they showed that these differences can "catch up to" the maximum single-item value when there are sufficiently many items. However, achieving such a decomposition becomes challenging for assignment valuations due to the combinatorial structure of the matchings. In our setting, the marginal contribution of an item to different bundles may vary, and the domination property that an item chosen in an earlier round has a value greater than or equal to an item chosen later no longer holds. Consequently, unlike the additive case, the round-robin may not yield a CEF1 matching (see Example 1 in [20]).

To overcome these challenges, our proof critically leverages techniques from random assignment theory [2, 13, 18, 36, 37]. This theory considers a bipartite graph with random edge weights, primarily focusing on analyzing the expected value of a minimum weight perfect matching (which is essentially equivalent to a maximum-weight perfect matching).

Instead of examining individual pairs of items allocated to two classes, p and q, in each round, we focus on the marginal utility that class p receives in each round. This approach helps us analyze the expected value of the items allocated to class p. First, we derive a lower bound on the expected total utility for class p by applying the novel techniques introduced by Wästlund [37, 38] and Frieze and Johansson [18], which involve introducing a special vertex whose behavior explicitly determines the expected marginal gain. Next, we evaluate the expected value of the items allocated to class q from class p's perspective, using a randomly selected bundle of the same size as class p's. Finally, we discuss the concentration of probability around these expected values, showing that the edge weights chosen by the round-robin algorithm are sufficiently large. We discuss our proof techniques in Sections 4.

Related work. Our work is closely related to the growing literature on asymptotic fair division [4, 6, 7, 15, 25, 28–30, 35]. Dickerson et al. [15] initiated the study of asymptotic fair division. Although the non-existence of an envy-free allocation also holds in this setting, Dickerson et al. [15] demonstrated that a welfare-maximizing algorithm for additive agents produces an envy-free allocation with a probability that approaches 1 as m goes to infinity when $m = \Omega(n \log n)$. Following [15], Manurangsi and Suksompong [30] showed that under the assumption that utilities are drawn from a PDF-bounded distribution, $m = \Omega(n \log n / \log \log n)$, and agents have additive valuations, the round-robin algorithm returns an envy-free allocation with a probability that approaches 1 as $n \to \infty$. Apart from requiring fewer items for establishing asymptotic envy-freeness, the round-robin algorithm has another advantage over the welfare-maximizing algorithm; it achieves envy-freeness up to one item (EF1), for additive agents [14].

Bai and Gölz [7] extend these results to the case where agents have asymmetric distributions when distributions are PDF-bounded. Benadè et al. [10] demonstrated that the round-robin algorithm produces an SD envy-free allocation with a probability that approaches 1 as $m \to \infty$ when agents have order-consistent valuation functions, items are renamed by a uniformly random permutation, m is divisible by n, and $m = \omega(n^2)$.

Omitted proofs can be found in the full version of the paper [40].

2 MODEL

We use $[k]$ to denote the set $\{1, 2, \ldots, k\}$. Let $N = [n]$ be the set of n agents, and $I = [m]$ be the set of m items. The set of agents N is partitioned into k classes, labeled as $N_1, N_2, \ldots, N_k$. Let $n_p = |N_p|$ for each class p. We assume that $n_p \geq 1$ for every $p \in [k]$. Each N_p is referred to as class p. We call a subset of I a *bundle*.

We consider a matching problem where each item in I is matched to at most one agent in N, and each agent receives at most one item. Each agent $i \in N$ is endowed with a non-negative utility $u_i(j)$ for every item $j \in I$ where $u_i(j)$ ranges within the interval $[0, 1]$. We assume that $u_i(j)$ is drawn from a distribution over $[0, 1]$. Detailed assumptions on distributions are presented later in the section.

We define a complete bipartite graph $G = (N \cup I, E)$, where the set of agents in N forms the left vertices and the set of items in I forms the right vertices. Here, E denotes the set of edges. We consider the weights of edges where the weight of edge $\{i, j\} \in E$ is given by $u_i(j)$ for each $i \in N$ and $j \in I$. A *matching* M of a bipartite graph is defined as a set of edges wherein each vertex appears in at most one edge of M. For $S \subseteq N$, let $M(S)$ be the set of items which are assigned to some agent in S by matching M, i.e.,

$M(S) = \{j \in I \mid \exists\, i \in S : \{i, j\} \in M\}$. Each matching M induces an allocation that assigns the bundle $M(N_p)$ to every class p. For bundle $I' \subseteq I$, let $\mathcal{M}(N_p, I')$ denote the set of possible matchings between N_p and I' in G. We define the total utility obtained by class p under matching M as $u_p(M) = \sum_{\{i,j\} \in M, i \in N_p} u_i(j)$.

To define envy between classes, we introduce an *assignment valuation*, which determines how much hypothetical value each class can derive from a bundle allocated to another class.

Definition 2.1 (Assignment valuation). An *assignment valuation* $v_p(I')$ of class p for a bundle $I' \subseteq I$ is defined as the maximum total weight of a matching between the agents in N_p and the items in I'. Namely, $v_p(I')$ is given by $\max_{M \in \mathcal{M}(N_p, I')} \sum_{\{i,j\} \in M} u_i(j)$.

It is worth noting that the assignment valuation $v_p(I')$ for bundle $I' \subseteq I$ is upper bounded by the size n_p of each class $p \in [k]$ since $u_i(j) \le 1$ for each i, j. Here, each $v_p(I')$ can be computed in polynomial time by computing a maximum-weight matching in the given bipartite graph with edge weights; see Section 9 in [27].

Next, we introduce a concept of fairness among classes—*class envy-freeness*. This notion requires that the total utility each class receives must be greater than or equal to the maximum total utility that the class can derive from the items allocated to other classes.

Definition 2.2 (Class envy-freeness). For a matching M, we say that class p *envies* class q if $u_p(M) < v_p(M(N_q))$. A matching M is called *class envy-free* if no class envies another class, i.e., $u_p(M) \ge v_p(M(N_q))$ holds for every pair $p, q \in [k]$ of distinct classes.

If we allow each class to optimally reassign items within the members of the class, then the class would select a maximum-weight matching between the members of the class and their bundle. In such a scenario, the class envy-freeness requirement is equivalent to the above-mentioned definition, where the left-hand side $u_p(M)$ is replaced by $v_p(M(N_p))$.

As is observed in [9, 20], unfortunately, there exists an input where a class envy-free matching may not exist without allowing us to dispose items. Thus, the following approximation of class envy-freeness has been considered in [20]. A matching M is *α-class envy-free matching up to one item (CEF1)* if for every pair of classes $p, q \in [k]$, either class p does not envy class q, or there exists an item $j \in M(N_q)$ such that $\alpha^{-1} \cdot u_p(M) \ge v_p(M(N_q) \setminus \{j\})$. If $\alpha = 1$, we call such a matching *CEF1* [20].

Next, we define a measure of efficiency, called *non-wastefulness* [9]. Non-wastefulness requires that valuable items are not wasted.

Definition 2.3 (Non-wastefulness). For a matching M, an item $j \in I$ is said to be *wasted* if either

(a) item j is an unallocated and can increase the total utility of some class, i.e., $j \notin M(N)$ and $v_p(M(N_p) \cup \{j\}) - v_p(M(N_p)) > 0$ for some class p, or
(b) item j can be reallocated from class q to class p in a way that increases the total utility of class p without reducing the total utility of class q, i.e., there exist classes p, q such that $j \in M(N_q)$, $v_p(M(N_q)) - v_p(M(N_q) \setminus \{j\}) = 0$, and $v_p(M(N_p) \cup j) - v_p(M(N_p)) > 0$.

A matching is *non-wasteful* if no item is wasted.

As mentioned in Introduction, a class envy-free matching that satisfies non-wastefulness may not exist. For example, consider the case of a single item being allocated between two classes. We also remark that if we do not impose non-wastefulness, a CEF1 matching that allocates all items always exists and can be found in polynomial time using the envy-graph algorithm introduced by Lipton et al. [26]. However, the matching produced by the envy-graph algorithm may not satisfy non-wastefulness as pointed out by Benabbou et al. [9].

Distributions. For each agent $i \in N$ and item $j \in I$, the utility $u_i(j)$ is independently drawn from a given distribution $\mathcal{D}$ supported on $[0, 1]$. Let $f_{\mathcal{D}}$ and $F_{\mathcal{D}}$ denote the probability density function (PDF) and the cumulative distribution function (CDF) of $\mathcal{D}$, respectively. A distribution is said to be *non-atomic* if it does not assign a positive probability to any single point.

We say that a distribution $\mathcal{D}$ is (α, β)*-PDF-bounded* for constants $0 < \alpha \le \beta$ if it is non-atomic and $\alpha \le f_{\mathcal{D}}(x) \le \beta$ for all $x \in [0, 1]$. When $\alpha = \beta = 1$, $\mathcal{D}$ represents the uniform distribution over $[0, 1]$ since $f_{\mathcal{D}}(x) = 1$ for all $x \in [0, 1]$. The PDF-boundedness assumption is introduced by Manurangsi and Suksompong [30] as a natural class of distributions, which includes, for example, the uniform distribution and the truncated normal distribution.

Let $\mathrm{Exp}(\lambda)$ denote the exponential distribution with rate λ over $[0, \infty)$. Furthermore, let $\mathrm{ReExp}(\lambda)$ denote a distribution with the density function $f_\lambda(x) = \lambda \mathrm{e}^{-\lambda(1-x)}$ on the interval $(-\infty, 1]$. We call this probability distribution the *reversed exponential distribution*, which mirrors the exponential distribution $\mathrm{Exp}(\lambda)$ across the line $x = 1/2$. The cumulative distribution of $\mathrm{ReExp}(\lambda)$ is given by $F_{\mathrm{ReExp}(\lambda)}(x) = \mathrm{e}^{-\lambda(1-x)}$. We say that an event occurs *almost surely* if it occurs with probability 1.

Known results on maximum-weight matchings. We present several known results on maximum-weight matchings in a bipartite graph with random edge weights. Let H be a complete bipartite graph with bipartition (A, B). For $A' \subseteq A$ and $B' \subseteq B$, let $H[A', B']$ denote the subgraph of H induced by A' and B'. We first present the following lemma, which can be proven by a proof similar to that of the isolation lemma [22, 34]. Note that after edge weights on H have been sampled, we can select A' and B' while referring to the values of those edge weights.

Lemma 2.4. *Let H be a complete bipartite graph with bipartition (A, B) whose edge weights are drawn independently from a non-atomic distribution on $[0, 1]$. Let $A' \subseteq A$ and $B' \subseteq B$. Then, no pair of distinct matchings has the same total weight in $H[A', B']$ almost surely.*

Lemma 2.4 implies that, almost surely, a maximum-weight matching of a fixed size in $H[A', B']$ is uniquely determined. We next explain the nesting lemma, which follows from Lemma 3 in [13] or Lemma 2.1 in [38].

Lemma 2.5 (The nesting lemma). *Let H be a complete bipartite graph with bipartition (A, B) whose edge weights are drawn independently from a non-atomic distribution on $[0, 1]$. Let $A' \subseteq A$ and $B' \subseteq B$. Then, for each r with $1 < r \le \min(|A'|, |B'|)$, every vertex that appears as an element of the maximum-weight matching with $r - 1$ edges in $H[A', B']$ also appears as an element of the maximum-weight matching with r edges in $H[A', B']$ almost surely.*

We introduce notations and definitions related to matchings in a bipartite graph. For a bipartite graph H and each vertex i that appears in a matching M of H (namely $\{i, j\} \in M$ for some j), we denote by $M(i)$ the vertex matched to i under M. An *alternating* path P (resp. a cycle C) of matching M in bipartite graph H is a path (resp. a cycle) in H where, for every pair of consecutive edges on P, one of them is in M and the other one is not in M.

3 MAXIMUM-WEIGHT MATCHING

We introduce our first result, which states that if the number of items is quadratically large, then the probability that a maximum-weight matching is class envy-free approaches 1 as $m \to \infty$.

Theorem 3.1. *Suppose that $\mathcal{D}$ is non-atomic and there exists a constant $c > 0$ with $m/\log m \geq c \cdot (k^2 \max_{p\in[k]} n_p)^2$. Then, as $m \to \infty$, the probability that a maximum-weight matching is class envy-free approaches* 1.

To prove Theorem 3.1, we rely on the following simple observation: if each class desires a disjoint bundle, then the maximum-weight matching is class envy-free.

Lemma 3.2. *Suppose that there exist k disjoint bundles $I_1, I_2, \dots, I_k$ such that such that for every $p \in [k]$, I_p is a most favorite bundle of class p, i.e., $I_p \in \operatorname{argmax}_{I' \subseteq I, |I'|=n_p} v_p(I')$. Then, any maximum-weight matching in G is class envy-free and non-wasteful.*

Proof of Theorem 3.1. By Lemma 2.4, a maximum-weight matching in G of size n is uniquely determined almost surely. Since each edge weight is drawn from the same non-atomic distribution, we have that $\mathbf{Pr}\left[\operatorname{argmax}_{I''\subseteq I, |I''|=n_p} v_p(I'') = I'\right] = \frac{1}{\binom{m}{n_p}}$ for all $p \in [k]$ and $I' \subseteq I$ with $|I'| = n_p$.

Let $\mathcal{A}$ be an event that there exists a class envy-free and non-wasteful matching. Let $\mathcal{B}$ be an event that there exist k disjoint bundles $I_1, I_2, \dots, I_k$ such that $I_p \in \operatorname{argmax}_{I'\subseteq I, |I'|=n_p} v_p(I')$ for every $p \in [k]$. By Lemma 3.2, if the most favorite bundles of any two classes are disjoint, then there exists a class envy-free and non-wasteful matching. Then, we have $\mathbf{Pr}[\mathcal{A}] \geq \mathbf{Pr}[\mathcal{B}]$. Let P denote the set of partitions of the m items into disjoint bundles of sizes $n_1, n_2, \dots, n_k$. We provide a lower bound for $\mathbf{Pr}[\mathcal{B}]$ as follows.

$$
\begin{aligned}
&\mathbf{Pr}[\mathcal{B}] \\
&= \sum_{(I_1,I_2,\dots,I_k)\in P} \mathbf{Pr}\left[\operatorname*{argmax}_{I'\subseteq I,\, |I'|=n_p} v_p(I') = I_p \text{ for every } p \in [k]\right] \\
&= \sum_{(I_1,I_2,\dots,I_k)\in P} \frac{1}{\binom{m}{n_1}} \cdot \frac{1}{\binom{m}{n_2}} \cdot \dots \cdot \frac{1}{\binom{m}{n_k}} \\
&= \frac{\binom{m}{n_1}}{\binom{m}{n_1}} \cdot \frac{\binom{m-n_1}{n_2}}{\binom{m}{n_2}} \cdot \dots \cdot \frac{\binom{m-\sum_{i=1}^{k-1} n_i}{n_k}}{\binom{m}{n_k}} \\
&\geq \left(1 - \frac{n_2}{m - n_1 + 1}\right)^{n_1} \left(1 - \frac{n_3}{m - n_1 - n_2 + 1}\right)^{n_1+n_2} \\
&\quad \cdots \left(1 - \frac{n_k}{m - \sum_{i=1}^{k-1} n_i + 1}\right)^{\sum_{p=1}^{k-1} n_p} \\
&\geq \exp\left(-\frac{n_1 n_2}{m - n_1 + 1} - \frac{(n_1+n_2)n_3}{m - n_1 - n_2 + 1} - \dots - \frac{(\sum_{i=1}^{k-1} n_i) n_k}{m - \sum_{p=1}^{k-1} n_p + 1}\right) \\
&\geq \exp\left(-\frac{k^2(\max_{p\in[k]} n_p)^2}{m - \sum_{p=1}^{k-1} n_p + 1}\right) \\
&\geq \exp\left(-\frac{1/c \cdot m \log m}{m - \sqrt{1/c \cdot m \log m} + 1}\right).
\end{aligned}
$$

For the last inequality, we use $k^2(\max_{p\in[k]} n_p)^2 \leq 1/c \cdot m \log m$ and $\sum_{p=1}^{k-1} n_p \leq \sqrt{1/c \cdot m \log m}$. From this, we have $\mathbf{Pr}[\mathcal{A}] \to 1$ as $m \to \infty$. □

We note that the asymptotic existence of a class envy-free matching where every agent obtains exactly one item can be readily derived from the existing result on the one-to-one house allocation problem given by Manurangsi and Suksompong [30]. They showed that an envy-free assignment can be obtained by considering a greedy algorithm. This algorithm selects, in each step, an agent who has not yet been assigned an item. This agent then chooses their favorite item from those that have not been discarded (including items that have already been assigned to other agents). If the selected item was previously chosen by another agent in an earlier step, it is removed from further consideration. When there are sufficiently many items, this algorithm asymptotically produces a matching where each agent receives their most preferred item among those that were not discarded, resulting in a class envy-free matching.

As previously mentioned, both the maximum-weight matching and the matching produced by the greedy algorithm of [30] can be inherently unfair for worst-case inputs. This raises the question of whether there exists a matching mechanism that is fair for both worst-case and average-case inputs.

4 ROUND-ROBIN ALGORITHM

We next present our second result that a round-robin algorithm, presented as Algorithm 1, produces a class envy-free asymptotically. Moreover, based on recent results of [3, 31], we show that the resulting matching satisfies 1/2-CEF1. Also, we prove that it is non-wasteful.

Theorem 4.1. *Suppose that $\mathcal{D}$ is (α, β)-PDF-bounded, and the following three conditions* (a), (b) *and* (c) *hold.*

(a) *The number of items m is sufficiently large such that $m \geq k \cdot \max_{p\in[k]}(n_p + 2)$,*
(b) *the class sizes are almost proportional to the total population; more precisely, there exists a constant $C > 0$ such that $n \leq C \cdot (\min_{p\in[k]} n_p)^{5/4}$, and*
(c) *the number of classes k satisfies that $k > \max\left(\frac{1}{2\alpha}, \frac{\beta}{\alpha^2}\right)$ and $k^2 = O(n^{1/6})$.*

Then, as $n \to \infty$, the probability that Algorithm 1 produces a class envy-free and non-wasteful matching converges to 1.

Proposition 4.2. *The matching produced by the round-robin algorithm is 1/2-CEF1 and non-wasteful.*

Note that the first condition (a) of Theorem 4.1 is slightly stronger than the condition where $m \geq n$. The third condition (c) is very mild for some distributions, e.g. for uniform distributions, $k >$

Algorithm 1 The round-robin algorithm for classes with assignment valuations

Input: $N = N_1 \cup N_2 \cup \cdots \cup N_k$, I, $\{u_i(j)\}_{i \in N, j \in I}$
Output: Matching M

1: $M \leftarrow \emptyset, I_0 \leftarrow I, r \leftarrow 1$, and $M_p^0 \leftarrow \emptyset \; \forall p \in [k]$
2: **while** there is a remaining item which some class desires **do**
3:   **for** $p = 1, 2, \ldots, k$ **do**
4:     **if** there is a remaining item which class p desires **then**
5:       Let $j_p^r \in \operatorname{argmax}_{j \in I_0} v_p(M(N_p) \cup \{j\}) - v_p(M(N_p))$
6:       $M_p^r \leftarrow$ the maximum-weight matching between N_p and $M(N_p) \cup \{j_p^r\}$
7:       $M \leftarrow (M \setminus M_p^{r-1}) \cup M_p^r$
8:       $I_0 \leftarrow I_0 \setminus \{j_p^r\}$
9:     **end if**
10:   **end for**
11:   $r \leftarrow r + 1$
12: **end while**
13: **return** M.

$\max\left(\frac{1}{2\alpha}, \frac{\beta}{\alpha^2}\right)$ is equivalent to $k > 1$ since $\alpha = \beta = 1$. A perhaps more intuitive but stronger condition of (b) is the case where the total number of agents is within a constant factor of the minimum size of a class, i.e., $n = c \cdot \min_{p \in [k]} n_p$, where $c \geq 1$ is a constant. This is relevant in scenarios where the class sizes under consideration are proportional to the total population, such as gender groups, ethnic groups, and groups of people with the same political interests.

Algorithmic description. Let us now explain the round-robin algorithm (Algorithm 1). Each iteration of the **while** loop (Lines 2–12) in Algorithm 1 is referred to as a *round*. In each round r, each class selects its most preferred item, which has the highest marginal utility to the current bundle (Line 5) and updates its matching to create a new maximum-weight matching with an additional edge (Line 6). If a class encounters several items with the highest marginal utility, it selects one item arbitrarily among them. In Line 7, the matching between N and I is updated to reflect the new item acquired by the class p. Observe that in Algorithm 1, we allow each class to optimally reassign items within its members in each round, thereby selecting a maximum-weight matching between its members and the items allocated to them. Consequently, the total utility $u_p(M)$ that class p receives under Algorithm 1 is $v_p(M(N_p))$.

In contrast to the round-robin algorithm for additive valuations, Algorithm 1 may not produce a CEF1 matching; in fact, the factor of $1/2$ is the best that can be achieved by the round-robin algorithm; see Example 1 of [20]. Below, we illustrate the behavior of the algorithm.

Example 4.3. Consider an instance with two classes, each consisting of two agents, and four items in Table 1. The unique maximum-weight matching for this instance is given by $\{\{i_1, j_1\}, \{i_2, j_4\}, \{i_3, j_2\}, \{i_4, j_3\}\}$. However, this matching does not satisfy class envy-freeness as the second class receives a total utility of 3 despite having a maximum-weight matching of value 5 with the bundle allocated to the first class. In contrast, Algorithm 1 produces matching $\{\{i_1, j_1\}, \{i_2, j_2\}, \{i_3, j_4\}, \{i_4, j_3\}\}$, which can be easily checked to be both class envy-free and non-wasteful.

Table 1: Maximum-weight matching (left) and the matching produced by round-robin algorithm (right).

		j_1	j_2	j_3	j_4
N_1	i_1	5	0	0	0
	i_2	0	1	0	5
N_2	i_3	2	1	0	3
	i_4	1	0	2	0

		j_1	j_2	j_3	j_4
N_1	i_1	5	0	0	0
	i_2	0	1	0	5
N_2	i_3	2	1	0	3
	i_4	1	0	2	0

Outline of the proof. The remainder of this section is devoted to proving Theorem 4.1. Throughout this section, let M denote the matching produced by Algorithm 1 and M^r denote the matching at the end of round r in the algorithm. Furthermore, $M^r(N_p)$ is defined as the set of items matched to an agent in class p under matching M^r. Let I_p^r denote the set of remaining items just before class p selects an item in round r, i.e., $I_p^r = I \setminus \left(M^r(N_1) \cup \cdots \cup M^r(N_{p-1}) \cup M^{r-1}(N_p) \cup \cdots \cup M^{r-1}(N_k)\right)$. We fix two classes p and q and analyze the behavior of the following random variables:

$$X_p = v_p(M(N_p)) \quad \text{and} \quad X_{pq} = v_p(M(N_q)).$$

To prove Theorem 4.1, we first examine expected marginal weights of maximum-weight matchings (Lemma 4.4) in Section 4.1. This lemma claims that the difference in expected values can be expressed in terms of the probability that the special vertex belongs to a maximum matching in a modified graph. In the proof of Lemma 4.4, we adopt a technique pioneered by [37, 38]. These works introduced an additional vertex and connected it to every vertex on the other side by an edge with weights following an exponential distribution to analyze the expected marginal weight of minimum weight matchings; we adapt this technique to the context of maximum-weight matchings.

By utilizing Lemma 4.4, we establish bounds on the expected values of X_p and X_{pq} (Lemmas 4.5 and 4.6). In the proof, we analyze the probabilities that special vertices belong to maximum matchings in the graphs between N_p and $M(N_p)$, as well as between N_p and $M(N_q)$. Due to the uniqueness of the maximum-weight matching (Lemma 2.4), we consider the unique augmenting path updating it in each round.

In Section 4.4, by Lemmas 4.5 and 4.6, and by demonstrating that the edge weights in maximum-weight matchings between N_p and $M(N_q)$ are sufficiently "heavy" (the second bullet in Lemma 4.7), we prove that the difference in the expected values of X_p and X_{pq} is lower-bounded (Lemma 4.8). Finally, we achieve stochastic concentrations on the expectations of X_p and X_{pq}, establishing Theorem 4.1.

4.1 Expected Marginal Weight of a Maximum-Weight Matching

In this section, we show Lemma 4.4 by examining the difference in expected values arising from a maximum-weight matching of consecutive sizes. Consider a complete bipartite graph H with the left set A of vertices and the right set B of vertices. The weights of all edges in H are derived from non-atomic distributions over the interval $[0, 1]$, without the requirement for these weights to be drawn independently. We make the following assumptions:

Assumption 1. *No two distinct matchings have the same total weight almost surely in H.*

Assumption 2. *For every size $r = 1, 2, \dots, \min\{|A|, |B|\}$ and for any pair of vertices $i, i' \in A$, the probability that i belongs to the maximum-weight matching of size r in H is the same as that for i'.*

Now, we modify H by introducing a new vertex $\hat{j}$ to B, and create edges $\{i, \hat{j}\}$ for all $i \in A$ (see Figure 1). The weight of each edge $\{i, \hat{j}\}$ for $i \in A$ is independently drawn from $\mathrm{ReExp}(\lambda)$ on $(-\infty, 1]$ where $0 < \lambda \leq 1$. Let $\hat{H}$ denote the modified bipartite graph. Later, we will consider the limit probability of $\hat{j}$ being included in a maximum-weight matching when λ converges to 0. Let $\hat{B} = B \cup \{\hat{j}\}$. Note that edges with negative weights between A and $\{\hat{j}\}$ will not be included in any maximum-weight matching of size up to $\min\{|A|, |B|\}$ in $\hat{H}$ since the edges in H are non-negative. Let $\hat{B}^r$ represent the set of vertices in $\hat{B}$ under the maximum-weight matching of size r between A and $\hat{B}$. Note that no two distinct matchings have the same total weight almost surely in $\hat{H}$ by non-atomicity of $\mathrm{ReExp}(\lambda)$, by the assumption that the weights of edges incident to $\hat{j}$ are drawn independently, and by the assumption that no two distinct matchings in H have the same total weight.

Let X^r denote the maximum weight of a matching of size r in H. Lemma 4.4 states that the difference in expected values of X^r and X^{r-1} can be expressed in terms of the probability of $\hat{j}$ being included in $\hat{B}^r$.

Lemma 4.4. *Under Assumptions 1 and 2,*

$$\mathbb{E}[X^r] - \mathbb{E}[X^{r-1}] = 1 - \frac{1}{r} \lim_{\lambda \to 0} \frac{1}{\lambda} \mathbf{Pr}\left[\hat{j} \in \hat{B}^r\right], \tag{1}$$

for every size $r = 1, 2, \dots, \min(|A|, |B|)$.

Proof. Let $W(i, \hat{j})$ denote the random variable representing the weight of the edge $\{i, \hat{j}\}$ for each $i \in A$. Let A^r denote the set of vertices in A that are matched under the maximum-weight matching of size r in H. By The maximum-weight matching of size r between A^r and $\hat{B}$ is denoted by $\hat{M}$, and for each $i \in A^r$, let X_i^{r-1} be the maximum weight of matchings of size $r-1$ between $A^r \setminus \{i\}$ and B.

Select i from A^r uniformly at random. By definition, the maximum weight of a matching of size r between A^r and $\hat{B}$ under the constraint that the edge $\{i, \hat{j}\}$ is included is $X_i^{r-1} + W(i, \hat{j})$. We claim that $\mathbf{Pr}\left[X_i^{r-1} + W(i, \hat{j}) > X^r\right] = \mathbf{Pr}\left[\{i, \hat{j}\} \in \hat{M}\right] + O(\lambda^2)$. If edge $\{i, \hat{j}\}$ is included in $\hat{M}$, we have $X_i^{r-1} + W(i, \hat{j}) > X^r$ by Assumption 1. Thus, $\mathbf{Pr}[\{i, \hat{j}\} \in \hat{M}] \leq \mathbf{Pr}[X_i^{r-1} + W(i, \hat{j}) > X^r]$.

Next, suppose that $W(i, \hat{j}) > X^r - X_i^{r-1}$. Then, $\hat{M}$ must include an edge incident to $\hat{j}$. If no $i' \in A^r$ other than i satisfies the inequality $W(i', \hat{j}) > X^r - X_{i'}^{r-1}$, then we get $\{i, \hat{j}\} \in \hat{M}$. Let $\mathcal{F}_i$ denote the event that there exists no vertex $i' \neq i$ such that it satisfies $X_{i'}^{r-1} + W(i', \hat{j}) > X^r$. By the above argument, we have

$$\begin{aligned}
&\mathbf{Pr}\left[X_i^{r-1} + W(i, \hat{j}) > X^r\right] \\
&= \mathbf{Pr}\left[(X_i^{r-1} + W(i, \hat{j}) > X^r) \wedge \mathcal{F}_i\right] + \mathbf{Pr}\left[(X_i^{r-1} + W(i, \hat{j}) > X^r) \wedge \overline{\mathcal{F}_i}\right] \\
&\leq \mathbf{Pr}\left[\{i, \hat{j}\} \in \hat{M}\right] + \mathbf{Pr}\left[(X_i^{r-1} + W(i, \hat{j}) > X^r) \wedge \overline{\mathcal{F}_i}\right].
\end{aligned}$$

The right term of the above inequality can be bounded as follows:

$$\begin{aligned}
&\mathbf{Pr}\left[X_i^{r-1} + W(i, \hat{j}) > X^r \ \wedge\ \overline{\mathcal{F}_i}\right] \\
&\leq \sum_{i' \in A^r \setminus \{i\}} \mathbf{Pr}\left[X_i^{r-1} + W(i, \hat{j}) > X^r \ \wedge\ X_{i'}^{r-1} + W(i', \hat{j}) > X^r\right] \\
&= \sum_{i' \in A^r \setminus \{i\}} \mathbb{E}\left[(1 - \mathrm{e}^{-\lambda(1 - X^r + X_i^{r-1})})(1 - \mathrm{e}^{-\lambda(1 - X^r + X_{i'}^{r-1})})\right] \\
&= O(\lambda^2),
\end{aligned}$$

where we use the fact that $1 - \mathrm{e}^{-x} \leq x$ for any x for the last relation. Thus, summing over all $i \in A^r$, we get

$$\begin{aligned}
\mathbf{Pr}\left[\hat{j} \in \hat{B}^r\right] &= \sum_{i \in A^r} \mathbf{Pr}\left[\{i, \hat{j}\} \in \hat{M}\right] \\
&= \sum_{i \in A^r} \left(\mathbf{Pr}\left[X_i^{r-1} + W(i, \hat{j}) > X^r\right] - O(\lambda^2)\right) \\
&= \sum_{i \in A^r} \left(\mathbb{E}\left[1 - \mathrm{e}^{-\lambda(1 - X^r + X_i^{r-1})}\right] - O(\lambda^2)\right).
\end{aligned}$$

By Assumption 2, for all $i_1, i_2 \in A$, both $\mathbf{Pr}[i_1 \in A^r] = \mathbf{Pr}[i_2 \in A^r]$ and $\mathbf{Pr}[i_1 \notin A^{r-1}] = \mathbf{Pr}[i_2 \notin A^{r-1}]$ hold. This implies that $\mathbb{E}[X^{r-1}] = \mathbb{E}[X_i^{r-1}]$ since i is selected from A^r uniformly at random. Thus, $\mathbb{E}[X^{r-1}] = \frac{1}{r} \sum_{i' \in A^r} \mathbb{E}[X_{i'}^{r-1}]$. Thus, $\lim_{\lambda \to 0} \frac{1}{\lambda} \mathbf{Pr}\left[\hat{j} \in \hat{B}^r\right] = \sum_{i' \in A^r} \mathbb{E}\left[1 - X^r + X_{i'}^{r-1}\right] = r - r(\mathbb{E}[X^r] - \mathbb{E}[X^{r-1}])$. □

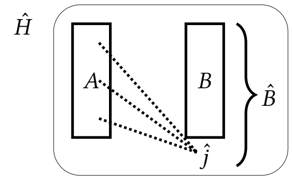

Figure 1: The modified graph $\hat{H}$ with the additional vertex $\hat{j}$.

4.2 Bounds on Expected Values of X_p and X_{pq}

Next, we establish a lower bound on the expected value of X_p, as well as an upper bound on the expected value of X_{pq}. By utilizing the linearity of expectation, we can decompose the expected value into the expected difference accumulated in each round. We then leverage Lemma 4.4 to analyze the difference between the expected values achieved by each class in two consecutive rounds.

A key observation is that the augmenting path updating the maximum-weight matching in each round can be uniquely determined due to Lemma 2.4. This uniqueness allows us to identify an edge incident to a newly added vertex in the path and condition on the weight of such an edge. Consequently, for the expected value of X_p, we obtain an upper bound on the limiting probability of the special vertex being included in a maximum-weight matching. Moreover, following a similar proof strategy, we derive an upper bound on the expected value of X_{pq}.

It is important to note that calculating the exact expectation of X_p and X_{pq} proves challenging. While exact computations of the expected minimum total weight of matchings have been explored in random assignment theory, these studies typically assume that the edge weight is drawn from the exponential distribution, whereas in our setting, the edge weight is drawn from the (α, β)-PDF-bounded distribution that lacks the memorylessness property. Nevertheless, we are able to establish lower and upper bounds on X_p and X_{pq} by carefully applying Lemma 4.4.

We now present a lower bound on the expected value of X_p.

Lemma 4.5. *Suppose that $\mathcal{D}$ is (α, β)-PDF-bounded and condition* (a). *Then, we have*

$$\mathbb{E}[X_p] \geq n_p - \frac{1}{\alpha} \sum_{r=1}^{n_p} \frac{1}{r} \sum_{r'=1}^{r} \frac{1}{m - r' \cdot k}. \tag{2}$$

We next turn our attention to the upper bound on the expected value of $X_{pq} = v_p(M(N_q))$. Since there are almost surely no edges with zero weight between N_p and $M(N_q)$, the size of the maximum-weight matching between N_p and $M(N_q)$ is $\min(n_p, n_q)$ almost surely. Also, by Lemma 2.4, the maximum-weight matching of a fixed size between N_p and $M(N_q)$ is uniquely determined. Now, for each $r = 1, 2, \dots, \min(n_p, n_q)$, let j_q^r denote the item which appears in the maximum-weight matching of size r but does not appear in that of size $r - 1$. Here, j_q^r is also uniquely determined from Lemma 2.5. In addition, let $W_q^{(r)}$ be the random variable representing the weight of the edge that is adjacent to j_q^r and included in the maximum-weight matching of size r between N_p and $M(N_q)$.

With these in hand, we present an upper bound on the expected value of X_{pq}.

Lemma 4.6. *Suppose that $\mathcal{D}$ is (α, β)-PDF-bounded and $m \geq n$. Then, we have*

$$\mathbb{E}\left[X_{pq}\right] \leq \min(n_p, n_q) - \frac{\alpha}{\beta} \sum_{r=1}^{\min(n_p,n_q)} \frac{1}{r} \sum_{r'=1}^{r} \frac{\mathbb{E}\left[W_q^{(r')}\right]}{n_q - r' + 2}.$$

A natural way to obtain an upper bound is to estimate the expected value $\mathbb{E}[X_{pq}]$ directly. However, it is difficult to do this since we do not know how class p evaluates the marginal gain class q enjoys at each round. Instead, we consider a bundle B_u of size n_q, which is selected uniformly at random from I. By conditioning on $M(N_q) = B_u$ and analyzing the maximum matching between N_p and B_u, we circumvent the difficulty and apply a similar argument to that for Lemma 4.5 to establish Lemma 4.6.

4.3 No Light Edges in Maximum-Weight Matchings

Using Lemmas 4.5 and 4.6, we can obtain a lower bound for $\mathbb{E}[X_p] - \mathbb{E}[X_{pq}]$. To achieve probabilistic concentrations around expected values, we demonstrate that the edge weights in the maximum-weight matching between N_p and $M(N_p)$ and those in the maximum-weight matching between N_p and $M(N_q)$ are sufficiently heavy. This is formalized in Lemma 4.7.

Lemma 4.7. *Suppose that distribution $\mathcal{D}$ is (α, β)-PDF-bounded and condition* (b) *of Theorem 4.1 holds. Then, we have the followings:*

- *Let c_p be a constant such that $c_p > 7000\alpha^{-1} > 0$. For every $r = 1, 2, \dots, n_p$, every edge in the maximum-weight matching of size r between N_p and $M(N_p)$ has a weight at least $1 - c_p \frac{(\log n_p)^2}{n_p}$ with a probability of at least $1 - O(n_p^{-3})$.*
- *Let c_q be a constant such that $c_q > 60c_p + 400\alpha^{-1} > 0$. For every $r = 1, 2, \dots, \min(n_p, n_q)$, every edge in the maximum-weight matching of size r between N_p and $M(N_q)$ has a weight at least $1 - c_q \frac{(\log \min(n_p,n_q))^4}{\min(n_p,n_q)}$ with a probability of at least $1 - O(\min(n_p, n_q)^{-3})$.*

To prove this, we employ an approach inspired by a technique for expanding bipartite graphs in the random assignment theory [17, 18, 36]. Although [17] and [18] consider the assignment problem with minimum cost, we use a similar argument to analyze maximum-weight matchings. Specifically, we analyze edge weights of a maximum-weight matching in a bipartite graph by considering an alternating cycle of a maximum-weight matching in the bipartite graph restricted to the "heavy" edges and bounding its length. Furthermore, to demonstrate that such an alternating cycle can be found, we show "expander" properties of sub-bipartite graphs of certain size.

4.4 Putting All the Pieces Together

Under conditions (a) and (b) of Theorem 4.1, we establish a lower bound on the difference in expected values that class p assigns to the bundles allocated to classes p and q under the matching produced by the round-robin algorithm.

Lemma 4.8. *Suppose that $\mathcal{D}$ is (α, β)-PDF-bounded and two conditions* (a) *and* (b) *of Theorem 4.1 hold. Then,*

$$\mathbb{E}[v_p(M(N_p))] - \mathbb{E}[v_p(M(N_q))] \geq \left(1 - \frac{1}{2\alpha k}\right)\left(n_p - \min(n_p, n_q)\right)$$
$$+ \left(\frac{\alpha}{\beta} - \frac{n_q + 1}{\alpha(m - k)}\right) \frac{\min(n_p, n_q)}{n_q} - O(\min(n_p, n_q)^{-1}).$$

We explain implications of Lemma 4.8. If we have that $k > \max(\frac{1}{2\alpha}, \frac{\beta}{\alpha^2})$ and $m \geq k \max_p(n_p + 2)$, then we have $1 - \frac{1}{2\alpha k} > 0$ and $\frac{\alpha}{\beta} - \frac{n_q+1}{\alpha(m-k)} \geq \frac{\alpha}{\beta} - \frac{1}{\alpha k} > 0$. Therefore, when $n_p \geq n_q$, there exists a positive constant $c > 0$ such that the expected difference is at least $c - O(n_q^{-1})$. When $n_p < n_q$, the expected difference is at least $\left(\frac{\alpha}{\beta} - \frac{n_q+1}{\alpha(m-k)}\right) \frac{n_p}{n_q} - O(n_p^{-1})$, where the lower bound is mainly determined by the ratio of n_p to n_q.

Proof. By the second bullet in Lemma 4.7, for every $r = 1, 2, \dots, \min(n_p, n_q)$, there is a constant $c > 0$ such that

$$\mathbf{Pr}\left[W_q^{(r)} \geq 1 - c\frac{(\log \min(n_p, n_q))^4}{\min(n_p, n_q)}\right] = 1 - O(\min(n_p, n_q)^{-3}).$$

For every $r = 1, 2, \dots, \min(n_p, n_q)$, since $W_q^{(r)} \geq 0$, we obtain $\mathbb{E}[W_q^{(r)}] = 1 - O(\min(n_p, n_q)^{-1})$. Thus, combined with Lemma 4.6, we get

$$\mathbb{E}[X_{pq}] \leq \min(n_p, n_q) - \frac{\alpha}{\beta} \sum_{r=1}^{\min(n_p,n_q)} \frac{1}{r} \sum_{r'=1}^{r} \frac{1}{n_q - r' + 2}$$
$$+ \frac{\alpha}{\beta} \sum_{r=1}^{\min(n_p,n_q)} \frac{1}{r} \sum_{r'=1}^{r} \frac{1}{n_q - r' + 2} O(\min(n_p, n_q)^{-1}).$$

Here, we can show $\sum_{r=1}^{\min(n_p,n_q)} \frac{1}{r} \sum_{r'=1}^{r} \frac{1}{n_q - r' + 2} = O(1)$. Combining this and Lemma 4.5 implies that

$$\mathbb{E}[X_p] - \mathbb{E}[X_{pq}] \geq n_p - \min(n_p, n_q) - O(\min(n_p, n_q)^{-1})$$
$$+ \left(\frac{\alpha}{\beta} - \frac{n_q + 1}{\alpha(m - k)}\right) \sum_{r=1}^{\min(n_p,n_q)} \frac{1}{r} \sum_{r'=1}^{r} \frac{1}{n_q - r' + 2}$$

$$-\frac{1}{\alpha}\sum_{r=\min(n_p,n_q)+1}^{n_p}\frac{1}{r}\sum_{r'=1}^{r}\frac{1}{m-r'\cdot k}$$
$$\geq n_p-\min(n_p,n_q)-O(\min(n_p,n_q)^{-1})$$
$$+\left(\frac{\alpha}{\beta}-\frac{n_q+1}{\alpha(m-k)}\right)\sum_{r=1}^{\min(n_p,n_q)}\frac{1}{r}\sum_{r'=1}^{r}\frac{1}{n_q}-\frac{1}{\alpha k}\sum_{r=\min(n_p,n_q)+1}^{n_p}\frac{1}{r}\sum_{r'=1}^{r}\frac{1}{2}$$
$$=\left(1-\frac{1}{2\alpha k}\right)(n_p-\min(n_p,n_q))+\left(\frac{\alpha}{\beta}-\frac{n_q+1}{\alpha(m-k)}\right)\frac{\min(n_p,n_q)}{n_q}$$
$$-O(\min(n_p,n_q)^{-1}),$$

where we use $m \geq k(n_q+2)$, $\frac{1}{n_q-r'+2} \geq \frac{1}{n_q}$ and $\frac{1}{2} \geq \frac{1}{m-r'\cdot k}$ for the second inequality. □

Finally, we prove Theorem 4.1. We denote the right hand-side of the inequality in Lemma 4.8 by $D(n_p, n_q)$. We first introduce the Efron-Stein inequality and Chebyshev's inequality in the following.

Lemma 4.9 (Efron–Stein inequality [12, 16]). *Suppose that n random variables $X_1, X_2, \ldots, X_n$ are independent. Let $f : \mathbb{R}^n \to \mathbb{R}$ be an arbitrary measurable function of n random variables. Let $X = (X_1, X_2, \ldots, X_n)$ and $X^{(i)} = (X_1, X_2, \ldots, X_{i-1}, 0, X_{i+1}, \ldots, X_n)$. Then we have* $\mathrm{Var}[f(X_1, X_2, \ldots, X_n)] \leq \sum_{i=1}^{n} \mathbb{E}[(f(X)-f(X^{(i)}))_+^2]$. *Here,* $(x)_+ = \max(x, 0)$.

Lemma 4.10 (Chebyshev's inequality). *If X is any random variable, then for any $\varepsilon > 0$ we have* $\mathbf{Pr}[\,|X - \mathbb{E}[X]| \geq \varepsilon\,] \leq \frac{\mathrm{Var}[X]}{\varepsilon^2}$.

Proof sketch of Theorem 4.1. First, we investigate the probability that the value of the random variable X_p deviates from its expected value. Let $W(i,j) = u_i(j)$ denote the weight of edge $\{i,j\}$ for $i \in N$ and $j \in I$. Let $X_p(W)$ denote the total utility that class p obtains from the matching produced by the round-robin algorithm when the input is $(W(i,j))_{i,j}$. We now consider another weight function. Let $\delta_p = c_p\frac{(\log n_p)^2}{n_p}$, and let $\overline{W}(i,j) = \max(W(i,j), 1-\delta_p)$ for $i \in N$ and $j \in I$. We denote by $X_p(\overline{W})$ the utility attained by class p under the round-robin algorithm when weights of edges are $\{\overline{W}(i,j)\}_{i,j}$. By some calculations, we have $\mathrm{Var}[X_p(W)] \leq 2\mathrm{Var}[X_p(\overline{W})] + 2n_p^2 \cdot \mathbf{Pr}[X_p(W) \neq X_p(\overline{W})]$. Let $\overline{W}^{(i,j)}$ denote the weights obtained by setting the (i,j)-th entry of $\overline{W}$ to zero while keeping all other entries unchanged. By applying the Efron–Stein inequality, we obtain $\mathrm{Var}[X_p(\overline{W})] \leq \sum_{i=1}^{n}\sum_{j=1}^{m}\mathbb{E}[(X_p(\overline{W})-X_p(\overline{W}^{(i,j)}))_+^2]$. By the definition of $\overline{W}$, we have $(X_p(\overline{W})-X_p(\overline{W}^{(i,j)}))_+^2 \leq 2\delta_p^2$ for all $i \in N$ and $j \in I$. We claim that the number of edges involved in the matching produced by the algorithm is, with a probability at least $1-O(\min_{p\in[k]} n_p^{-3})$, at most $\sum_{p=1}^{k} O(\log n_p)n_p$. Thus, $\mathrm{Var}[X_p(\overline{W})] \leq \delta_p^2 \sum_{p=1}^{k} O(\log n_p)n_p + O(n_p^{-1}) \leq \delta_p^2 nO(\log n) + O(n_p^{-1})$. Moreover, from the first bullet in Lemma 4.7, we have $2n_p^2 \cdot \mathbf{Pr}[X_p(W) \neq X_p(\overline{W})] = 2n_p^2 \cdot O(n_p^{-3}) = O(n_p^{-1})$. Hence, we get $\mathrm{Var}[X_p] = \mathrm{Var}[X_p(W)] \leq 2\delta_p^2 nO(\log n) + O(n_p^{-1})$. From the Chebyshev's inequality, we get

$$\mathbf{Pr}\left[X_p - \mathbb{E}[X_p] < -\frac{1}{2}D(n_p,n_q)\right]$$
$$\leq \frac{4\mathrm{Var}\left[X_p\right]}{D(n_p,n_q)^2} \leq \frac{8\delta_p^2 nO(\log n) + O(n_p^{-1})}{D(n_p,n_q)^2}.$$

Subsequently, we investigate the concentration around the expected value of the random variable X_{pq}. Similarly to X_p, letting $\delta_{pq} = c_q\frac{(\log\min(n_p,n_q))^4}{\min(n_p,n_q)}$, we obtain

$$\mathbf{Pr}\left[X_{pq} - \mathbb{E}[X_{pq}] > \frac{1}{2}D(n_p,n_q)\right]$$
$$\leq \frac{4\mathrm{Var}\left[X_{pq}\right]}{D(n_p,n_q)^2} \leq \frac{8\delta_{pq}^2 nO(\log n) + O(\min(n_p,n_q)^{-1})}{D(n_p,n_q)^2}.$$

From Lemma 4.8, if class p envies class q, then $X_p < \mathbb{E}[X_p] - \frac{1}{2}D(n_p,n_q)$ or $X_{pq} > \mathbb{E}[X_{pq}] + \frac{1}{2}D(n_p,n_q)$ must hold. Thus, we obtain

$$\mathbf{Pr}\left[\text{Class } p \text{ envies class } q\right] \leq \mathbf{Pr}\left[X_p < \mathbb{E}[X_p] - \frac{1}{2}D(n_p,n_q)\right]$$
$$+\,\mathbf{Pr}\left[X_{pq} > \mathbb{E}[X_{pq}] + \frac{1}{2}D(n_p,n_q)\right]$$
$$\leq \left(8(\delta_p^2+\delta_{pq}^2)nO(\log n) + O(\min(n_p,n_q)^{-1})\right)D(n_p,n_q)^{-2}$$
$$\leq \Big(16c_q \cdot \min(n_p,n_q)^{-2}nO(\log^9 n)$$
$$+O(\min(n_p,n_q)^{-1})\Big)D(n_p,n_q)^{-2},$$

where we use $\delta_p^2 + \delta_{pq}^2 \leq 2c_q \min(n_p,n_q)^{-2}\log^8 n$ for the last inequality.

By condition (b) (i.e., $\max(n_p,n_q) \leq n \leq C\cdot\min(n_p,n_q)^{5/4}$) and Lemma 4.8, we can show that the probability that class p envies class q is at most $\tilde{O}(n^{-1/5})$ for every pair of $p, q \in [k]$, where $\tilde{O}$ is the big-O notation that ignores logarithmic factors. By condition (c), the probability that matching M is not class envy-free is at most $k^2 \cdot \tilde{O}(n^{-1/5}) = o(1)$. □

5 CONCLUDING REMARKS

This paper addressed the problem of achieving fairness in matching across different classes and suggests several open problems for future research. In Theorem 4.1, we made several assumptions. In particular, while Dickerson et al. [15] and Manurangsi and Suksompong [30], as well as our Theorem 3.1, provide asymptotic results when the number of items approaches infinity, Theorem 4.1 assumes the number of agents approaches infinity. This assumption is necessary for our analysis because the assignment valuation of each class is bounded by the number of agents in that class. However, it remains unclear whether the round-robin algorithm produces asymptotically a class envy-free matching when these assumptions do not hold. We leave this as an interesting open question for future work. Moreover, it would be interesting to explore the case of asymmetric agents, where each agent has a different probability distribution for their utilities.

ACKNOWLEDGMENTS

This work was partially supported by JST FOREST Grant Numbers JPMJPR20C1 and by JST ERATO under grant number JPMJER2301. We thank the anonymous AAMAS 2025 for their valuable comments.

REFERENCES

[1] Atila Abdulkadiroğlu and Tayfun Sönmez. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66, 3 (1998), 689–701.

[2] David Aldous. 2001. The $\zeta(2)$ limit in the random assignment problem. *Random Structures and Algorithms* 18 (2001), 381–418.

[3] Georgios Amanatidis, Georgios Birmpas, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. 2023. Round-Robin Beyond Additive Agents: Existence and Fairness of Approximate Equilibria. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*. 67–87.

[4] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi. 2017. Approximation algorithms for computing maximin share allocations. *ACM Transactions on Algorithms* 13, 4 (2017), 1–28.

[5] Itai Ashlagi, Mark Braverman, and Avinatan Hassidim. 2014. Stability in large matching markets with complementarities. *Operations Research* 62, 4 (2014), 713–732.

[6] Yushi Bai, Uriel Feige, Paul Gölz, and Ariel D. Procaccia. 2022. Fair Allocations for Smoothed Utilities. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*. 436–465.

[7] Yushi Bai and Paul Gölz. 2022. Envy-free and Pareto-optimal allocations for agents with asymmetric random valuations. In *Proceedings of the 31th International Joint Conference on Artificial Intelligence (IJCAI)*. 53–59.

[8] Siddhartha Banerjee and Ramesh Johari. 2019. Ride sharing. In *Sharing economy*. Springer, 73–97.

[9] Nawal Benabbou, Mithun Chakraborty, Edith Elkind, and Yair Zick. 2019. Fairness towards groups of agents in the allocation of indivisible items. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. 95–101.

[10] Gerdus Benadè, Daniel Halpern, Alexandros Psomas, and Paritosh Verma. 2024. On the existence of envy-free allocations beyond additive valuations. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC'24)*.

[11] Anna Bogomolnaia and Hervé Moulin. 2001. A new solution to the random assignment problem. *Journal of Economic Theory* 100, 2 (2001), 295–328.

[12] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. 2013. *Concentration Inequalities: A Nonasymptotic Theory of Independence.* Oxford University Press.

[13] Marshall W. Buck, Clara S. Chan, and David P. Robbins. 2002. On the expected value of the minimum assignment. *Random Structures and Algorithms* 21, 1 (2002), 33–58.

[14] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation* 7, 3 (2019), 1–12.

[15] John P. Dickerson, Jonathan Goldman, Jeremy Karp, Ariel D. Procaccia, and Tuomas Sandholm. 2014. The computational rise and fall of fairness. In *Proceedings of the 28nd AAAI Conference on Artificial Intelligence (AAAI)*. 1405–1411.

[16] Bradley Efron and Charles Stein. 1981. The jackknife estimate of variance. *The Annals of Statistics* 9, 3 (1981), 586 – 596.

[17] Alan M. Frieze. 2021. The effect of adding randomly weighted edges. *SIAM Journal on Discrete Mathematics* 35, 2 (2021), 1182–1200.

[18] Alan M. Frieze and Tony Johansson. 2017. Minimum-cost matching in a random graph with random costs. *SIAM Journal on Discrete Mathematics* 31, 1 (2017), 489–510.

[19] Jiarui Gan, Warut Suksompong, and Alexandros A. Voudouris. 2019. Envy-freeness in house allocation problems. *Mathematical Social Sciences* 101 (2019), 104–106.

[20] Hadi Hosseini, Zhiyi Huang, Ayumi Igarashi, and Nisarg Shah. 2024. Class fairness in online matching. *Artificial Intelligence* 335 (2024), 104177.

[21] Aanund Hylland and Richard Zeckhauser. 1979. The efficient allocation of individuals to positions. *Journal of Political Economy* 87, 2 (1979), 293–314.

[22] Stasys Jukna. 2011. *Extremal combinatorics: with applications in computer science* (2nd ed.). Springer Publishing Company, Incorporated.

[23] Fuhito Kojima and Parag A. Pathak. 2009. Incentives and stability in large two-sided matching markets. *American Economic Review* 99, 3 (2009), 608–627.

[24] Fuhito Kojima, Parag A. Pathak, and Alvin E. Roth. 2013. Matching with couples: Stability and incentives in large markets. *The Quarterly Journal of Economics* 128, 4 (2013), 1585–1632.

[25] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. 2016. When can the maximin share guarantee be guaranteed?. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*. 523–529.

[26] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*. 125–131.

[27] László Lovász and Michael D. Plummer. 2009. *Matching Theory.* American Mathematical Society.

[28] Pasin Manurangsi and Warut Suksompong. 2017. Asymptotic existence of fair divisions for groups. *Mathematical Social Sciences* 89 (2017), 100–108 (2017), 100–108.

[29] Pasin Manurangsi and Warut Suksompong. 2020. When do envy-free allocations exist? *SIAM Journal on Discrete Mathematics* 34, 3 (2020), 1505–1521.

[30] Pasin Manurangsi and Warut Suksompong. 2021. Closing gaps in asymptotic fair division. *SIAM Journal on Discrete Mathematics* 35, 2 (2021), 668–706.

[31] Luisa Montanari, Ulrike Schmidt-Kraepelin, Warut Suksompong, and Nicholas Teh. 2024. Weighted envy-freeness for submodular valuations. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 38. 9865–9873.

[32] Arkadii Slinko. 2002. On asymptotic strategy-proofness of the plurality and the run-off rules. *Social Choice and Welfare* 19, 2 (2002), 313–324.

[33] Tayfun Sönmez and M. Utku Ünver. 2017. Market design for living-donor organ exchanges: an economic policy perspective. *Oxford Review of Economic Policy* 33, 4 (2017), pp. 676–704.

[34] Joel Spencer. 1996. Probabilistic methods. In *Handbook of Combinatorics (Vol. 2)*. MIT Press, Chapter 33, 1785–1817.

[35] Warut Suksompong. 2016. Asymptotic existence of proportionally fair allocations. *Mathematical Social Sciences* 81 (2016), 62–65.

[36] Michel Talagrand. 1995. Concentration of measure and isoperimetric inequalities in product spaces. *Publications Mathmatiques de L'IHÉS* (1995).

[37] Johan Wästlund. 2005. A simple proof of the Parisi and Coppersmith-Sorkin formulas for the random assignment problem. *Linköping Studies in Mathematics* 6 (2005).

[38] Johan Wästlund. 2009. An easy proof of the $\zeta(2)$ limit in the random assignment problem. *Electronic Communications in Probability* 14 (2009), 261–269.

[39] Lirong Xia. 2020. The smoothed possibility of social choice. In *Proceedings of the 33rd Annual Conference on Neural Information Processing Systems (NeurIPS)*. 11044–11055.

[40] Tomohiko Yokoyama and Ayumi Igarashi. 2025. Asymptotic existence of class envy-free matchings. *arXiv preprint arXiv:2502.14371* (2025).

[41] Lin Zhou. 1990. On a conjecture by gale about one-sided matching problems. *Journal of Economic Theory* 52, 1 (1990), 123–135.