

Class Fairness in Online Matching

HADI HOSSEINI, Penn State University

ZHIYI HUANG, University of Hong Kong

AYUMI IGARASHI, National Institute of Informatics

NISARG SHAH, University of Toronto

In the classical version of online bipartite matching, there is a given set of offline vertices (aka agents) and another set of vertices (aka items) that arrive online. When each item arrives, its incident edges—the agents who like the item—are revealed and the algorithm must irrevocably match the item to such agents.

We initiate the study of *class fairness* in this setting, where agents are partitioned into a set of classes and the matching is required to be fair with respect to the classes. We adopt popular fairness notions from the fair division literature such as envy-freeness (up to one item), proportionality, and maximin share fairness to our setting. Our class versions of these notions demand that all classes, regardless of their sizes, receive a fair treatment. We study deterministic and randomized algorithms for matching indivisible items (leading to integral matchings) and for matching divisible items (leading to fractional matchings).

We design and analyze three novel algorithms. 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. Finally, we build upon *EQUAL-FILLING* to design a randomized algorithm for matching indivisible items, *EQUAL-FILLING-OCS*, which achieves 0.593-approximation of class proportionality. The algorithm and its analysis crucially leverage the recently introduced technique of online correlated selection (OCS) [Fahrback et al., 2020].

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 et al. [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 equally, 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).

Consider, for example, the appealing notion of envy-freeness, which, when applied to individual agents, demands that no agent envy the resources given to another agent. When applied to classes, our *class envy-freeness* (CEF) notion requires that no class of agents be able to increase their total value by taking the items matched to another class, *even if* it assigns these items optimally among its members. With indivisible items (which must be assigned entirely to a single agent), a class envy-free matching may not always exist: consider a single item to be divided between two classes with one agent each liking the item. In the standard fair division model, this impossibility 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 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. In the offline setting wherein all items are




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

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 a classical algorithm that is blind to the class information can easily violate CEF1. For example, if there are two classes containing two agents each, and two items arrive that are liked by all four agents, the algorithm may end up assigning both items to agents from the same class, rendering the other class envious even if we remove one of the items. This simple example is easy to fix via a “class-aware” algorithm that pays attention to the classes: simply assign the second item to an agent from the class that did not receive the first item. Alas, a slightly larger example shows that even class-aware online algorithms cannot always achieve CEF1.

Example 1. Consider the example shown in Figure 1, in which six agents are partitioned into two classes $N_1 = \{a_1, a_2, a_3\}$ and $N_2 = \{b_1, b_2, b_3\}$, and four items arrive sequentially in the order (o_1, o_2, o_3, o_4) . An edge between an agent and an item indicates that the agent likes the item; thick edges indicate the matching. Let us assume that we do not wish to throw away any item as long as there is an unmatched agent who likes it; we later formalize this as *non-wastefulness*.

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 and b_1 ; without loss of generality, suppose it is matched to $a_1 \in N_1$. When the second item o_2 arrives, note that 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, the fourth item o_4 arrives, and the algorithm learns that 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?

1.1 Our Results

We initiate the study of fairness among classes of agents in online bipartite matching. Our first contribution (Section 2) is developing a detailed mathematical framework in which we adopt classical fairness concepts to online matching. We consider two types of online matching models,

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.

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.

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:** When items are indivisible, 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 $1/2$ -USW, for any $\alpha > 1/2$ (Theorem 2), establishing our algorithm to be simultaneously optimal for each guarantee.
- **Divisible matching:** When items are divisible, 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). Furthermore, 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).
- **Randomized algorithms:** Finally, we propose a randomized algorithm, EQUAL-FILLING-OCS, for matching indivisible algorithms that breaks the $1/2$ barrier. We run a variant of EQUAL-FILLING to obtain a guiding divisible matching, and round it into an indivisible matching using a technique called online correlated selection (OCS). We prove that it is simultaneously 0.593 -CPROP and $1/2$ -USW (Theorem 5).

1.2 Related Work

Online matching. 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. The RANKING algorithm of Karp et al. [1990] assigns each item in its entirety; in our model, this corresponds to a randomized algorithm for matching indivisible items that achieves $(1 - 1/e)$ -USW. The case of divisible items is often called *fractional online matching* in the matching literature.² For this, Kalyanasundaram and Pruhs [2000] gave a deterministic $(1 - 1/e)$ -competitive algorithm, which achieves $(1 - 1/e)$ -USW in our framework; different papers refer to this algorithm as Balance,

²It is closely related to another model called online b -matching in which each offline agent may be matched up to b times. Since the algorithms and analyses are usually interchangeable in these two models, we phrase both models as the case of divisible items.

Water-filling, or Water-level. The RANKING algorithm and its analysis were generalized to the vertex-weighted case by Aggarwal et al. [2011]. Feldman et al. [2009] introduced the free disposal model of edge-weighted online matching and gave a $(1 - 1/e)$ -competitive algorithm for divisible items. The series of works by Fahrback et al. [2020], Shin and An [2021], Gao et al. [2021], and Blanc and Charikar [2021] led to the state-of-the-art 0.536-competitive algorithm for edge-weighted online matching with indivisible items. These works developed a new technique called online correlated selection which we also use in this paper.

The literature also considers stochastic models of online matching problems to break the $1 - 1/e$ barrier. Mahdian and Yan [2011] and Karande et al. [2011] showed that the competitive ratio of RANKING is between 0.696 and 0.727 if online vertices arrive by a random order. Huang et al. [2019] introduced a variant of RANKING that breaks the $1 - 1/e$ barrier in vertex-weighted online matching under random-order arrivals; the ratio was further improved to 0.668 [Jin and Williamson, 2020]. If items are drawn from a distribution known to the algorithm, it is called *online stochastic matching* [Feldman et al., 2009]. The best known competitive ratios for unweighted and vertex-weighted online stochastic matching are 0.711 and 0.700, respectively [Huang and Shu, 2021].

Fair division. There is a rich body of literature on fair allocation of indivisible or divisible items. A common assumption in most fair division studies is that there is no constraint on how many items each agent can receive, and agents receive increasing value when receiving more items.

In this literature, envy-freeness and proportionality (and approximations thereof) have been used as the primary criteria of fairness. For divisible items, an allocation satisfying both envy-freeness and an economic efficiency notion called Pareto optimality is known to exist [Varian, 1974] and can be computed via convex programming when agents have additive valuations [Eisenberg and Gale, 1959]. For indivisible items, two relaxations of envy-freeness are commonly studied: envy-freeness up to one item (EF1) [Lipton et al., 2004] and maximin share fairness (MMS) [Budish, 2011]. An EF1 allocation is guaranteed to exist with monotone valuations [Lipton et al., 2004], and can be achieved together with Pareto optimality when agents have additive valuations [Caragiannis et al., 2016]. On the other hand, MMS allocations are not guaranteed to exist, even for additive valuations, though constant factor approximation algorithms [Garg and Taki, 2020, Ghodsi et al., 2018, Kurokawa et al., 2018] and ordinal approximations [Hosseini and Searns, 2021, Hosseini et al., 2021] exist and can be computed in polynomial time.

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 observed that the EF1 algorithm of Lipton et al. [2004] may result in a wasteful allocation; nevertheless, they showed that an allocation satisfying EF1 and non-wastefulness exists and can be computed in polynomial time. Subsequent papers [Babaioff et al., 2021, Barman and Verma, 2021, Benabbou et al., 2020] 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.

Fairness in online matching. Our paper is also related to the growing line of work on *online* fair division [Aleksandrov et al., 2015, Benade et al., 2018, Gorokh et al., 2022, Walsh, 2011, Zeng and Psomas, 2020], but a majority of this work focuses on additive valuations, and hence, their techniques do not apply to our matching setting. Several recent papers are concerned with group fairness in online matching [Ma et al., 2020, Sankar et al., 2021]. Ma et al. [2020] studied a stochastic setting wherein the *agents* arrive online (as opposed to the items in our model), following an independent Poisson process with known homogeneous rate; the objective is to maximize the

minimum ratio of the number of agents served to the number of agents in each group. Sankar et al. [2021] studied an online matching problem where the items arrive online. Here, the *items* are grouped into classes (as opposed to the agents in our model), and each agent specifies capacity constraints, which they referred to as *group fairness constraints*, restricting the number of items from each class that can be assigned to the agent. Due to these crucial differences between their models and ours, their techniques and results do not overlap with ours.

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 a set of items $S \subseteq M$, let $\mathbf{e}^S \in \{0, 1\}^M$ denote the incidence vector such that $\mathbf{e}_o^S = 1$ if $o \in S$ and $\mathbf{e}_o^S = 0$ otherwise; we may write $V_i^*(\mathbf{e}^S)$ 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

Our goal is to ensure that items are matched to agents in a manner that is fair to agents belonging to different classes. To that end, we consider classical fairness notions from the fair division literature, such as envy-freeness [Foley, 1967, George and Marvin, 1958], proportionality [Steinhaus, 1948], and maximin share guarantee [Budish, 2011], which are typically used to ensure fairness between individual agents. We 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. For example, when one desirable item has to be allocated among two classes, the class which does not receive the item necessarily envies the other class which receives it. 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 to some readers.

(Approximate) class proportionality and maximin share fairness. Another classical fairness concept is *proportionality*. In the traditional fair division model where agent valuations are additive and there is no limit to how many items can be assigned to an agent, proportionality is typically stated as requiring that each agent receive 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 receive at least the maximum value she can receive from the worst bundle among all fractional partitions of the items into n bundles. While these two versions are equivalent under additive valuations, they are significantly different under non-additive valuations. For subadditive valuations (like our optimistic valuations), the latter version is stronger. Further, the latter version continues to imply its indivisible counterpart, called maximin share fairness, whereas the former version no longer implies it in our model. For these reasons, we use the latter version as the appropriate definition of proportionality in our model.

Since we are interested in fairness at the class level, 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. This naturally adopts the well-studied notion of maximin share fairness to our setting. 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).

For fair division with additive valuations, Segal-Halevi and Suksompong [2019] proved that, subject to allocating every item, EF1 is equivalent to MMS. In contrast, in our model neither implies even an approximation of the other (see Appendix B.2).

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. Non-wasteful integral matchings are also known as *maximal* matchings.

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 a likes o (i.e., $(a, o) \in E$), 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 $\text{usw}(X) \geq \alpha \cdot \text{usw}(X^*)$ for all divisible (resp., indivisible) matchings X^* . When $\alpha = 1$, we refer to X as the USW-optimal matching. Note that the benchmarks for the divisible and indivisible cases are identical as the indivisible matching with the highest USW also has the highest USW among all divisible matchings.

The following is a known relation between maximal (non-wasteful) and maximum matchings in both divisible and indivisible cases. We provide a proof in the appendix for completeness.

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 edges indicate matching. Figure 2a shows an empty matching, which is class envy-free (CEF) but wasteful. Subject to non-wastefulness (NW), CEF is impossible to achieve, but CEF1 can be achieved (Figure 2b). Finally, Figure 2c shows 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 .

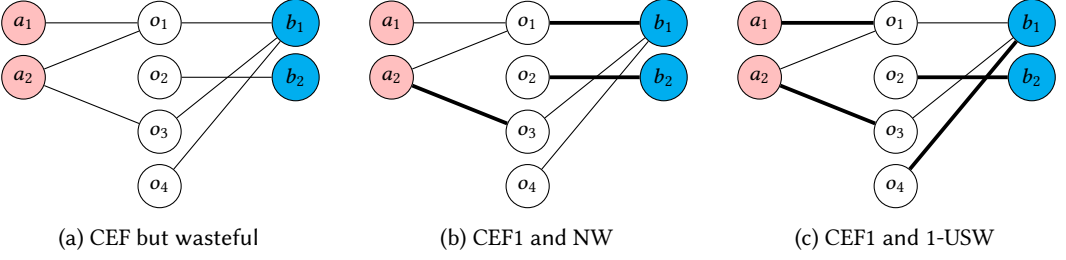


Fig. 2. Class envy-freeness (CEF), non-wastefulness (NW), and utilitarian social welfare approximation (USW).

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 both algorithms which set these values *deterministically* and ones which set them in a *randomized* fashion (but must fix them before the next item arrives). For randomized algorithms, we seek the desired guarantees in expectation.

For the algorithms we design in this paper, 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.

Definition 8. For $\alpha \in (0, 1]$, a randomized online algorithm for matching indivisible items is

- α -CEF if, when all items have arrived, it produces a matching X such that for every pair of classes $i, j \in [k]$, $\mathbb{E}[V_i(X)] \geq \alpha \cdot \mathbb{E}[V_j^*(Y_j(X))]$;
- α -CPROP if, when all items have arrived, it produces a matching X such that for every class $i \in [k]$, $\mathbb{E}[V_i(X)] \geq \alpha \cdot \text{prop}_i$; and
- α -USW if, when all items have arrived, it produces a matching X such that $\mathbb{E}[\text{usw}(X)] \geq \alpha \cdot \text{usw}(X^*)$, where $\mathbb{E}[\text{usw}(X)] = \sum_{a \in N} \sum_{o \in M: (a,o) \in E} \mathbb{E}[x_{a,o}]$ and X^* is a matching with the highest utilitarian social welfare.

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. We provide a formal proof in the appendix.

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 DETERMINISTIC ALGORITHMS FOR 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.

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

```

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.³

Example 1 illustrated that CEF1 and non-wastefulness are incompatible in the online setting.⁴ 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, 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 at the end.

³In fact, discarding all items—an empty matching—is vacuously class envy-free.

⁴In Appendix B.1, 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*.

NW. 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.

USW. 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 j . 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 $t = V_i(X) < (1/2) \cdot \text{mms}_i$. 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 $O \setminus \bigcup_{h \in [k]} Y_h$. 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 on its own by throwing away every item.

The constructions are based on creating instances in which a subset of agents in one class get 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 already 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. For completeness, consider the following simple instance.

There are two items, o_1 and o_2 , arriving in the increasing order of their indices. There is a single class containing two agents. Item o_1 is liked by both agents. The algorithm matches it to one of the two agents. Item o_2 then arrives, and is liked only by the agent who did not receive item o_1 . The optimal utilitarian social welfare is 2, but that of the algorithm is only 1. \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; this is precisely what we achieve in Section 5.

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-wasteful⁵ 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 DETERMINISTIC ALGORITHMS FOR DIVISIBLE ITEMS

We now turn our attention to deterministic 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 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.⁶ 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 7 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 item 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.

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   Define the demand of each class  $i \in [k]$  as  $d_{i,o} = \sum_{a \in N_{i,o}} (1 - \sum_{o' \in M} x_{a,o'})$ 
6   Find the largest  $\beta_o \leq 1$  satisfying  $\sum_{i \in [k]} \min\{\beta_o, d_{i,o}\} \leq 1$ 
7   Set  $y_{i,o} = \min\{\beta_o, d_{i,o}\}$  for each  $i \in [k]$ 
8   for  $i = 1$  to  $k$  do
9     /*individual-phase*/
10    Let  $N_{i,o}$  denote the set of neighbours of item  $o$  in class  $i$ , i.e.,  $N_{i,o} = \{a \in N_i : (a, o) \in E\}$ 
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  $a \in N_{i,o}$ 

```

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.*

PROOF. We prove that EQUAL-FILLING satisfies each of the desirable properties.

NW. Non-wastefulness follows by the algorithm's definition.

⁵Seeking the weaker requirement makes sense only with non-wastefulness since the empty matching vacuously satisfies it.

⁶We 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.

$1/2$ -USW. This is implied by non-wastefulness (Proposition 1).

$(1 - 1/e)$ -CEF. Consider two arbitrary classes i and j . We want to prove that class i 's value for its matching is at least $1 - 1/e$ times its optimistic value for class j 's matching, i.e., $V_i(X) \geq (1 - 1/e) \cdot V_i^*(Y_j)$.

For $\theta \in [0, 1]$, let $f(\theta)$ denote the number of agents in class i who have value ("water level") at least θ under X . Let $N_i(\theta)$ be the set of these $f(\theta)$ agents and $\bar{N}_i(\theta) = N_i \setminus N_i(\theta)$. One can check that for any $\theta \in [0, 1]$, $\int_0^\theta f(z) dz = \sum_{a \in N_i} \min(\theta, \sum_{o \in M} x_{a,o})$.

Let us now rewrite both $V_i(X)$ and $V_i^*(Y_j)$ in terms of $f(y)$. Plugging in $\theta = 1$ above, we see that the total value of the agents in class i is given by

$$V_i(X) = \int_0^1 f(z) dz.$$

Next, fix an arbitrary $\theta \in (0, 1]$. In order to upper bound $V_i^*(Y_j)$, we consider the value derived from Y_j by the agents in $N_i(\theta)$ and those in $\bar{N}_i(\theta)$.

Since agents in $\bar{N}_i(\theta)$ remain unsaturated till the end, for every item o liked by any such agent, the fraction $y_{i,o}$ of the item given to class i must be at least as much as the fraction $y_{j,o}$ of it given to class j . Further, the portion given to class i must be assigned to agents who, at the time of the assignment, had value less than θ . Hence, the total fraction of items given to class j that are liked by at least one agent in $\bar{N}_i(\theta)$, which is an upper bound on the contribution of the agents in $\bar{N}_i(\theta)$ to $V_i^*(Y_j)$, is at most $\int_0^\theta f(z) dz$. Note that the $f(\theta)$ agents in $N_i(\theta)$ contribute at most 1 each to $V_i^*(Y_j)$. Combining these observations, the optimistic value of class i for the items assigned to class j satisfies

$$V_i^*(Y_j) \leq \int_0^\theta f(z) dz + f(\theta), \quad \forall 0 < \theta \leq 1.$$

Multiplying the above inequality by $e^{\theta-1}$ and integrating over $\theta \in (0, 1]$, we get:

$$\begin{aligned} \left(1 - \frac{1}{e}\right) V_i^*(Y_j) &= \int_{\theta=0}^1 e^{\theta-1} V_i^*(Y_j) d\theta \\ &\leq \int_{\theta=0}^1 e^{\theta-1} \left(\int_{z=0}^\theta f(z) dz + f(\theta) \right) d\theta \\ &= \int_{z=0}^1 f(z) \left(\int_{\theta=z}^1 e^{\theta-1} d\theta \right) dz + \int_{\theta=0}^1 e^{\theta-1} f(\theta) d\theta \\ &= \int_{z=0}^1 (1 - e^{z-1}) f(z) dz + \int_{z=0}^1 e^{z-1} f(z) dz \\ &= \int_{z=0}^1 f(z) dz = V_i(X), \end{aligned}$$

where the third transition follows from breaking the integral over the two terms and exchanging the order of integrals in the first part; and during the fourth transition, we rename the index from θ to z in the second part.

$(1 - 1/e)$ -CPROP. Consider an arbitrary class i . We want to prove that class i 's value for the matching is at least $1 - 1/e$ times its proportional share, i.e., $V_i(X) \geq (1 - 1/e) \cdot \text{prop}_i$. Consider an arbitrary divisible partition of the items \tilde{Y} , consisting of non-negative vectors $\tilde{Y}_i = (\tilde{y}_{i,o})_{o \in M}$ for $i \in [k]$ satisfying $\sum_{i \in [k]} \tilde{y}_{i,o} = 1$ for each $o \in M$. It suffices to prove that:

$$k \cdot V_i(X) \geq \left(1 - \frac{1}{e}\right) \cdot \sum_{j \in [k]} V_i^*(\tilde{Y}_j).$$

Recall that $f(\theta)$ denotes the number of agents in class i who have value at least θ under X , $N_i(\theta)$ is the set of these $f(\theta)$ agents, and $\bar{N}_i(\theta) = N_i \setminus N_i(\theta)$. Fix an arbitrary $\theta \in (0, 1]$.

Since the agents in $\bar{N}_i(\theta)$ remain unsaturated till the end, for each item o liked by at least one such agent, the algorithm gives $y_{i,o} \geq 1/k$ fraction of the item to class i (but not necessarily to the agents in $\bar{N}_i(\theta)$). Further, as argued above, this portion of the item must be assigned to the agents in the class who, at the time of the assignment, have value less than θ . Hence, the total number of items liked by at least one agent in $\bar{N}_i(\theta)$, which is an upper bound on the contribution of these agents to $\sum_{j \in [k]} V_i^*(\tilde{Y}_j)$, is at most $k \int_0^\theta f(z) dz$.

Also, each of $f(\theta)$ many agents in $N_i(\theta)$ can contribute a value of at most 1 to $V_i^*(\tilde{Y}_j)$ for each $j \in [k]$. Hence, the total contribution of these agents to $\sum_{j \in [k]} V_i^*(\tilde{Y}_j)$ is at most $k \cdot f(\theta)$.

Combining the two observations, we get that

$$\sum_{j \in [k]} V_i^*(\tilde{Y}_j) \leq k \cdot \left(\int_0^\theta f(z) dz + f(\theta) \right), \quad \forall 0 < \theta \leq 1.$$

Multiplying the inequality by $e^{\theta-1}$, integrating over $\theta \in [0, 1]$, and following the same steps as in the $(1 - 1/e)$ -CEF proof above, we have:

$$\left(1 - \frac{1}{e}\right) \cdot \sum_{j \in [k]} V_i^*(\tilde{Y}_j) \leq k \cdot \int_0^1 f(z) dz = k \cdot V_i(X),$$

as needed. \square

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

PROOF. We argue each impossibility separately.

CEF and NW. Consider any deterministic online algorithm that satisfies non-wastefulness. Consider an instance that consists of two classes, $N_1 = \{a_1, a_2, a_3\}$ and $N_2 = \{b_1, b_2, b_3\}$, and four items o_1, o_2, o_3, o_4 arriving in that order. We denote by X the matching that will be produced by the algorithm on this instance.

Agents a_1, a_2, b_1 , and b_2 like the first two items o_1 and o_2 . By non-wastefulness, the algorithm must fully divide o_1 and o_2 between $\{a_1, a_2, b_1, b_2\}$. Without loss of generality, suppose that the total fraction of these items assigned to class N_1 is at least the total fraction assigned to class N_2 , i.e., $\sum_{a \in N_1} \sum_{o \in \{o_1, o_2\}} x_{a,o} \geq \sum_{b \in N_2} \sum_{o \in \{o_1, o_2\}} x_{b,o}$. Further, we assume, without loss of generality, that agent b_1 obtains at least as much total fraction of these items as agent b_2 , i.e., $\sum_{o \in \{o_1, o_2\}} x_{b_1,o} \geq \sum_{o \in \{o_1, o_2\}} x_{b_2,o}$. Finally, all agents of class N_1 as well as agent b_1 like the remaining two items o_3 and o_4 ; agents b_2 and b_3 do not like them. We will prove that $V_2(X) \leq (3/4) \cdot V_2^*(Y_1)$.

First, we show that $V_2(X) \leq 3/2$. Observe that the value derived by b_2 under X is at most $1/2$. This holds because the total fraction of o_1 and o_2 assigned to b_2 is at most $1/2$ by the assumptions above,

and the agent does not like items o_3 and o_4 . Further, agent b_3 does not like any of the items. Thus, the total value class N_2 can achieve under X is $V_2(X) \leq 1 + 1/2 = 3/2$.

Next, we show that $V_2^*(Y_1) \geq 2$. Note that N_1 must receive a total fraction of at least 1 from each of $\{o_1, o_2\}$ and $\{o_3, o_4\}$. Since b_2 likes every item in $\{o_1, o_2\}$ and b_1 likes every item in $\{o_3, o_4\}$, class N_2 can optimistically derive a total value of at least 2 by assigning Y_{1,o_1} and Y_{1,o_2} fractions of o_1 and o_2 to b_2 (capped by 1), and Y_{1,o_3} and Y_{1,o_4} fractions of o_3 and o_4 to b_1 (capped by 1).

This shows that the algorithm does not achieve α -CEF for any $\alpha > 3/4$.

USW. Note that the utilitarian social welfare is simply the size of the (divisible) matching, which is independent of the class information. Hence, the $1 - 1/e$ upper bound on USW follows from the classical $1 - 1/e$ upper bound on the competitive ratio of any online divisible matching algorithm; see, e.g., the work of Kalyanasundaram and Pruhs [2000].

CPROP. Consider an instance of a single class. In this case, the proportional share of the class coincides with the value $usw(X^*)$ of a USW-optimal matching X^* . Thus, the $1 - 1/e$ upper bound on CPROP approximation follows from the $1 - 1/e$ upper bound on USW approximation. \square

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 based on the divisible matchings of the allocated items (instead of all items). In Proposition 7 in Appendix C, 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 immediately 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 Proposition 7 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 RANDOMIZED ALGORITHMS FOR INDIVISIBLE ITEMS

Recall from Section 3 that for indivisible items, no deterministic online algorithm can achieve α -CMMS for any $\alpha > 1/2$. When moving to randomized algorithms, one can naturally hope to approximate CPROP instead of CMMS because the value to a class is evaluated in expectation. However, a priori it is not clear whether a randomized algorithm can achieve α -CPROP for any $\alpha > 1/2$.

By applying a recently introduced rounding technique, called *Online Correlated Selection* (OCS) [Fahrbach et al., 2020], to the divisible matching given by EQUAL-FILLING (Algorithm 2), we are able to design a randomized algorithm for indivisible items that achieves 0.593-CPROP.

We start by introducing a recent result about OCS that forms the backbone of our approach.

Lemma 1 (c.f., Gao et al. 2021). *There is a polynomial-time online algorithm which works as follows. In each step, it takes as input a non-negative vector $(\tilde{x}_{a,o})_{a \in N}$ for some $o \in M$ satisfying $\sum_{a \in N} \tilde{x}_{a,o} \leq 1$ and selects an agent a with positive $\tilde{x}_{a,o}$. Further, by the end, each agent a is selected at least once with*

probability at least:

$$p(\tilde{x}_a) = 1 - \exp\left(-\tilde{x}_a - \frac{1}{2} \cdot \tilde{x}_a^2 - \frac{4-2\sqrt{3}}{3} \cdot \tilde{x}_a^3\right),$$

where $\tilde{x}_a = \sum_{o \in M} \tilde{x}_{a,o}$.

Technically, such an algorithm is called (multi-way) semi-OCS instead of OCS. But the nomenclature is unimportant for our application, so we will call it OCS for brevity, and refer interested readers to the works of Fahrbach et al. [2020] and Gao et al. [2021] for a detailed comparison.

How good is the guarantee in Lemma 1? For comparison, consider the simpler independent randomized rounding algorithm, which, upon receiving the vector $(\tilde{x}_{a,o})_{a \in N}$, selects each agent a with probability $\tilde{x}_{a,o}$, independently of the rounding outcomes in the previous steps. By the end, each agent a is selected at least once with probability $1 - \prod_{o \in M} (1 - \tilde{x}_{a,o}) \geq 1 - \exp(-\sum_{o \in M} \tilde{x}_{a,o}) = 1 - \exp(-\tilde{x}_a)$. Readers can verify that using this weaker bound in the proof of Theorem 5 only yields $1/2$ -CPROP. The improved guarantee in Lemma 1 is critical for achieving an approximation better than $1/2$.

Our algorithm, EQUAL-FILLING-OCS (presented as Algorithm 3), runs a variant of EQUAL-FILLING in the background to get a guiding divisible matching $\tilde{X} = (\tilde{x}_{a,o})_{a \in N, o \in M}$. The only difference is that unlike EQUAL-FILLING, this variant does not cap the value (total fraction of all items) assigned to an agent at 1. This is because the algorithm will perform rounding to compute an indivisible matching, and by Lemma 1, the probability that an agent a is matched depends on the value \tilde{x}_a of the agent in the divisible matching in such a manner that even reaching a value of 1 would not guarantee being matched with certainty.

Upon receiving a new item o , the algorithm first continues running this variant of EQUAL-FILLING to obtain the guiding division $(\tilde{x}_{a,o})_{a \in N}$ (Lines 5-12), and then lets OCS select an agent a^* accordingly (Line 14). If the selected agent a^* is not yet matched, the algorithm matches item o to this agent. If a^* is already matched, the algorithm matches item o to an arbitrary unmatched agent who likes it, and discards the item if there is no such agent (Line 15).

ALGORITHM 3: EQUAL-FILLING-OCS

- 1 Initialize an empty indivisible matching $X = (x_{a,o})_{a \in N, o \in M}$
 - 2 Initialize an empty divisible matching $\tilde{X} = (\tilde{x}_{a,o})_{a \in N, o \in M}$
 - 3 Maintain a class-level divisible matching $\tilde{Y} = (\tilde{y}_{i,o} = 0)_{i \in [k], o \in M}$ such that $y_{i,o} = \sum_{a \in N_i} \tilde{x}_{a,o}$
 - 4 **when** item $o \in M$ arrives **do**
 - 5 **/*class-phase divisible matching*/**
 - 6 For each class i , let $N_{i,o}$ be the set of agents in class i who like item o
 - 7 Let k_o be the number of classes i such that $N_{i,o} \neq \emptyset$
 - 8 Let $\tilde{y}_{i,o} = \frac{1}{k_o}$ for each of these k_o classes
 - 9 **/*individual-phase divisible matching*/**
 - 10 **for** each class i with $y_{i,o} > 0$ **do**
 - 11 Find γ_o such that $\sum_{a \in N_{i,o}} \max(\gamma_o - \tilde{x}_{a,o}, 0) = \tilde{y}_{i,o}$
 - 12 Let $\tilde{x}_{a,o} = \max(\gamma_o - \tilde{x}_{a,o}, 0)$ for all $a \in N_{i,o}$
 - 13 **/*indivisible matching rounded by OCS*/**
 - 14 Send $(\tilde{x}_{a,o})_{a \in N}$ to the OCS in Lemma 1 and let it select an agent a^*
 - 15 Match o to a^* if a^* is not yet matched, and to an arbitrary unmatched neighbor (if any) otherwise
-

Theorem 5. *For randomized matching of indivisible items, EQUAL-FILLING-OCS (Algorithm 3) satisfies non-wastefulness, 0.593-CPROP, and $1/2$ -USW.*

PROOF. Non-wastefulness is clear from Line 15 of Algorithm 3. Proposition 1 implies $1/2$ -USW. Hence, we focus on the interesting 0.593-CPROP guarantee.

Fix an arbitrary class i . The first part of the analysis bounds the proportional value of class i using the guiding divisible matching \tilde{X} . This part is almost verbatim to its counterpart in the proof of Theorem 3, except we do not bound the value threshold θ by 1. We include this part to be self-contained.

For $\theta \geq 0$, let $f(\theta)$ denote the number of agents in class i who have value at least θ under \tilde{X} . Let $N_i(\theta)$ denote the set of these $f(\theta)$ agents, and let $\bar{N}_i(\theta) = N_i \setminus N_i(\theta)$.

Fix any $\theta > 0$. For each item o liked by at least one agent in $\bar{N}_i(\theta)$, Algorithm 3 assigns a fraction $\bar{y}_{i,o} \geq 1/k_o$ to class i in the guiding divisible matching (but not necessarily to the agents in $\bar{N}_i(\theta)$). Further, any agent in N_i receiving a positive share of item o must have value less than θ right after receiving it. Hence, the total number of items liked by at least one agent in $\bar{N}_i(\theta)$ is at most $k \int_0^\theta f(z) dz$.

On the other hand, the total value that agents in $N_i(\theta)$ can obtain from any set of items is at most $f(\theta)$ (at most 1 per agent).

Therefore, for any divisible partition of the items, denoted by non-negative vectors $\hat{Y}_i = (\hat{y}_{i,o})_{o \in M}$ for $i \in [k]$ such that $\sum_{i \in [k]} \hat{Y}_{i,o} = 1$ for each $o \in M$, we have:

$$\sum_{j \in [k]} V_i^*(\hat{Y}_j) \leq k \cdot \left(\int_0^\theta f(z) dz + f(\theta) \right), \quad \forall \theta > 0.$$

This implies that the proportional share of i is bounded by:

$$\text{prop}_i \leq \int_0^\theta f(z) dz + f(\theta), \quad \forall \theta > 0. \quad (1)$$

Next, we lower bound the expected value of class i for the randomized indivisible matching X . OCS ensures that for each agent a in class i , its probability of being matched is at least $p(\bar{x}_a)$. Hence, the expected value of class i for X is:

$$\begin{aligned} \mathbb{E}[V_i(X)] &\geq \sum_{a \in N_i} p(\bar{x}_a) && \text{(Lemma 1)} \\ &= - \int_0^\infty p(\theta) df(\theta) && \text{(definition of } f(\theta)) \\ &= \int_0^\infty p'(\theta) f(\theta) d\theta. && \text{(integration by parts, } p(0) = f(\infty) = 0) \end{aligned}$$

Multiplying inequality (1) by non-negative coefficients $c(\theta)$ (to be determined later), and integrating over $\theta > 0$ gives that:

$$\begin{aligned} \text{prop}_i \cdot \int_0^\infty c(\theta) d\theta &\leq \int_0^\infty c(\theta) \left(\int_0^\theta f(z) dz + f(\theta) \right) d\theta \\ &= \int_0^\infty c(\theta) \int_0^\theta f(z) dz d\theta + \int_0^\infty c(\theta) f(\theta) d\theta \\ &= \int_0^\infty \left(\int_z^\infty c(\theta) d\theta + c(z) \right) f(z) dz, \end{aligned}$$

where, during the last transition, we exchange the order of integrals in the first part and change the index from θ to z in the second part.

We choose $c(\theta) = -e^\theta \int_\theta^\infty p''(y)e^{-y} dy$, so that $\int_z^\infty c(\theta) d\theta + c(z) = p'(z)$ for all $z > 0$. Hence, we get that:

$$\text{prop}_i \cdot \int_0^\infty c(\theta) d\theta \leq \int_0^\infty p'(z)f(z) dz \leq \mathbb{E} V_i(X).$$

The theorem then follows by numerically calculating the integral:

$$\int_0^\infty c(\theta) d\theta \approx 0.5936 > 0.593.$$

This concludes the proof of the theorem. \square

In Appendix D, we briefly discuss other randomized algorithms and their obstacles in achieving better than $1/2$ approximation to CPROP. We also present a randomized algorithm based on the classical RANKING algorithm, which achieves $(1 - 1/e)$ -CEF. While it achieves this guarantee non-vacuously (i.e., it does not simply return the empty matching), it still violates non-wastefulness. It would be interesting to analyze its efficiency.

6 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?

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.

All of these fall under the umbrella of online fair allocation of *private goods*, which is a literature still in its infancy with many exciting research directions in sight. Studying its counterpart, online fair allocation of *public goods*, is another worthy goal, which may bring its own set of challenges.

REFERENCES

- Gagan Aggarwal, Gagan Goel, Chinmay Karande, and Aranyak Mehta. 2011. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1253–1264.
- Martin Aleksandrov, Haris Aziz, Serge Gaspers, and Toby Walsh. 2015. Online fair division: Analysing a food bank problem. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*. 2540–2546.
- Yossi Azar and Yossi Richter. 2005. Management of multi-queue switches in QoS networks. *Algorithmica* 43, 1 (2005), 81–96.
- Moshe Babaioff, Tomer Ezra, and Uriel Feige. 2021. Fair and truthful mechanisms for dichotomous valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5119–5126.
- Siddhartha Banerjee and Ramesh Johari. 2019. Ride sharing. In *Sharing Economy*. Springer, 73–97.
- Siddharth Barman and Paritosh Verma. 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.
- 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.
- Nawal Benabbou, Ayumi Igarashi, Mithun Chakraborty, and Yair Zick. 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.

- Gerdus Benade, Aleksandr M. Kazachkov, Ariel D. Procaccia, and Christos-Alexandros Psomas. 2018. How to make envy vanish over time. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. 593–610.
- Guy Blanc and Moses Charikar. 2021. Multiway online correlated selection. In *Proceedings of the 62nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- Steven J Brams and Philip D Straffin. 1979. Prisoner’s dilemma and professional sports drafts. *Amer. Math. Monthly* (1979), 80–88.
- Steven J Brams and Alan D Taylor. 2000. *The Win-Win solution: Guaranteeing fair shares to everybody*. WW Norton & Company.
- Eric Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- Eric Budish and Estelle Cantillon. 2012. The multi-unit assignment problem: Theory and evidence from course allocation at Harvard. *The American Economic Review* 102, 5 (2012), 2237–71.
- Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2016. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*. ACM, 305–322.
- Marc Demange and Tinaz Ekim. 2008. Minimum maximal matching is NP-hard in regular bipartite graphs. In *Proceedings of the 4th International Conference on Theory and Applications of Models of Computation (TAMC)*. Springer, 364–374.
- Edmund Eisenberg and David Gale. 1959. Consensus of subjective probabilities: the Pari-Mutuel method. *The Annals of Mathematical Statistics* 30, 1 (1959), 165–168.
- Matthew Fahrbach, Zhiyi Huang, Runzhou Tao, and Morteza Zadimoghaddam. 2020. Edge-weighted online bipartite matching. In *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE, 412–423.
- Jon Feldman, Nitish Korula, Vahab Mirrokni, Shanmugavelayutham Muthukrishnan, and Martin Pál. 2009. Online ad assignment with free disposal. In *Proceedings of the 5th International Workshop on Internet and Network Economics (WINE)*. Springer, 374–385.
- Duncan K. Foley. 1967. Resource allocation and the public sector. *Yale Economic Essays* 7 (1967), 45–98.
- Ruiquan Gao, Zhongtian He, Zhiyi Huang, Zipei Nie, Bijun Yuan, and Yan Zhong. 2021. Improved online correlated selection. In *Proceedings of the 62nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- Jugal Garg and Setareh Taki. 2020. An improved approximation algorithm for maximin shares. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 379–380.
- Gamow George and Stern Marvin. 1958. Puzzle-math.
- Mohammad Ghodsi, MohammadTaghi HajiAghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2018. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. 539–556.
- Artur Gorokh, Siddhartha Banerjee, Billy Jin, and Vasilis Gkatzelis. 2022. Online Nash social welfare via promised utilities. In *Proceedings of the 33rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Forthcoming.
- Hadi Hosseini and Andrew Searns. 2021. Guaranteeing maximin shares: Some agents left behind. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*. 238–244.
- Hadi Hosseini, Andrew Searns, and Erel Segal-Halevi. 2021. Ordinal maximin share approximation for goods. arXiv:2109.01925 [cs.GT]
- Zhiyi Huang, Ning Kang, Zhihao Gavin Tang, Xiaowei Wu, Yuhao Zhang, and Xue Zhu. 2020. Fully online matching. *Journal of the ACM (JACM)* 67, 3 (2020), 1–25.
- Zhiyi Huang and Xinkai Shu. 2021. Online stochastic matching, poisson arrivals, and the natural linear program. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 682–693.
- Zhiyi Huang, Zhihao Gavin Tang, Xiaowei Wu, and Yuhao Zhang. 2019. Online vertex-weighted bipartite matching: Beating $1-1/e$ with random arrivals. *ACM Transactions on Algorithms (TALG)* 15, 3 (2019), 1–15.
- Billy Jin and David P Williamson. 2020. Improved analysis of ranking for online vertex-weighted bipartite matching. *arXiv preprint arXiv:2007.12823* (2020).
- Bala Kalyanasundaram and Kirk R Pruhs. 2000. An optimal deterministic algorithm for online b -matching. *Theoretical Computer Science* 233, 1-2 (2000), 319–325.
- Chinmay Karande, Aranyak Mehta, and Pushkar Tripathi. 2011. Online bipartite matching with unknown distributions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC)*. 587–596.
- Richard M Karp, Umesh V Vazirani, and Vijay V Vazirani. 1990. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*. ACM, 352–358.
- Bernhard Korte and Jens Vygen. 2006. *Combinatorial Optimization: Theory and Algorithms*. Springer-Verlag Berlin Heidelberg.
- David Kurokawa, Ariel D. Procaccia, and Junxing Wang. 2018. Fair enough: Guaranteeing Aapproximate maximin shares. *J. ACM* 65, 2, Article 8 (feb 2018), 27 pages.

- Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See, Ritesh Noothigattu, Siheon Lee, Alexandros Psomas, et al. 2019. WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–35.
- 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*. 125–131.
- Will Ma, Pan Xu, and Yifan Xu. 2020. Group-level fairness maximization in online bipartite matching. *CoRR* abs/2011.13908 (2020). arXiv:2011.13908 <https://arxiv.org/abs/2011.13908>
- Mohammad Mahdian and Qiqi Yan. 2011. Online bipartite matching with random arrivals: An approach based on strongly factor-revealing lps. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*. 597–606.
- Aranyak Mehta. 2013. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science* 8, 4 (2013), 265–368.
- Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. 2007. Adwords and generalized online matching. *Journal of the ACM (JACM)* 54, 5 (2007), 22–es.
- Hervé Moulin. 2019. Fair division in the Internet age. *Annual Review of Economics* 11, 1 (2019), 407–441.
- Renato Paes Leme. 2017. Gross substitutability: An algorithmic survey. *Games and Economic Behavior* 106 (2017), 294 – 316.
- Govind S. Sankar, Anand Louis, Meghana Nasre, and Prajakta Nimbhorkar. 2021. Matchings with group fairness constraints: Online and offline algorithms. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, Zhi-Hua Zhou (Ed.). 377–383.
- Erel Segal-Halevi and Warut Suksompong. 2019. Democratic fair allocation of indivisible goods. *Artificial Intelligence* 277 (2019), 103167.
- Yongho Shin and Hyung-Chan An. 2021. Making three out of two: Three-way online correlated selection. In *Proceedings of the 32nd International Symposium on Algorithms and Computation (ISAAC)*, Vol. 212. Schloss Dagstuhl, 49:1–49:17.
- Hugo Steinhaus. 1948. The problem of fair division. *Econometrica* 16 (1948), 101–104.
- Hal R Varian. 1974. Equity, envy, and efficiency. *Journal of Economic Theory* 9, 1 (1974), 63–91.
- Toby Walsh. 2011. Online cake cutting. In *International Conference on Algorithmic Decision Theory*. Springer, 292–305.
- Mihalis Yannakakis and Fanica Gavril. 1980. Edge dominating sets in graphs. *SIAM J. Appl. Math.* 38, 3 (1980), 364–372.
- David Zeng and Alexandros Psomas. 2020. Fairness-efficiency tradeoffs in dynamic fair division. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 911–912.

APPENDIX

A OMITTED MATERIAL FROM SECTION 2

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

PROOF OF PROPOSITION 1. Let X^* be a matching maximizing the utilitarian social welfare. Without loss of generality, we can pick X^* to be integral. Let X be any non-wasteful (divisible or indivisible) matching. Hence, for every $(a, o) \in E$, we have $\sum_{o' \in M} x_{a,o'} = 1$ or $\sum_{a' \in N} x_{a',o} = 1$. Then, we have

$$\begin{aligned} \text{usw}(X^*) &= \sum_{(a,o):x_{a,o}^*=1} 1 \leq \sum_{(a,o):x_{a,o}^*=1} \left(\sum_{o' \in M} x_{a,o'} + \sum_{a' \in N} x_{a',o} \right) \\ &\leq \sum_{a \in N} \sum_{o' \in M} x_{a,o'} + \sum_{o \in M} \sum_{a' \in N} x_{a',o} = 2 \cdot \text{usw}(X), \end{aligned}$$

where the second transition holds because $x_{a,o}^* = 1$ implies $(a, o) \in E$ and X is non-wasteful, and the third transition holds because X^* is an indivisible matching (i.e., if $x_{a,o}^* = 1$, $x_{a',o'}^* = 1$, and $(a, o) \neq (a', o')$, then $a \neq a'$ and $o \neq o'$). This proves that X is $1/2$ -USW. \square

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.*

PROOF OF PROPOSITION 2. Let us first consider indivisible items. Let A be any deterministic online algorithm that may be wasteful. Consider a non-wasteful version of it, denoted as A' , that works as follows. It runs A in the background and treats A 's output as an advice. Importantly, A keeps its own internal state and is oblivious to the actual matching decisions made by A' . For an item o , suppose that A matches o to agent a . Algorithm A' would follow A 's advice and match o to a if a is not yet matched, and would otherwise match o to any unmatched agent who likes item o .

By definition, A' is non-wasteful. Further, we can prove by induction over the steps that the set of agents matched by A' is a superset of the set of agents matched by A . Since CMMS is a monotone property (i.e., increasing agent values preserves its approximation), A' achieves at least as good an approximation of CMMS as A does.

For divisible items, the same proof works for CPROP, except A' now gives a fraction of o to each agent a that is the minimum of the fraction of o matched to a under the advice given by A and the remaining capacity of a in the current matching maintained by A' . \square

B OMITTED MATERIAL FROM SECTION 3

B.1 Pessimistic class envy-freeness

One may wonder whether relaxing the way each class measures its hypothetical value for a set of items could help alleviating the incompatibility between class envy-freeness and non-wastefulness. We show that even if each class considers a pessimistic value for a set of items (in other words, considers worst-case scenario for matching the items), the clash between envy-freeness and non-wastefulness persists.

Given a vector $\mathbf{y} = (y_o)_{o \in M} \in \{0, 1\}^M$ representing a set of items, the *pessimistic valuation* $V_i^\ominus(\mathbf{y})$ of class i for \mathbf{y} is the value of a *minimum-cardinality maximal matching* between the agents of N_i and the set $\{o \in M \mid y_o = 1\}$. This problem has shown to be NP-hard for graphs with maximum degree 3 and k -regular bipartite graphs for $k \geq 3$ [Demange and Ekim, 2008, Yannakakis and Gavrill, 1980].




Fig. 3. An allocation that is PEF but not CEF. The red group pessimally considers the worst-case matching of items o_1 and o_3 with the value of 1.

We compare the value $V_i(X)$ derived by class i from matching X with class i 's pessimistic valuation for the items matched to another class j , i.e. $V_i^\ominus(Y_j(X))$.

Definition 9 (Pessimal class envy-freeness). A matching X is α -pessimal class envy-free (α -PEF) if for every pair of classes $i, j \in [k]$, $V_i(X) \geq \alpha \cdot V_i^\ominus(Y_j(X))$. When $\alpha = 1$, we simply refer to it as pessimal class envy-freeness (PEF).

Similar to its optimistic counterpart, CEF, a PEF matching may not always exist. Therefore, we consider the following relaxation of PEF for integral matchings.

Definition 10 (Pessimal class envy-freeness up to one item). An integral matching X is α -pessimal class envy-free up to one item (α -PEF1) 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 $V_i(X) \geq \alpha \cdot V_i^\ominus(Y_j(X) \setminus \{o\})$. When $\alpha = 1$, we simply refer to it as class envy-freeness up to one item (PEF1).

It is easy to verify that PEF1 is weaker than CEF1. Intuitively, a class values its matching compared to the items assigned to another class if it has a pessimistic view of the items arrival and matched items, should the items were exchanged. Clearly, a CEF matching is also PEF, and similarly CEF1 implies PEF1.

Example 3. In the example given in Figure 3, there are two classes $N_1 = \{a_1, a_2\}$ and $N_2 = \{b_1, b_2\}$. The bold edges indicate the matched items. This matching is not CEF, since class N_1 envies class N_2 should it able to optimally match items o_1 and o_3 within its members. However, the same matching is PEF because class N_1 considers a pessimal matching of the same items, that is o_1 and o_3 , where item o_1 is matched to a_1 upon its arrival, and thus, o_3 remains unmatched (Since there is no edge from a_2 to o_3).

The following proposition strengthens our previous results on the incompatibility between non-wastefulness and CEF1 by showing that non-wastefulness remains incompatible with a weaker fairness notion of PEF1.

Proposition 4. *No deterministic algorithm for matching indivisible items can guarantee non-wastefulness and PEF1.*

PROOF. Consider the example given in Figure 1. It is easy to verify that the matching is non-wasteful. However, in this scenario the pessimal value of class N_1 for the items assigned to the class N_2 is 3, implying that the matching is not PEF1. \square

B.2 Relationships Between CEF1 and CMMS

Proposition 5 (CEF1+NW \Rightarrow CMMS). *Given an indivisible instance, a CEF1+NW matching does not imply any α -CMMS for any $\alpha > 0$.*

PROOF. We construct an instance for which a α -CEF1+NW matching with $\alpha = 1$ gives only a 0-CMMS approximation.

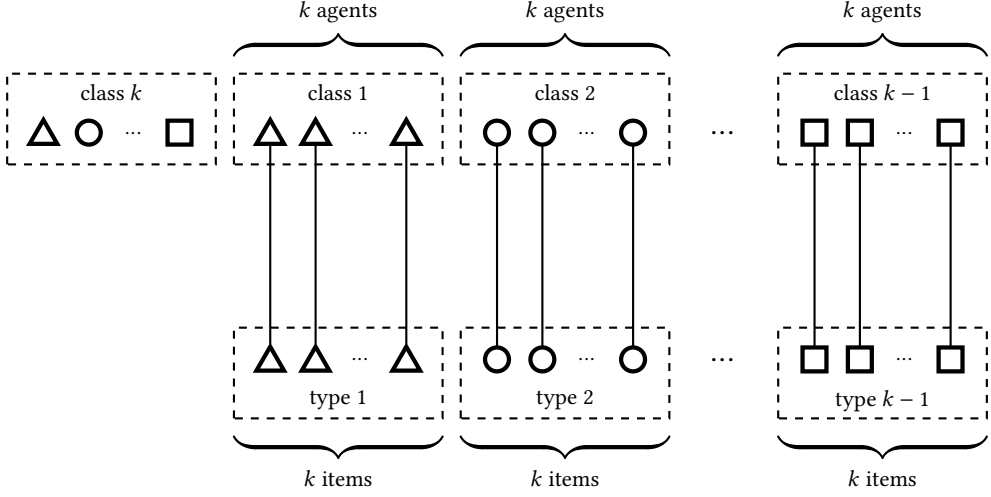


Fig. 4. A CEF1+NW matching that does not imply any approximation for CMMS.

Suppose there are k classes N_1, N_2, \dots, N_k . Each N_i for $i \in [k-1]$ consists of k agents. The last class N_k consists of $k-1$ agents a_1, a_2, \dots, a_{k-1} . There are $k(k-1)$ items that are partitioned into $k-1$ subsets C_1, C_2, \dots, C_{k-1} . For $j \in [k-1]$, C_j consists of k items, $o_{1j}, o_{2j}, \dots, o_{kj}$, each of which is referred to as a *type j* item. For each $j \in [k-1]$, every agent in class N_j likes every item in C_j . For class N_k , each agent a_j for $j \in [k-1]$ likes every item in C_j . For example, agent a_2 likes k items $o_{12}, o_{22}, \dots, o_{k2}$ but does not like none of the other items.

Now, consider a matching X that gives no item to class N_k and matches arbitrarily each of the k items in C_j to one of the k agents in each class N_j for $j \in [k-1]$ (as illustrated in Figure 4). Since each of the $k(k-1)$ items are fully assigned to an agent who likes it, the matching X is clearly non-wasteful. Further, this matching is CEF1. In fact, all classes except N_k receive a perfect matching and are not envious of any other class. Also, for $j \in [k-1]$, there is at most one agent a_j in N_k who likes an item in C_j . Thus, class N_k is not envious for more than one item since $V_k^*(Y_j) \leq 1$ for any $j \in [k-1]$. Thus, the matching is CEF1.

In contrast, consider a partition (L_1, L_2, \dots, L_k) of the items where $L_i = \{o_{i1}, o_{i2}, \dots, o_{ik-1}\}$ for each $i \in [k]$. Observe that for each $i = 1, 2, \dots, k$, each agent a_j in N_k likes exactly one item o_{ij} in L_i , i.e., $L_i \cap C_j = \{o_{ij}\}$ for $j \in [k-1]$. This means that there is a perfect matching of size $k-1$ between N_k and the items of each L_i , yielding $V_1^*(L_i) \geq k-1$ for $i \in [k]$. We thus establish that $\text{mms}_k \geq k-1$. Given that class N_1 's value for X is $V_1(X) = 0$, X provides 0-CMMS approximation, which proves the claim. \square

Proposition 6 (CMMS \Rightarrow CEF1+NW). *Given an indivisible instance, a CMMS matching does not imply α -CEF1 for any $\alpha > 0$.*

PROOF. Consider an instance with k classes each with k agents. There are $k-1$ items liked by every agent in each class. A matching that assigns all $k-1$ of items to a single class, say N_1 , satisfies CMMS. This is because the CMMS value for each class is obtained by partitioning the $k-1$ items into k bundles, yielding $\text{mms}_i = 0$ for $i = 1, 2, \dots, k$. However, this matching is not CEF1 (nor any

α approximation of it for $\alpha > 0$) because every class values the matching assigned to N_1 as $k - 1$ while only receiving 0 valuation. \square

C OMITTED MATERIAL FROM SECTION 4

C.1 Proportionality with respect to allocated items

Our objective of this section is to show that $(1 - 1/e)$ -bound is tight even for CPROP with respect to the allocated items. Formally, we define the *proportional share* of class i with respect to a set S of items as

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

where $\mathcal{X}(S)$ is the set of (divisible) matchings of the set of items S to the set of agents N . For $\alpha \in (0, 1]$, we say that matching X is α -class proportional (α -CPROP) with respect to a set S of items if for every class $i \in [k]$, $V_i(X) \geq \alpha \cdot \text{prop}_i^S$. For $\alpha \in (0, 1]$, a deterministic online algorithm for matching divisible items is α -class proportional (α -CPROP) with respect to the allocated items if when all items have arrived, it produces a matching that is α -class proportional with respect to the items that have been fully assigned by the algorithm.

Proposition 7. *No deterministic algorithm for matching divisible items satisfies α -CPROP with respect to the allocated items for any $\alpha > 1 - 1/e$.*

PROOF. We will prove that no deterministic online algorithm satisfying non-wastefulness can achieve α -CPROP with respect to the allocated items for any $\alpha > 1 - 1/e$. By the proof of Proposition 2, this implies that no deterministic algorithm can guarantee α -CPROP with respect to the allocated items for any $\alpha > 1 - 1/e$.

Take any non-wasteful algorithm for divisible item allocation and consider the following adversarial instance. There are two classes of $3n$ agents each, $N_1 = \{a_1, \dots, a_n, d_1, \dots, d_{2n}\}$ and $N_2 = \{a'_1, \dots, a'_n, d'_1, \dots, d'_{2n}\}$. We call the agents $d_1, \dots, d_{2n}, d'_1, \dots, d'_{2n}$ *dummy agents*. There are $2n$ items, labeled o_i and o'_i for $i \in [n]$.

The construction of the instance works in rounds as follows.

- We start with $t = 1$, $R_1^0 = \{a_1, a_2, \dots, a_n\}$, and $R_2^0 = \{a'_1, a'_2, \dots, a'_n\}$.
- In round t , items o_t arrives, followed immediately by item o'_t . Both these items are liked by agents in R_1^{t-1} and R_2^{t-1} .
- Let $V^t(a)$ denote the value that agent a derives at the end of round t when the algorithm finishes allocating both items. Find the lowest valuation agent in each class. WLOG, say $a_t \in \arg \min_{a \in R_1^{t-1}} V^t(a)$ and $a'_t \in \arg \min_{a' \in R_2^{t-1}} V^t(a')$. Set $R_1^t \leftarrow R_1^{t-1} \setminus \{a_t\}$, $R_2^t \leftarrow R_2^{t-1} \setminus \{a'_t\}$, and $t \leftarrow t + 1$.

We stop this process after the first round t^* such that at the end of that round every agent in $R_1^{t^*}$ and every agent in $R_2^{t^*}$ is fully saturated.

Without loss of generality, assume that at the end of round t^* , the total value of agents in N_1 is at most the total value of agents in N_2 , i.e., $\sum_{a \in N_1} V^{t^*}(a) \leq \sum_{a' \in N_2} V^{t^*}(a')$. For shorthand, let us denote $V^t(A) = \sum_{a \in A} V^t(a)$ for a set of agents A .

Then, the remaining $2(n - t^*)$ items that arrive are liked by agents in $N_2 \cup R_1^{t^*}$. Note that by non-wastefulness and by the fact that N_2 contains $2n$ dummy agents, the $2(n - t^*)$ items are fully assigned to some agent.

We claim the following properties at the end of round t^* .

- The agents $n - t^*$ agents in $R_1^{t^*}$ and the $n - t^*$ agents in $R_2^{t^*}$ are all fully saturated.
- $V^{t^*}(N_1) \leq t^*$, $V^{t^*}(N_2) \geq t^* - 1$.
- $t^* \leq (1 - 1/e) \cdot n$ (in particular, the process will stop after no more than n rounds).

The first claim follows immediately due to the definition of t^* . For the second claim, note that the total value of both classes after t rounds must be at most $2t$ since only $2t$ items have arrived. Also, the total value of both classes after t rounds must be at least $2(t-1)$; this is because the t th round only happens if some agent in $R_1^{t-1} \cup R_2^{t-1}$ was not fully saturated after $t-1$ rounds, and since this agent was part of $R_1^{t'} \cup R_2^{t'}$ for all $t' \leq t-1$, non-wastefulness implies that the algorithm must have assigned the $2(t-1)$ items from the first $t-1$ rounds fully. These two claims, along with the convention that $V^{t^*}(N_1) \leq V^{t^*}(N_2)$ implies the second claim.

Before we prove the third claim, we show why these claims imply the desired bound on the envy ratio. At the end of the algorithm, the total value of class N_1 is at most t^* because of the second claim and the fact that they do not receive any items from the last $2(n-t^*)$ items (as all agents in $R_1^{t^*}$ are saturated after round t^*).

In contrast, the proportional fair share prop_1^S of class N_1 with respect to the allocated items S is at least $n-1$. Note that all the items except for o_{t^*} and o'_{t^*} are fully assigned. Thus, $M \setminus \{o_{t^*}, o'_{t^*}\} \subseteq S$. Further, consider two sets $P_1 = \{o_1, \dots, o_{t^*-1}, o_{t^*+1}, \dots, o_n\}$ and $P_2 = \{o'_1, \dots, o'_{t^*-1}, o'_{t^*+1}, \dots, o'_n\}$. From P_1 , the t^*-1 items $o_1, o_2, \dots, o_{t^*-1}$ can be matched to t^*-1 agents $a_1, a_2, \dots, a_{t^*-1}$ and the remaining $n-t^*$ items can be matched to $n-t^*$ agents in $R_1^{t^*}$. Similarly, from P_2 , t^*-1 items $o'_1, o'_2, \dots, o'_{t^*-1}$ can be matched to t^* agents $a'_1, a'_2, \dots, a'_{t^*-1}$ and the remaining $n-t^*$ items can be matched to $n-t^*$ agents in $R_1^{t^*}$. Thus, $\text{prop}_1^S \geq n-1$. From the third claim, if $V_1(X) \geq \alpha \cdot \text{prop}_1^S$, then $(1-1/e)n \geq \alpha(n-1)$, meaning that $(1-1/e)\frac{n}{n-1} \geq \alpha$.

Finally, we show that $t^* \leq (1-1/e) \cdot n$. To see this, we first show that after t rounds,

$$V^t(N_1 \setminus R_1^t) + V^t(N_2 \setminus R_2^t) \leq \frac{2t}{n} + \frac{2(t-1)}{n-1} + \dots + \frac{2 \cdot 1}{n-t+1}.$$

For the base case, note that after the first round, $V^1(a_1) + V^1(a'_1) \leq 2/n$ follows from the pigeonhole principle. Suppose this claim holds after $t-1$ rounds. Then, after round t , we have

$$V^t(a_t) + V^t(a'_t) \leq \frac{2t - (V^{t-1}(N_1 \setminus R_1^{t-1}) + V^{t-1}(N_2 \setminus R_2^{t-1}))}{n-t+1}.$$

Adding $V^{t-1}(N_1 \setminus R_1^{t-1}) + V^{t-1}(N_2 \setminus R_2^{t-1}) = V^t(N_1 \setminus R_1^t) + V^t(N_2 \setminus R_2^t)$ to both sides, we obtain

$$V^t(N_1 \setminus R_1^t) + V^t(N_2 \setminus R_2^t) \leq \frac{2t}{n-t+1} + \frac{n-t}{n-t+1} \cdot (V^{t-1}(N_1 \setminus R_1^{t-1}) + V^{t-1}(N_2 \setminus R_2^{t-1})).$$

Using the induction hypothesis, we get the desired result. Consider the smallest \hat{t} such that

$$2\hat{t} - 2 - \left(\frac{2\hat{t}}{n} + \frac{2(\hat{t}-1)}{n-1} + \dots + \frac{2 \cdot 1}{n-\hat{t}+1} \right) \geq 2(n-\hat{t}).$$

Note that the process must stop at $t^* \leq \hat{t}$. This is because the total value of both classes after \hat{t} round is at least $2\hat{t} - 2$, but the value to the removed agents is at most the expression in the brackets. Hence, the remaining allocation must have saturated the remaining $2(n-\hat{t})$ agents. After simple algebra, we can see that the left hand side is equal to $2 \cdot (n-\hat{t}) \cdot (H_n - H_{n-\hat{t}}) - 2$. If this is at least $2(n-\hat{t})$, then $H_n - H_{n-\hat{t}} \geq 1 + 1/(n-\hat{t})$. The smallest \hat{t} when this is satisfied is roughly $(1-1/e) \cdot n + o(n)$. \square

C.2 Upper bounds for Algorithm 2

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 $\alpha > 1/2$.

PROOF. The fact that Algorithm 2 cannot achieve α -CPROP for $\alpha > 1 - 1/e$ immediately follows from Theorem 4.

For each of the fairness or efficiency guarantees, we provide an instance for which Algorithm 2 cannot achieve the corresponding bound.

CEF. Consider the following instance with two classes $N_1 = \{a_1, a_2, \dots, a_n\}$ and $N_2 = \{a'_1, a'_2, \dots, a'_{2n}\}$. There are $2n$ items $o_1, o'_1, o_2, o'_2, \dots, o_n, o'_n$. There are n rounds: in round $t \in [n]$, item o_t arrives, followed immediately by item o'_t . Each agent a'_i ($i \in [2n]$) likes every item. Each agent a_i ($i \in [n]$) likes the items o_t and o'_t with $t = 1, 2, \dots, i$; namely, agent a_1 likes the items o_1, o'_1 , agent a_2 likes items o_1, o'_1, o_2, o'_2 , and so on.

Note that since N_2 has $2n$ agents who like all $2n$ items, for each item, there is at least one agent in N_2 who is not saturated and likes that item. Thus, until the agents in N_1 who like o_t and o'_t are fully saturated, the equal-filling algorithm splits the item into halves among the two classes. The algorithm assigns the amount $\frac{1}{2n}$ of $\{o_t, o'_t\}$ to each agent in N_2 . On the other hand, it assigns the amount $\frac{1}{n-(t-1)}$ of o_t and o'_t to each agent i of class N_1 with $i \geq t$; for example, agent a_1 receives $\frac{1}{n}$ of $\{o_1, o'_1\}$; agent a_2 receives $\frac{1}{n}$ of $\{o_1, o'_1\}$ and $\frac{1}{n-1}$ of $\{o_2, o'_2\}$; agent a_3 receives $\frac{1}{n}$ of $\{o_1, o'_1\}$, $\frac{1}{n-1}$ of $\{o_2, o'_2\}$, and $\frac{1}{n-2}$ of $\{o_3, o'_3\}$; and so on.

Let X denote the matching returned by Algorithm 2. We will establish that $V_1(X) \leq (1 - 1/e)V_1^*(Y_2)$. First, it is not difficult to see that under X , class N_2 is assigned to at least 1 for each item set of $\{o_t, o'_t\}$ ($t \in [n]$). Thus, $V_1^*(Y_2) \geq n$. Now, let $t^* = n - \lceil \frac{n}{e} \rceil$. It can be easily checked by the integral test that $\sum_{t=1}^{t^*} \frac{1}{n+1-t}$ is between $1 - \frac{5}{n}$ and 1. Thus, after the algorithm assigns o_{t^*+5}, o'_{t^*+5} , the set N_{1,o_t} becomes empty, i.e., there is no agent in N_1 who is not saturated and likes new items o_t, o'_t for $t > t^* + 5$. Thus, the value $V_1(X)$ derived by class N_1 from X is at most

$$t^* + 5 < (1 - \frac{1}{e})n + 5 \leq (1 - \frac{1}{e})V_1^*(Y_2) + 5,$$

which proves the claim.

USW. Let n be a positive integer. Consider $n + 1$ classes: There are n classes N_j , each of which consists of a single agent c_j for $j = 1, 2, \dots, n$. The last class N_{n+1} consists of n agents $\{a_1, a_2, \dots, a_n\}$. There are $2n$ items: n red items r_1, r_2, \dots, r_n and n blue items b_1, b_2, \dots, b_n . Each red item is liked by every agent. Each blue item b_i is liked by the single agent c_i in N_i . Now the instance admits a perfect matching of size $2n$ that matches every agent c_i for $i \in [n]$ to the blue item b_i and the remaining n agents in N_{n+1} arbitrarily to the remaining n red items.

Now suppose that the items arrive in the order of $r_1, r_2, \dots, r_n, b_1, b_2, \dots, b_n$. For each red item r_i ($i \in [n]$), the equal-filling algorithm assigns an equal amount $\frac{1}{n+1}$ of fractions among the $n + 1$ classes. Thus, after the algorithm matches the last red item r_n , the total amount of fractions each class N_i for $i \in [n+1]$ has received is $\frac{n}{n+1}$. For each blue item b_i ($i \in [n]$), the equal-filling algorithm assigns an amount of $\frac{1}{n+1}$ to the agent c_i in N_i since c_i is the only agent who likes the blue item b_i but has already been saturated up to $\frac{n}{n+1}$. Thus, the utilitarian social welfare of the resulting matching X is given as follows:

$$\sum_{i=1}^n V_i(X) + V_{n+1}(X) = \sum_{i=1}^n 1 + \frac{n}{n+1} = n + \frac{n}{n+1}.$$

This proves the claim. \square

D OMITTED MATERIAL FROM SECTION 5

D.1 Discussion on Other Randomized Algorithms

Readers familiar with the online matching literature may wonder why can't we use the Ranking algorithm of Karp et al. [1990] to decide how to match items *within* each class, and combine it with some fair class-level matching approach. While we believe this is an interesting direction for future research, there is a concrete technical difficulty in analyzing such algorithms. Naturally, the class-level matching must take into account which agents are already matched to previous items. This means that the realization of randomness used by Ranking within some class i will influence what items are allocated to the class!

How about applying Ranking directly, ignoring how agents are partitioned into classes? While this approach circumvents the above challenge, it fails on two classes with lopsided sizes. In the extreme, consider a class with only one agent, and another class with $n \gg 1$ agents, and only one item. The second class will get the item with probability $\frac{n}{n+1}$ while the first class gets it only with probability $\frac{1}{n+1}$.

Finally, we observe that it is necessary to have randomness in both the class-level matching and the individual-level matching, in order to exploit the power of randomized algorithms.

Proposition 8. *If an algorithm assigns deterministically at the class-level, it is at best $\frac{1}{2}$ -CPROP.*

PROOF. Consider two classes $N_1 = \{a_1, a_2, a_3\}$ and $N_2 = \{b_1, b_2, b_3\}$. For $1 \leq i \leq 3$, the i -th item is liked by a_i and b_i . If the algorithm assigns all three items to the same class, it is only 0-CPROP. Otherwise, assume without loss of generality that 2 items go to class 2. Let the next item be only liked by the matched agent in class 1 and the unmatched agent in class 2, as in Figure 1. The algorithm is then at best $\frac{1}{2}$ -CPROP. \square

Proposition 9. *If an algorithm assigns deterministically within each class, it is at best $\frac{1}{2}$ -CPROP.*

PROOF. It becomes apparent when we consider a single class. The proposition then reduces to the fact that deterministic online matching algorithms are at best $\frac{1}{2}$ -competitive. We can extend this hard instance to k classes by making k copies of the class and each item. \square

D.2 Discussion on Randomized Algorithms and CEF

As discussed in the last subsection, if the the class-level matching depends on which agents are already matched, i.e., if it is adaptive to the realization of randomness in the agent-level matching, then the realization of randomness in an online algorithm, e.g., Ranking, within each class would affect what items get assigned to the class. How about using a class-level matching algorithm that is oblivious to the randomness in the agent-level matching? Although such algorithms must violate non-wastefulness in general, we find an algorithm that isn't blatantly wasteful and looks interesting enough to be a stepping stone towards stronger algorithms in future works.

We call this algorithm EQUAL-RANKING. For each item, it randomly assigns the item to a class with at least one agent who likes the item. Within each class, it runs a separate Ranking algorithm to match items to agents therein.

Proposition 10. *Given an online indivisible instance, EQUAL-RANKING guarantees $(1 - 1/e)$ -CEF.*

PROOF. Consider any class i and any other class j . Let $y_i = (y_{io})_{o \in M} \in \{0, 1\}^M$ be the vector that represents the subset of items assigned to i by EQUAL-RANKING at the class-level, regardless of whether such the items are matched to agents successfully. Define y_j similarly. Note that both y_i and y_j are random variables that depend on the class-level random assignments of items. Finally, let $X = (x_{ao})_{a \in N, o \in M} \in \{0, 1\}^{N \times M}$ be the matrix that represents the matching by EQUAL-RANKING.

We seek to prove that:

$$\mathbb{E}[V_i(X)] \geq (1 - 1/e) \mathbb{E}[V_i^*(y_j)] .$$

Conditioned on the subset of items assigned to i , i.e., y_i , the Ranking algorithm ensures that (see, e.g., Karp et al. [1990]):

$$\mathbb{E}[V_i(X) \mid y_i] \geq (1 - 1/e) V_i^*(y_i) .$$

It remains to show that:

$$\mathbb{E}[V_i^*(y_i)] \geq \mathbb{E}[V_i^*(\hat{y}_j)] .$$

Define \hat{y}_j be such that $\hat{y}_{jo} = y_{jo}$ if class i has at least one agent who likes item o , and $\hat{y}_{jo} = 0$ otherwise. By definition $V_i^*(y_j) = V_i^*(\hat{y}_j)$ and therefore it suffices to prove:

$$\mathbb{E}[V_i^*(y_i)] \geq \mathbb{E}[V_i^*(\hat{y}_j)] .$$

Note that for any item o , EQUAL-RANKING ensures that the probability that $y_{io} = 1$ is greater than or equal to the probability that $\hat{y}_{jo} = 0$. Further, the assignment of items at the class-level are independent. Hence we get that random variable y_i stochastically dominates \hat{y}_j . The above inequality now follows by the monotonicity of V_i^* . \square