

Class Fairness in Online Matching

Hadi Hosseini¹, Zhiyi Huang², Ayumi Igarashi³, Nisarg Shah⁴

¹ Pennsylvania State University

² University of Hong Kong

³ University of Tokyo

⁴ University of Toronto

hadi@psu.edu, zhiyi@cs.hku.hk, igarashi@mist.i.u-tokyo.ac.jp, nisarg@cs.toronto.edu

Abstract

We initiate the study of fairness among *classes* of agents in online bipartite matching where there is a given set of offline vertices (aka agents) and another set of vertices (aka items) that arrive online and must be matched irrevocably upon arrival. In this setting, agents are partitioned into classes and the matching is required to be fair with respect to the classes. We adopt popular fairness notions (e.g. envy-freeness, proportionality, and maximin share) and their relaxations to this setting and study deterministic algorithms for matching indivisible items (leading to integral matchings) and for matching divisible items (leading to fractional matchings). For matching indivisible items, we propose an adaptive-priority-based algorithm, MATCH-AND-SHIFT, prove that it achieves $1/2$ -approximation of both class envy-freeness up to one item and class maximin share fairness, and show that each guarantee is tight. For matching divisible items, we design a water-filling-based algorithm, EQUAL-FILLING, that achieves $(1 - 1/e)$ -approximation of class envy-freeness and class proportionality; we prove $1 - 1/e$ to be tight for class proportionality and establish a $3/4$ upper bound on class envy-freeness.

1 Introduction

The one-sided matching problem is a fundamental subject within economics and computation that deals with the matching of a set of items to a set of agents. Its primary objective is to ensure desirable normative properties such as *economic efficiency* and *fairness*. The advent of Internet economics along with the introduction of novel marketplaces has posed new challenges in designing desirable solutions for which, as noted by Moulin (2019), “*we need division rules that are both transparent and agreeable, in other words, fair.*” A wide array of these applications are inherently online, that is, items (or goods) arrive in an online fashion, and need to be matched immediately and irrevocably to the participating agents: consider the examples of allocating advertisement slots to Internet advertisers (Mehta et al. 2007), assigning packets to output ports in switch routing (Azar and Richter 2005), distributing food donations among nonprofit charitable organizations (Lee et al. 2019), and matching riders to drivers in ridesharing platforms (Banerjee and Johari 2019).

Over the past few decades, a large body of literature—within the field of online algorithm design—is devoted to the study of *online bipartite matching* problems. Their primary goal is to satisfy some notion of economic efficiency—e.g. maximizing the size of the final matching—with no knowledge of which items will arrive in the future and in what order. Algorithms designed for this problem are judged by their *competitive ratio*, which is the worst-case approximation ratio of the size of the matching produced to the maximum possible size in hindsight. It is well known that the best deterministic algorithm can only achieve a $1/2$ -approximation of this efficiency goal, e.g., by using a greedy algorithm to get a maximal matching. Notably, the seminal work of Karp, Vazirani, and Vazirani (1990) provides a randomized algorithm called RANKING with the best possible $(1 - 1/e)$ -approximation.

While the literature offers online algorithms with optimal efficiency guarantees, little work has been done in ensuring that these algorithms treat agents, or rather, classes of agents *fairly*. Consider the example of a food bank that wishes to distribute the donated items among nonprofit organizations and homeless shelters. The perishable food items donated to the food bank must be assigned upon their arrival. How should an online matching algorithm distribute these donations to the nonprofits and shelters in such a manner that the communities they serve are treated equitably?

Class fairness. We initiate the study of class fairness in online matching, where a set of items arriving online must be assigned to agents who are partitioned into known classes, with the goal of achieving fairness among classes. Agents either like an item (value 1) or don’t like it (value 0). We adopt classical notions from the fair division literature that typically apply to individual agents—such as *envy-freeness* (EF), *proportionality* (Prop), and *maximin share guarantee* (MMS)—to classes of agents. Our extensions ensure that different classes are treated fairly, regardless of their sizes (e.g., in the food bank example above, different communities are treated equally, even if some have many more organizations serving them).

In the standard fair division model, the impossibility of achieving envy-freeness has motivated relaxations such as envy-freeness up to one item (EF1), which can be guaranteed (Lipton et al. 2004). When applied to classes, our *class envy-freeness up to one item* (CEF1) requires that envy of




Figure 1: An adversarial instance where CEF1 cannot be achieved together with non-wastefulness.

any class towards another class to be eliminated after the removal of at most one item that is matched to an agent within the envied class. When all items are available up front, it is known that CEF1 can be achieved without unnecessarily throwing away items (Benabbou et al. 2020).¹ Can it still be achieved in the online setting?

Impossibility of CEF1 in online matching. First, note that “class-awareness” is necessary for any algorithm; otherwise an algorithm that is blind to the class information may violate CEF1 by assigning two arriving items to the same class, when there is another class that likes both items. Unfortunately, a slightly larger example shows that even class-aware online algorithms cannot always achieve CEF1.

Example 1. Consider the example in Figure 1, in which six agents are partitioned into two classes $N_1 = \{a_1, a_2, a_3\}$ and $N_2 = \{b_1, b_2, b_3\}$. Four items arrive sequentially in the order (o_1, o_2, o_3, o_4) . An edge indicates that an agent likes an item; thick edges indicate the matching. Assume that we do not wish to throw away any item as long as there is an unmatched agent who likes it. For $i \in \{1, 2, 3\}$, item o_i is liked by agents a_i and b_i . The first item o_1 can be matched to either a_1 or b_1 ; without loss of generality, suppose it is matched to $a_1 \in N_1$. When o_2 arrives, it must be matched to $b_2 \in N_2$ in order to satisfy CEF1. The third item o_3 can again be matched to either of a_3 and b_3 ; without loss of generality, suppose it is matched to $b_3 \in N_2$. Now, o_4 arrives, and it is liked only by a_1 (who is already matched) and b_1 (who is unmatched). The algorithm must assign it to b_1 due to non-wastefulness, which leaves class N_1 envious of class N_2 , even if we ignore any one of the items assigned to N_2 .

Given this impossibility, we seek online matching algorithms that achieve the fairness notions *approximately*, often in conjunction with approximate efficiency guarantees. We aim to answer the following theoretical questions:

Can we design deterministic algorithms for matching indivisible or divisible items that achieve approximate class fairness while adhering to efficiency requirements? And, can we surpass their guarantees by using randomization?

Our Results Our first contribution (Section 2) is developing a detailed mathematical framework in which we adopt classical fairness concepts to online matching. We consider two

¹We later formalize the latter restriction as *non-wastefulness* (NW). This is required because CEF1, on its own, can be achieved vacuously via an empty matching by throwing away all the items.

types of online matching models, one with indivisible items, wherein an item must be matched in its entirety to a single agent, and one with divisible items, wherein an item may be fractionally divided between multiple agents.

For both settings, we design online algorithms that achieve approximate fairness and efficiency guarantees, and also provide upper bounds on the approximations that can be achieved by any online algorithm. Our algorithms satisfy non-wastefulness, which implies $1/2$ -approximation of the optimal utilitarian social welfare (USW); the utilitarian social welfare, i.e., the sum of agent utilities, is effectively the size of the matching. Specifically, we make the following contributions (summarized in Table 1):

- **Indivisible matching:** For indivisible items, we develop a deterministic algorithm, MATCH-AND-SHIFT, that simultaneously achieves non-wastefulness, $1/2$ -CEF1, $1/2$ -CMMS, and $1/2$ -USW (Theorem 1). The algorithm uses an adaptive priority queue over classes, in which a class is shifted to the end of the queue immediately upon receiving an item. Further, we prove that no deterministic algorithm can achieve any of α -CEF1 (subject to non-wastefulness), α -CMMS, or α -USW, for any $\alpha > 1/2$ (Theorem 2), establishing our algorithm to be simultaneously optimal for each guarantee.
- **Divisible matching:** For divisible items, we improve the above bounds via a different algorithm, EQUAL-FILLING. This algorithm divides items equally between the classes, but uses water-filling to divide the portion of an item assigned to a class between the agents in that class. This algorithm simultaneously achieves non-wastefulness, $(1 - 1/e)$ -CEF, $(1 - 1/e)$ -CPROP, and $1/2$ -USW (Theorem 3). Further, no deterministic algorithm can achieve α -CEF for any $\alpha > 3/4$, or α -USW for any $\alpha > 1 - 1/e$, and $(1 - 1/e)$ -CPROP is tight (Theorem 4).

Related Work We refer readers to Mehta (2013) for a survey of the vast literature on online matching, and summarize some results that are the most related to this paper. Our problem can be seen as a fair division problem by considering each class to be a meta-agent; the value of this meta-agent for a bundle of items is the maximum total value obtained by matching the items to the agents in the class, which induces OXS valuations (Paes Leme 2017) (these are not additive). Benabbou et al. (2019) studied a model similar to ours in the offline setting, and showed that an allocation satisfying EF1 and non-wastefulness exists and can be computed in polynomial time. Subsequent papers (Benabbou et al. 2020; Babaioff, Ezra, and Feige 2021; Barman and Verma 2021) considered a more general class of submodular valuations with dichotomous marginals and proved that EF1 and optimal USW can be achieved together; Barman and Verma (2021) proved a similar result for MMS and optimal USW. Our paper is also related to the growing line of work on *online* fair division (Benade et al. 2018; Banerjee et al. 2022; Zeng and Psomas 2020; Walsh 2011; Aleksandrov et al. 2015), but a majority of this work focuses on additive valuations, and hence, their techniques do not apply to our matching setting. In the full version (Hosseini et al. 2022), we provide an extended review on the related literature.

Indivisible			Divisible		
Fairness	Algorithm	Upper Bound	Fairness	Algorithm	Upper Bound
α -CEF1 + NW	$1/2$	$1/2$	α -CEF + NW	$1 - \frac{1}{e}$	$3/4$
α -CMMS	$1/2$	$1/2$	α -CPROP	$1 - \frac{1}{e}$	$1 - \frac{1}{e}$
α -USW	$1/2$	$1/2$	α -USW	$1/2$	$1 - \frac{1}{e}$

Table 1: The summary of our results on deterministic algorithms for matching indivisible and divisible items. Each algorithm achieves its three guarantees simultaneously, while the upper bound holds for any algorithm, separately for each guarantee.

2 Model

For $t \in \mathbb{N}$, define $[t] = \{1, \dots, t\}$. First, let us introduce an offline version of our model and the solution concepts we seek. Later, we will discuss the online model and algorithms in that model.

Consider a bipartite graph $G = (N, M, E)$, where N represents a set of vertices called agents, M a set of vertices called items, and E the set of edges. We say that agent a likes item o if a is adjacent to o , i.e., $(a, o) \in E$. The set of agents N is partitioned into k known classes N_1, \dots, N_k so that $N_i \cap N_j = \emptyset$ for all $i \neq j$ and $\cup_{i=1}^k N_i = N$. For simplicity, we refer to class N_i simply as class i .

Matching. We consider the cases of *divisible* items (where each item can be matched to multiple agents fractionally) and *indivisible* items (where each item must be matched to a single agent integrally). A (divisible) *matching* is a matrix $X = (x_{a,o})_{a \in N, o \in M} \in [0, 1]^{N \times M}$ satisfying $\sum_{a \in N} x_{a,o} \leq 1$ for each item $o \in M$, and $\sum_{o \in M} x_{a,o} \leq 1$ for each agent $a \in N$. We say that matching X is *indivisible* if $x_{a,o} \in \{0, 1\}$ for each agent $a \in N$ and item $o \in M$. Given a matching X , we say that agent a is *saturated* if $\sum_{o \in M} x_{a,o} = 1$, and item o is *fully assigned* if $\sum_{a \in N} x_{a,o} = 1$.

For a matching X , we write $Y(X) = (\sum_{a \in N_i} x_{a,o})_{i \in [k], o \in M}$ as the matrix containing the total fraction of each item assigned to agents in each class. Let $Y_i(X)$ denote the row of $Y(X)$ corresponding to class i . For an indivisible matching X , we may abuse the notation and use $Y_i(X)$ to refer to the set of items matched to agents in class i , i.e., $\{o \in M \mid x_{a,o} = 1 \text{ for some } a \in N_i\}$. We may omit the argument X from $Y(X)$ and $Y_i(X)$ if it is clear from the context.

Class valuations. The value derived by agent a from matching X is $V_a(X) = \sum_{o \in M: (a,o) \in E} x_{a,o}$. We define the value of class i from matching X as the utilitarian social welfare of the agents in class i under matching X , denoted $V_i(X) = \sum_{a \in N_i} V_a(X)$.

In order to define fairness at the level of classes, we need to also define how much hypothetical value agents in class i could derive from the items matched to agents in another class j . However, it is not obvious how one should define this value because it depends on how the items matched to agents in N_j would be matched to agents in N_i in this hypothetical scenario. Following (Benabbou et al. 2019), we use the following optimistic valuations.

Given a vector $\mathbf{y} = (y_o)_{o \in M} \in [0, 1]^M$ representing fractions of different items, the *optimistic valuation* $V_i^*(\mathbf{y})$ of

class i for \mathbf{y} is the size of the maximum fractional matching between the agents of N_i and \mathbf{y} ; namely, $V_i^*(\mathbf{y})$ is given by the optimal value of the following LP:

$$\begin{aligned} \max \quad & \sum_{a \in N_i} \sum_{o \in M: (a,o) \in E} x_{a,o} \\ \text{s.t.} \quad & \sum_{a \in N_i} x_{a,o} \leq y_o & \forall o \in M, \\ & \sum_{o \in M} x_{a,o} \leq 1 & \forall a \in N_i, \\ & x_{a,o} \geq 0 & \forall a \in N_i, o \in M. \end{aligned}$$

For $S \subseteq M$, let $e^S \in \{0, 1\}^M$ denote the incidence vector such that $e_o^S = 1$ if $o \in S$ and $e_o^S = 0$ otherwise; we may write $V_i^*(e^S)$ as $V_i^*(S)$ for ease of notation. For an integral vector \mathbf{y} , it is known that there is an integral optimal solution to the above LP (see, e.g., Section 5 of (Korte and Vygen 2006)); thus, $V_i^*(S)$ coincides with the maximum size of an integral matching between S and the agents in N_i .

2.1 Solution Concepts

We consider classical fairness notions from the fair division literature and extend these notions to ensure fairness between classes of agents.

(Approximate) class envy-freeness. Envy-freeness between individual agents demands that every agent values the resources allocated to her at least as much as she values the resources allocated to another agent. When applied to classes, we compare the value $V_i(X)$ derived by class i for its matched items with class i 's optimistic valuation for the items matched to another class j , i.e. $V_i^*(Y_j(X))$. Note that this results in a strong class envy-freeness notion: *even if*, hypothetically, class i were to be matched to the items currently matched to class j under X in an optimal manner, they would still not be any happier overall.

Definition 1 (Class envy-freeness). A matching X is α -class envy-free (α -CEF) if for all classes $i, j \in [k]$, $V_i(X) \geq \alpha \cdot V_i^*(Y_j(X))$. When $\alpha = 1$, we simply refer to it as class envy-freeness (CEF).

It is impossible to achieve exact CEF with an indivisible matching in general, e.g., consider when one desirable item has to be allocated among two classes. Hence, we consider the following relaxation of CEF for integral matchings.

Definition 2 (Class envy-freeness up to one item). An integral matching X is α -class envy-free up to one item (α -CEF1) if for every pair of classes $i, j \in [k]$, either $Y_j(X) = \emptyset$ or there exists an item $o \in Y_j(X)$ such that $V_i(X) \geq \alpha \cdot V_i^*(Y_j(X) \setminus \{o\})$. When $\alpha = 1$, we simply refer to it as class envy-freeness up to one item (CEF1).

We remark that CEF1 is called type-wise EF1 (TEF1) by (Benabbou et al. 2019); we use the terminology “class” instead of “type” because letting agents of the same “type” have different incident edges may be confusing.

(Approximate) class proportionality and maximin share fairness. Another classical fairness concept is *proportionality*. In the traditional fair division model where agent valuations are additive, proportionality is typically stated as requiring that each agent receives value that is at least $1/n$ -th of her value for the set of all items, where n is the number of agents. This can be equivalently viewed as demanding that each agent receives at least the maximum value she can receive from the worst bundle among all fractional partitions of the items into n bundles. We use the latter version as the appropriate definition of proportionality in our model. We define the *proportional share* of class i as

$$\text{prop}_i = \max_{X \in \mathcal{X}} \min_{j \in [k]} V_i^*(Y_j(X)),$$

where \mathcal{X} is the set of (divisible) matchings of the set of items M to the set of agents N .

Definition 3 (Class proportionality). We say that matching X is α -class proportional (α -CPROP) if for every class $i \in [k]$, $V_i(X) \geq \alpha \cdot \text{prop}_i$. When $\alpha = 1$, we simply refer to it as class proportionality (CPROP).

As in the case with class envy-freeness, class proportionality is impossible to guarantee via indivisible matchings. Nevertheless, we can naturally relax the notion of proportionality by only taking into account indivisible matchings in the definition of proportional share above. Formally, the *maximin share* of class i is defined as

$$\text{mms}_i = \max_{X \in \mathcal{I}} \min_{j \in [k]} V_i^*(Y_j(X)),$$

where \mathcal{I} is the set of indivisible matchings of the set of items M to the set of agents N .

Definition 4 (Class maximin share fairness). We say that matching X is α -class maximin share fair (α -CMMS) if for every class $i \in [k]$, $V_i(X) \geq \alpha \cdot \text{mms}_i$. When $\alpha = 1$, we simply refer to it as class maximin share fairness (CMMS).

Efficiency. We consider two notions of efficiency. *Non-wastefulness* demands that each item to be fully assigned, unless all the agents who like it are saturated.

Definition 5 (Non-wastefulness). We say that matching X is *non-wasteful* (NW) if there is no pair of agent a and item o such that (i) o is allocated to a (i.e., $x_{a,o} > 0$) but a does not like o , or (ii) a likes o , a is not saturated (i.e., $\sum_{o' \in M} x_{a,o'} < 1$), and o is not fully assigned (i.e., $\sum_{a' \in N} x_{a',o} < 1$).

A more quantitative notion of efficiency is the utilitarian social welfare, which, in our context, is the size of the (divisible) matching. Note that this is the classical objective that the literature on online matching optimizes, in the absence of any fairness constraints.

Definition 6 (Utilitarian social welfare). The *utilitarian social welfare* (USW) of a matching X is given by $\text{usw}(X) = \sum_{a \in N} \sum_{o \in M: (a,o) \in E} x_{a,o}$. We say that a divisible (resp., indivisible) matching X is α -USW if $\text{usw}(X) \geq \alpha \cdot$




Figure 2: Class envy-freeness (CEF), non-wastefulness (NW), and utilitarian social welfare approximation (USW): an empty matching is CEF1 but wasteful; wiggly lines show a CEF1 and NW matching; thick lines indicate a CEF1 and 1-USW matching.

$\text{usw}(X^*)$ for all divisible (resp., indivisible) matchings X^* . When $\alpha = 1$, we refer to X as the USW-optimal matching.

The following is a known relation between maximal (non-wasteful) and maximum matchings in both divisible and indivisible cases.

Proposition 1. *Every non-wasteful (divisible or indivisible) matching is $1/2$ -USW.*

Let us illustrate the above concepts of fairness and efficiency using examples.

Example 2. Consider the example given in Figure 2, where there are four items (o_1, o_2, o_3 , and o_4), agents a_1 and a_2 belong to one class, and agents b_1 and b_2 belong to another class. An edge between an agent and an item indicates that the agent likes the item; thick and wiggly lines indicate matchings. An empty matching is class envy-free (CEF) but wasteful. The wiggly lines show a CEF1 and non-wasteful matching. Finally, thick lines show a matching that achieves CEF1 along with optimal utilitarian social welfare.

2.2 Online Model

Let us now introduce our online model. In this model, the items in M arrive one-by-one in an arbitrary order. We refer to the step in which item $o \in M$ arrives as step o .

When item o arrives, all agents reveal whether or not they like the item. In other words, the edges incident to item o are revealed in graph G . At this point, an online algorithm must make an immediate and irrevocable decision to “match” the item to the agents in N , i.e., set the values of $(x_{a,o})_{a \in N}$. We consider algorithms which set these values *deterministically*.

For the algorithms we design, we prove that they achieve the desired guarantees (approximate CEF, CEF1, CPROP, CMMS, USW, or non-wastefulness) at every step. However, a key property of our algorithms is that they do not need to know in advance the number of items that will arrive, which means that proving the desired guarantees at the end implies that they hold at every step. In contrast, our upper bounds (impossibility results) will hold even if the desired guarantees are required to hold only at the end.

Definition 7. For $\alpha \in (0, 1]$, a deterministic online algorithm for matching divisible or indivisible items is α -CEF (resp., α -CEF1, α -CPROP, α -CMMS, α -USW, or NW) if it produces an α -CEF (resp., α -CEF1, α -CPROP, α -CMMS, α -USW, or NW) matching when all items have arrived.

Because CMMS and CPROP place only a lower bound on the utility of every agent, there is no tension between them and non-wastefulness. Any algorithm achieving an approximation of these notions can be made non-wasteful without losing the said fairness approximation.

Proposition 2. *For $\alpha \in (0, 1]$, if there is a deterministic online algorithm satisfying α -CMMS (resp., α -CPROP), then there is a non-wasteful deterministic online algorithm satisfying α -CMMS (resp., α -CPROP). This holds for matching both divisible and indivisible items.*

3 Indivisible Items

We start by focusing on deterministic algorithms for matching indivisible items. We study possible approximations of two fairness concepts, CEF1 and CMMS, along with efficiency guarantees in terms of non-wastefulness and the utilitarian social welfare. When matching indivisible items, CEF1 may seem trivial to achieve: only match an item to some agent in some class if this preserves CEF1, and discard the item otherwise. However, this algorithm may ‘waste’ too many items and lose significant efficiency.² Example 1 illustrated that CEF1 and non-wastefulness are incompatible in the online setting.³ In this light, for arbitrary classes, it is natural to ask what approximation of CEF1 can be achieved subject to non-wastefulness.

3.1 Algorithm MATCH-AND-SHIFT

One way to achieve approximate CEF1 is to ensure a balanced treatment of all classes by providing them approximately equal ‘opportunity’ for receiving an item. This approach is inspired by the well-studied *Round-Robin* algorithm in fair division (Caragiannis et al. 2016) and its widely-adopted cousin, *Draft*, that is used in sports for selecting players (Brams and Straffin 1979; Brams and Taylor 2000) or assigning courses to college students (Budish and Cantillon 2012).

However, running such algorithms naïvely in our online setting, where not all items are available upfront, can be problematic: if we do a round-robin over classes, a class can be disadvantaged if the item arriving in its turn is not liked by any unmatched agent in the class. Further, non-wastefulness requires that any arriving item be matched as long as there is an unsaturated agent who likes it, even if this agent does not belong to the class whose turn it is. Keeping these observations in mind, we design MATCH-AND-SHIFT (Algorithm 1), which provides equal treatment to the different classes while achieving non-wastefulness.

Algorithm description. Fix an arbitrary priority ordering $\pi = (\pi_1, \pi_2, \dots, \pi_k)$ over the k classes, where π_1 is the class with the highest priority. Upon arrival of each item,

²In fact, discarding all items—an empty matching—is vacuously class envy-free.

³In Appendix B.1 of the full version (Hosseini et al. 2022), we show that this incompatibility holds even after weakening the CEF1 requirement to account for ‘pessimistic’ valuations, i.e. when each class evaluates the items matched to another class through a *minimum-cardinality maximal matching*.

ALGORITHM 1: MATCH-AND-SHIFT

```

1 Fix a priority ordering over classes,  $\pi = (\pi_1, \dots, \pi_k)$ 
2 when item  $o \in M$  arrives do
3   for  $i = 1$  to  $k$  do
4     Let  $N_{\pi_i, o}$  be the set of unmatched agents  $a \in N_{\pi_i}$ 
       such that  $(a, o) \in E$ 
5     if  $N_{\pi_i, o} \neq \emptyset$  then
6       Arbitrarily match  $o$  to an agent in  $N_{\pi_i, o}$ 
7        $\pi \leftarrow (\pi_1, \dots, \pi_{i-1}, \pi_{i+1}, \dots, \pi_k, \pi_i)$ 
8       break

```

pick the first class N_{π_i} in the priority ordering that contains an unmatched agent who likes the item. Match the item to any unmatched agent—there may be several such agents—in N_{π_i} who likes the item. Update the priority ordering π by moving class π_i to the end.

The following theorem establishes approximate fairness and efficiency guarantees of MATCH-AND-SHIFT; later, in Theorem 2, we prove that these guarantees are tight.

Theorem 1. *For deterministic matching of indivisible items, MATCH-AND-SHIFT (Algorithm 1) satisfies non-wastefulness, $1/2$ -CEF1, $1/2$ -CMMS, and $1/2$ -USW.*

Proof. Let X be the matching returned by the algorithm.

NW & $1/2$ -USW. Non-wastefulness of X follows immediately from the description of the algorithm: at each step, the arriving item is matched to an agent who likes it whenever such an agent exists. Because X is non-wasteful, due to Proposition 1 it also satisfies $1/2$ -USW.

Now, we turn our attention to the fairness guarantees. Recall that for each $i \in [k]$, Y_i denotes the set of items matched to agents in class i . Fix any class i . Let $t = |Y_i|$ denote the number of items matched to the agents in class i under X . Due to non-wastefulness, we have $V_i(X) = t$.

$1/2$ -CEF1. Consider any class $j \in [k] \setminus \{i\}$. Let $Y_j^* \subseteq Y_j$ be the set of items matched to class j that are liked by at least one unmatched agent in class i . The claim immediately holds when $Y_j^* = \emptyset$: in this case, the optimistic value of class i for Y_j is $V_i^*(Y_j) \leq t = V_i(X)$, implying that X satisfies CEF for i . Thus, we assume that at least one item in Y_j is liked by at least one unmatched agent of class i .

By construction of the algorithm, we have $|Y_j^*| \leq t + 1$. This is because every time class j receives an item in Y_j^* (that is liked by an agent in class i who remains unmatched till the end, and, therefore, is unmatched at the time of the item’s arrival), class j must have a higher priority than class i . Hence, the algorithm must match an item to class i before it can match another item in Y_j^* to class j . Thus, $|Y_j^*| \leq 1 + |Y_i| = t + 1$. Fix an arbitrary item $o \in Y_j^* \subseteq Y_j$. We claim that $V_i^*(Y_j \setminus \{o\}) \leq 2t$, which establishes the $1/2$ -CEF1 claim. Note that the t matched agents in class i can derive a maximum total utility of t from these items. Further, the total utility that the unmatched agents in class i can derive from these items is upper bounded by $|Y_j^* \setminus \{o\}| \leq t$. Hence, $V_i^*(Y_j \setminus \{o\}) \leq 2t$.

$1/2$ -CMMS. Assume for contradiction that $t = V_i(X) < (1/2) \cdot mms_i$. Because mms_i is an integer, this implies $2t +$

$1 \leq \text{mms}_i$. Let (S_1, S_2, \dots, S_k) be a maximin partition of the items for class i such that $V_i^*(S_j) \geq \text{mms}_i$ for every $j \in [k]$. By our assumption, we have $V_i^*(S_j) \geq 2t + 1$ for every $j \in [k]$. For each $j \in [k]$, we let S_j^* denote the set of items in S_j that are liked by at least one unmatched agent in class i . Note that $V_i^*(S_j) \leq t + |S_j^*|$: the t matched agents in class i can derive total utility at most t , and the unmatched agents can derive total utility at most $|S_j^*|$.

Recall that $|Y_i| = t$ and we have already established $|Y_j^*| \leq t + 1$ for every class $j \in [k] \setminus \{i\}$. Further, by non-wastefulness, none of the unmatched agents of class i likes any item in $M \setminus \bigcup_{h \in [k]} Y_h$. Thus, we have $|\bigcup_{j \in [k]} S_j^*| \leq |Y_i \cup (\bigcup_{j \in [k] \setminus \{i\}} Y_j^*)| \leq t + (k - 1)(t + 1)$, meaning that there exists some $h \in [k]$ such that $|S_h^*| \leq t$. Thus, we have $V_i^*(S_h) \leq 2t < 2t + 1$, a contradiction. \square

Before we turn to proving these guarantees to be the best possible in our online setting, we remark that in the offline setting, it is known that (exact) CEF1 and NW can be achieved simultaneously (Benabbou et al. 2019). However, whether they can be achieved together with α -CMMS, for any $\alpha > 0$, is an interesting open question.

3.2 Impossibility Results

In this section, we show that each of the fairness and efficiency guarantees achieved by MATCH-AND-SHIFT (Theorem 1) is tight; no deterministic online algorithm for matching indivisible items can achieve a better approximation. Note that our CEF1 upper bound is subject to non-wastefulness because an algorithm can trivially achieve CEF1 by throwing away every item. The constructions are based on creating instances in which a subset of agents in one class gets saturated early on, rendering the class envious of another class at the end since all the remaining items can only be matched to the agents in that other class.

Theorem 2. *No deterministic online algorithm for matching indivisible items can achieve any of the following guarantees:*

- α -CEF1 for any $\alpha > 1/2$ and non-wastefulness,
- α -CMMS for any $\alpha > 1/2$,
- α -USW for any $\alpha > 1/2$.

Proof. We argue each impossibility result separately.

CEF1 and NW. Consider Example 1 in the introduction. In that example, we argued that any deterministic online algorithm satisfying non-wastefulness ends up matching (without loss of generality) $Y_2 = \{o_2, o_3, o_4\}$ to class 2 and $Y_1 = \{o_1\}$ to class 1. One can check that $V_1^*(Y_2 \setminus \{o\}) = 2$ for any $o \in Y_2$, whereas $V_1(X) = 1$, implying that the algorithm cannot achieve α -CEF1 for any $\alpha > 1/2$.

CMMS. We will prove that no deterministic online algorithm satisfying non-wastefulness can achieve α -CMMS for any $\alpha > 1/2$. Proposition 2 implies that no deterministic algorithm, regardless of whether it satisfies non-wastefulness, can guarantee α -CMMS for any $\alpha > 1/2$.

Since we have assumed non-wastefulness, we can repeat the construction used above for the CEF1 upper bound. Consider the same example again, and consider the partition

the items into $(\tilde{Y}_1 = \{o_1, o_2\}, \tilde{Y}_2 = \{o_3, o_4\})$. Note that $V_1^*(\tilde{Y}_1) = V_1^*(\tilde{Y}_2) = 2$, implying that the maximin share of class 1 is $\text{mms}_1 \geq 2$. Since the value derived by class 1 is $V_1(X) = 1$, we see that the algorithm cannot achieve α -CMMS for any $\alpha > 1/2$.

USW. Note that the USW guarantee does not depend on the class structure; hence, the well-known upper bound of $1/2$ on the approximation of a maximum matching by any deterministic algorithm carries over to our model, and implies the desired $1/2$ -USW upper bound. \square

Following Theorem 2, a natural question is whether there is any way to circumvent this impossibility result. We show that two such approaches do not work, demonstrating robustness of Theorem 2.

Remark 1 (Reshuffling items within each class cannot help.). One idea is to only require the online algorithm to match each item to a class, and allow every class to optimally distribute the items matched to it among its members at the end. This effectively increases the utility of class i from $V_i(X)$ to $V_i^*(Y_i)$. However, in Example 1 used for the CEF1 and CMMS upper bounds in the proof above, the matching produced already assigns items optimally within each class (i.e., satisfies $V_i(X) = V_i^*(Y_i)$ for each class i). Hence, reshuffling items at the end cannot improve the value any further. This shows that we must use randomization when deciding which class should receive an item in order to achieve a better approximation.

Remark 2. Another natural direction is to weaken the requirements in Theorem 2. In our online setting, there is a weakening of our α -CMMS guarantee that also makes sense. Instead of computing the MMS values by partitioning the set of *all* items, we can first observe the matching X produced by an algorithm and then compute the MMS values by having each class partition only the set of items *allocated* under X . This produces smaller (or equal) values, making this CMMS with respect to *allocated* items a weaker requirement than our CMMS with respect to *all* items.

MATCH-AND-SHIFT achieves a $1/2$ -approximation of the stronger requirement. In contrast, the proof of Theorem 2 shows that no non-wasteful⁴ algorithm can achieve $(1/2 + \epsilon)$ -approximation of even the weaker requirement, for any $\epsilon > 0$, because all items are allocated in our construction.

4 Divisible Items

We now turn our attention to online matching of *divisible* items. First, we design an algorithm that simultaneously achieves non-wastefulness, $(1 - 1/e)$ -CEF, $(1 - 1/e)$ -CPROP, and $1/2$ -USW. Later, we prove upper bounds on the approximation ratio of each guarantee that hold for any algorithm.

4.1 Algorithm EQUAL-FILLING

We propose an algorithm, EQUAL-FILLING (presented as Algorithm 2), that divides items equally at the class level and performs water-filling to further divide the items assigned

⁴Seeking the weaker requirement makes sense only with non-wastefulness since the empty matching vacuously satisfies it.

ALGORITHM 2: EQUAL-FILLING

```
1 Initialize  $X = (x_{a,o})_{a \in N, o \in M}$  so that  $x_{a,o} = 0$  for every agent  $a$  and item  $o$ 
2 Initialize  $Y = (y_{i,o})_{i \in [k], o \in M}$  so that  $y_{i,o} = 0$  for every class  $i$  and item  $o$ 
3 when item  $o \in M$  arrives do
4   /*class-phase*/
5   Let  $N_{i,o}$  denote the set of neighbours of item  $o$  in class  $i$ , i.e.,  $N_{i,o} = \{a \in N_i : (a,o) \in E\}$ 
6   Define the demand of each class  $i \in [k]$  as
      $d_{i,o} = \sum_{a \in N_{i,o}} (1 - \sum_{o' \in M} x_{a,o'})$ 
7   Find the largest  $\beta_o \leq 1$  satisfying
      $\sum_{i \in [k]} \min\{\beta_o, d_{i,o}\} \leq 1$ 
8   Set  $y_{i,o} = \min\{\beta_o, d_{i,o}\}$  for each  $i \in [k]$ 
9   for  $i = 1$  to  $k$  do
10    /*individual-phase*/
11    Find the largest  $\gamma_{i,o} \leq 1$  satisfying
        $\sum_{j \in N_{i,o}} \max\{\gamma_{i,o} - \sum_{o' \in M} x_{a,o'}, 0\} \leq y_{i,o}$ 
12    Set  $x_{a,o} = \max\{\gamma_{i,o} - \sum_{o' \in M} x_{a,o'}, 0\}$  for all
        $a \in N_{i,o}$ 
```

to each class between the agents in that class. Recall that our model has a capacity constraint: $\sum_{o \in M} x_{a,o} \leq 1$ for each agent a . Agent a is saturated if $\sum_{o \in M} x_{a,o} = 1$, and unsaturated otherwise.

When item o arrives, EQUAL-FILLING continuously splits the item equally among classes with at least one unsaturated agent who likes the item.⁵ At the end of this process, each class either receives the same fraction β_o of the item, or has all of its agents who like item o saturated. This computation is performed in Line 8 of Algorithm 2. Then, to divide fraction of item o assigned to each class i within its members, we conduct water-filling among the members who like o , which continuously prioritizes agents with the lowest utility. At the end of this process, each member who likes item o either receives the same final utility $\gamma_{i,o}$ or is saturated. This computation is performed in Line 12 of Algorithm 2.

Theorem 3. *For deterministic matching of divisible items, EQUAL-FILLING (Algorithm 2) satisfies non-wastefulness, $(1 - 1/e)$ -CEF, $(1 - 1/e)$ -CPROP, and $1/2$ -USW.*

4.2 Impossibility Results

Our goal in this section is to provide upper bounds on the fairness and efficiency guarantees that hold for any deterministic online algorithm for matching divisible items. We prove that the $(1 - 1/e)$ -CPROP guarantee achieved by EQUAL-FILLING is tight, and establish a weaker upper bound on CEF and USW.

Theorem 4. *No deterministic online algorithm for matching divisible items can achieve any of the following guarantees:*

- α -CEF for any $\alpha > 3/4$ and non-wastefulness,
- α -CPROP for any $\alpha > 1 - 1/e$,

⁵We do not yet need to know how the fraction of item o assigned to a class is divided between its members; we can simply keep track of the total remaining capacity of the agents in the class who like the item.

- α -USW for any $\alpha > 1 - 1/e$.

Remark 3. Similar to Remark 2, one may wonder what we can say about a weaker notion of proportionality with respect to only the *allocated items*, i.e., if the proportional share of each class is defined on the divisible matchings of the allocated items (instead of all items). In Proposition 7 in the full version (Hosseini et al. 2022), we show that the upper bound of $1 - 1/e$ continues to hold even for this weaker version. However, unlike in the case of indivisible items, this does not follow from the proof above (which considers an instance with a single class, for which, trivially, the weaker version is exactly satisfied). The proof of this proposition is much more intricate.

While EQUAL-FILLING achieves the optimal $1 - 1/e$ approximation of CPROP, its guarantees with respect to CEF and USW identified in Theorem 3 are weaker than the upper bounds in Theorem 4. One might wonder if this is simply because our analysis in Theorem 3 is loose. We show that this is not the case. Hence, future work must focus either on proving better upper bounds, or on designing new algorithms which might surpass EQUAL-FILLING.

Proposition 3. *EQUAL-FILLING does not achieve any of the following guarantees:*

- α -CEF for any $\alpha > 1 - 1/e$,
- α -CPROP for any $\alpha > 1 - 1/e$,
- α -USW for any $\alpha > 1/2$.

5 Discussion

Our work introduces the novel framework of class fairness in online matching. We derive bounds on approximate fairness and efficiency guarantees that deterministic and randomized online algorithms can achieve in this framework for matching divisible and indivisible items, and leave open a number of exciting open questions. For example, can a deterministic algorithm for matching divisible items achieve a CEF approximation together with non-wastefulness better than $1 - 1/e$? (We conjecture the answer to be *no*.) Can it achieve any reasonable CEF or CPROP approximation together with a USW approximation better than $1/2$ (ideally, $1 - 1/e$)? Can a randomized algorithm for matching indivisible items achieve any reasonable CEF approximation together with either non-wastefulness or a USW approximation? In the full version (Hosseini et al. 2022), we explore randomized algorithms for allocating indivisible goods in hopes of breaking the $1/2$ barrier. We discuss the challenges in devising fair randomized algorithms and develop an extension to the EQUAL-FILLING algorithm that simultaneously achieves 0.593-CPROP and $1/2$ -USW in this setting.

More broadly, our basic framework paves the road for interesting extensions. For example, one can allow agents to have non-binary values for the items, consider class fairness notions that give more importance to bigger classes, consider both agents and items arriving online (Huang et al. 2020), study weaker adversarial models, or consider stochastic instead of adversarial arrivals.

Acknowledgments

Hadi Hosseini acknowledges support from NSF grants #2144413 (CAREER), #2052488, and #2107173. Zhiyi Huang was supported in part by an RGC grant #17201221. Ayumi Igarashi was supported by JST PRESTO under grant number JPMJPR20C1. Nisarg Shah was partially supported by an NSERC Discovery Grant.

References

- Aleksandrov, M.; Aziz, H.; Gaspers, S.; and Walsh, T. 2015. Online fair division: Analysing a food bank problem. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2540–2546.
- Azar, Y.; and Richter, Y. 2005. Management of multi-queue switches in QoS networks. *Algorithmica*, 43(1): 81–96.
- Babaioff, M.; Ezra, T.; and Feige, U. 2021. Fair and truthful mechanisms for dichotomous valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 5119–5126.
- Banerjee, S.; Gkatzelis, V.; Gorokh, A.; and Jin, B. 2022. Online Nash social welfare maximization with predictions. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1–19. SIAM.
- Banerjee, S.; and Johari, R. 2019. Ride sharing. In *Sharing Economy*, 73–97. Springer.
- Barman, S.; and Verma, P. 2021. Existence and computation of maximin fair allocations under matroid-rank valuations. In *Proceedings of the 20th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 169–177.
- Benabbou, N.; Chakraborty, M.; Elkind, E.; and Zick, Y. 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.
- Benabbou, N.; Igarashi, A.; Chakraborty, M.; and Zick, Y. 2020. Finding fair and efficient allocations when valuations don't add up. In *Proceedings of the 13th International Symposium on Algorithmic Game Theory (SAGT)*, 32–46.
- Benade, G.; Kazachkov, A. M.; Procaccia, A. D.; and Psomas, C.-A. 2018. How to make envy vanish over time. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 593–610.
- Brams, S. J.; and Straffin, P. D. 1979. Prisoner's dilemma and professional sports drafts. *American Mathematical Monthly*, 80–88.
- Brams, S. J.; and Taylor, A. D. 2000. *The Win-Win solution: Guaranteeing fair shares to everybody*. WW Norton & Company.
- Budish, E.; and Cantillon, E. 2012. The multi-unit assignment problem: Theory and evidence from course allocation at Harvard. *The American Economic Review*, 102(5): 2237–71.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2016. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, 305–322. ACM.
- Hosseini, H.; Huang, Z.; Igarashi, A.; and Shah, N. 2022. Class fairness in online matching. arXiv:2203.03751.
- Huang, Z.; Kang, N.; Tang, Z. G.; Wu, X.; Zhang, Y.; and Zhu, X. 2020. Fully online matching. *Journal of the ACM (JACM)*, 67(3): 1–25.
- Karp, R. M.; Vazirani, U. V.; and Vazirani, V. V. 1990. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*, 352–358. ACM.
- Korte, B.; and Vygen, J. 2006. *Combinatorial Optimization: Theory and Algorithms*. Springer-Verlag Berlin Heidelberg.
- Lee, M. K.; Kusbit, D.; Kahng, A.; Kim, J. T.; Yuan, X.; Chan, A.; See, D.; Noothigattu, R.; Lee, S.; Psomas, A.; et al. 2019. WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW): 1–35.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM conference on Electronic commerce*, 125–131.
- Mehta, A. 2013. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science*, 8(4): 265–368.
- Mehta, A.; Saberi, A.; Vazirani, U.; and Vazirani, V. 2007. Adwords and generalized online matching. *Journal of the ACM (JACM)*, 54(5): 22–es.
- Moulin, H. 2019. Fair division in the Internet age. *Annual Review of Economics*, 11(1): 407–441.
- Paes Leme, R. 2017. Gross substitutability: An algorithmic survey. *Games and Economic Behavior*, 106: 294 – 316.
- Walsh, T. 2011. Online cake cutting. In *International Conference on Algorithmic Decision Theory*, 292–305. Springer.
- Zeng, D.; and Psomas, A. 2020. Fairness-efficiency trade-offs in dynamic fair division. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 911–912.