

Two-Sided Matching Meets Fair Division

Rupert Freeman¹, Evi Micha², Nisarg Shah²

¹University of Virginia

²University of Toronto

FreemanR@darden.virginia.edu, {emicha,nisarg}@cs.toronto.edu

Abstract

We introduce a new model for two-sided matching which allows us to borrow popular fairness notions from the fair division literature such as envy-freeness up to one good and maximin share guarantee. In our model, each agent is matched to multiple agents on the other side over whom she has additive preferences. We demand fairness for each side separately, giving rise to notions such as double envy-freeness up to one match (DEF1) and double maximin share guarantee (DMMS). We show that (a slight strengthening of) DEF1 cannot always be achieved, but in the special case where both sides have identical preferences, the round-robin algorithm with a carefully designed agent ordering achieves it. In contrast, DMMS cannot be achieved even when both sides have identical preferences.

1 Introduction

Consider a group of agents seeking to divide some number of indivisible goods amongst themselves. Each agent has a utility function describing the value that they have for every possible bundle of goods, and each agent may have a different utility function. This is a canonical resource allocation problem that arises in estate division, partnership dissolution, and charitable donations, to name just a few. A central goal is to find an allocation of the goods that is *fair*.

One desirable notion of fairness is *envy-freeness* [Foley, 1967], which requires that no agent prefer another agent’s allocation of goods to her own. This is a compelling definition but, due to the discrete nature of the problem, cannot always be satisfied. Instead, we must consider relaxed versions, with one popular relaxation being *envy-freeness up to one good* (EF1) [Lipton *et al.*, 2004; Budish, 2011], which requires that any pairwise envy can be eliminated by removing a single good from the envied agent’s allocation. An allocation satisfying EF1 always exists for a broad class of agent utility functions [Lipton *et al.*, 2004].

While quite general, the resource allocation model fails to capture some allocation settings that we might be interested in. In particular, it does not allow for the possibility of *two-sided* preferences, in which agents have preferences over “goods,” but also “goods” have preferences over agents.

For instance, when college courses are allocated to students, it is reasonable to assume that students have preferences over the courses they take, and that teachers in charge of courses also have preferences over the students they accept (perhaps measured by prerequisites or GPA).¹ As another example, consider the problem of matching social services to vulnerable individuals², where individuals have preferences over the services they receive, and service providers have preferences over the individuals they serve (perhaps based on demographics, location, or synergy with existing clients).

Allowing for two-sided preferences is immediately reminiscent of the matching literature. In two sided matching, it is generally assumed that each agent has ordinal preferences over the other side, and a matching is sought that is in some sense *stable* to individual or group deviations. It is well-known that stability is closely related to envy-freeness in the sense that a one-to-one matching is stable if and only if it eliminates *justified envy* [Abdulkadiroğlu and Sönmez, 2003], a requirement specifying that any envy that i may feel for j ’s match is “justified” by j ’s match preferring j to i . In many-to-one settings, the notions remain tightly connected, with a stable matching being one that eliminates justified envy and has no waste. Justified envy-freeness has also been studied as in its own right in the many-to-one setting [Wu and Roth, 2018; Yokoi, 2020].

Justified envy, and therefore stability, fundamentally rely on the idea that the less an agent is preferred by the other side, the lower her own entitlement should be. However, in some applications, it may be desirable to provide *equal* entitlements or moral claims to agents regardless of how valued they are by the agents on the other side. For example, instructors may prefer students with high GPA over students with low GPA, but it is not clear that universities should adopt such a policy in their course scheduling (and, in fact, usually do not). Therefore, while stability is a valuable notion in many settings, in this work we consider two-sided preferences while incorporating traditional notions from fair division.

Our contributions. We introduce and study a two-sided resource allocation setting in which we have two groups of

¹Assigning students to courses has been studied before [Budish, 2011; Othman *et al.*, 2010; Budish and Cantillon, 2012], but these papers typically only consider the preferences of the students.

²<http://csse.utoronto.ca/social-needs-marketplace>

agents, where each agent has preferences over agents on the other side. Each agent must be “matched” to a subset of agents on the other side, subject to a maximum degree constraint. Our goal is to find a many-to-many matching that provides fairness to both sets of agents simultaneously. The standard resource allocation setting is a special case of our model in which each good can be matched to at most one agent, each agent can be matched to any number of goods, and goods are indifferent to which agent they are assigned to.

As a natural tradeoff between expressiveness and succinctness, we restrict our attention to additive preferences, in which the utility for being matched to a group of agents is equal to the sum of utilities for being matched to each agent in the group individually. While conceptually simple, additive preferences have led to a rich body of work in fair division. We focus primarily on the case in which all agents on the same side have the same degree constraint, and the total maximum degree on both sides is equal. In this case, it is reasonable to seek a *complete* matching, which saturates the degree constraints of all the agents on both sides.

We begin by considering *double envy-freeness up to one match* (DEF1), requiring that EF1 hold for both sets of agents simultaneously. We show that a complete matching satisfying (a slight strengthening of) DEF1 does not always exist, but in the special case where both sides have identical ordinal preferences, it exists and can be computed efficiently using a carefully designed round robin algorithm.

We also ask whether it is possible to find matchings that satisfy *double maximin share guarantee* (DMMS), a two-sided version of the maximin share guarantee. Even when both sides have identical preferences, a complete DMMS matching may not exist, in contrast to the one-sided setting in which an MMS allocation is guaranteed to exist when agents have identical preferences. In general, we show that approximate DMMS and approximate DEF1 are incompatible, although in the special case where the degree constraint is equal to two we can achieve exact versions of both simultaneously.

Related work. Most related to our work is that of Patro *et al.* [2020], who draw on the resource allocation literature to guarantee fairness for both producers and consumers on a two-sided platform. However, in their model, producers are indifferent between the customers; thus, only one side has interesting preferences. Other work [Sühr *et al.*, 2019] has focused on guaranteeing fairness in two-sided platforms over time, rather than in a one-shot setting. Of particular note is the work of Gollapudi *et al.* [2020], who consider two-sided EF1 in a dynamic setting, but obtain positive results primarily for symmetric binary valuations, a much more restrictive class of valuations than we consider. Tadenuma [2011] studies envy minimization in two-sided matching subject to other notions, including stability, but focuses on ordinal notions of envy and restricts attention to one-to-one matchings.

The theories of matching and fair division each have a rich history. Traditional work in matching theory has focused on one-to-one or many-to-one matchings, beginning with the seminal work of Gale and Shapley [1962] and finding applications in areas such as school choice [Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu *et al.*, 2005;

Hatfield *et al.*, 2011], kidney exchange [Saidman *et al.*, 2006], and the famous US resident-to-hospital match.³ We note that EF1 as a condition becomes vacuous whenever a set of agents has a maximum degree constraint of one, so we focus instead on the more general case of many-to-many matchings. This case has also been well-explored in the matching literature [Roth, 1984; Sotomayor, 1999; Roth and Sotomayor, 1992; Echenique and Oviedo, 2006], although that literature focuses on stability notions, which have a very different flavor to our guarantees.

Our work draws extensively on notions from the fair division literature, particularly envy-freeness and its relaxations [Foley, 1967; Budish, 2011; Lipton *et al.*, 2004] and the maximin share guarantee [Budish, 2011]. Prior work has studied the satisfiability of these properties in the resource allocation setting [Caragiannis *et al.*, 2019; Procaccia and Wang, 2014; Kurokawa *et al.*, 2016], including the house allocation setting in which each agent is “matched” to a single item [Aigner-Horev and Segal-Halevi, 2019; Beynier *et al.*, 2019; Gan *et al.*, 2019], but, to our knowledge, no work has considered satisfying them on both sides of a market simultaneously.

2 Preliminaries

For $n \in \mathbb{N}$, define $[n] = \{0, \dots, n-1\}$. There are two disjoint groups of agents, denoted N^ℓ (“left”) and N^r (“right”), of sizes n^ℓ and n^r , respectively. For simplicity of notation, we write $N^\ell = [n^\ell]$ and $N^r = [n^r]$; when referring to an agent by only its index, the group she belongs to will be clear from context. We use indices $i \in [n^\ell]$ and $j \in [n^r]$ to refer to agents on the left and right, respectively. We are given degree constraints d_i^ℓ and d_j^r such that each $i \in N^\ell$ and each $j \in N^r$ can be matched to at most d_i^ℓ and d_j^r agents on the opposite side, respectively. When $d_i^\ell = d_{i'}^\ell$ for any $i, i' \in N^\ell$ (resp. $d_j^r = d_{j'}^r$ for any $j, j' \in N^r$), we denote by d^ℓ (resp. d^r) the common degree constraint of all agents in N^ℓ (resp. N^r).

A (many-to-many) *matching* M is represented as a binary $n^\ell \times n^r$ matrix, where $M(i, j) = 1$ if $i \in N^\ell$ and $j \in N^r$ are matched, and $M(i, j) = 0$ otherwise. With slight abuse of notation, we denote $M_i^\ell = \{j \in N^r : M(i, j) = 1\}$ and $M_j^r = \{i \in N^\ell : M(i, j) = 1\}$ as the sets of agents on the opposite side that agents $i \in N^\ell$ and $j \in N^r$ are matched to, respectively. We say that M is *valid* if it respects the degree constraints, i.e., if $|M_i^\ell| \leq d_i^\ell$ for each $i \in N^\ell$ and $|M_j^r| \leq d_j^r$ for each $j \in N^r$. Hereinafter, we omit the term valid, but will always refer to valid matchings. We say that M is *complete* if $\sum_{i \in N^\ell} |M_i^\ell| = \sum_{j \in N^r} |M_j^r| = \min(\sum_{i \in N^\ell} d_i^\ell, \sum_{j \in N^r} d_j^r)$. That is, a complete matching is one in which either every agent on the left has their degree constraint met exactly, or every agent on the right does.

Each agent $i \in N^\ell$ has a valuation function $u_i^\ell : N^r \rightarrow \mathbb{R}_{\geq 0}$ and each agent $j \in N^r$ has a valuation function $u_j^r : N^\ell \rightarrow \mathbb{R}_{\geq 0}$. When agents $i \in N^\ell$ and $j \in N^r$ are matched, they simultaneously receive utilities $u_i^\ell(j)$ and $u_j^r(i)$, respectively. We assume that utilities are additive. Thus, with

³<https://www.nrmmp.org/>

slight abuse of notation, the utilities to agents $i \in N^\ell$ and $j \in N^r$ under matching M are $u_i^\ell(M_i^\ell) = \sum_{j \in M_i^\ell} u_i^\ell(j)$ and $u_j^r(M_j^r) = \sum_{i \in M_j^r} u_j^r(i)$, respectively.

Our main constructive results take only the agents' preference orders as input. For agent $i \in N^\ell$ (resp. $j \in N^r$), we denote by σ_i^ℓ (resp. σ_j^r) a linear order over N^r (resp. N^ℓ) which is consistent with the valuation function u_i^ℓ (resp. u_j^r), i.e., $j \succ_{\sigma_i^\ell} j'$ whenever $u_i^\ell(j) > u_i^\ell(j')$ (resp. $i \succ_{\sigma_j^r} i'$ whenever $u_j^r(i) > u_j^r(i')$).⁴ With a slight abuse of notation, we denote with $\sigma_i^\ell(p)$ (resp. $\sigma_j^r(p)$) the position of alternative p in σ_i^ℓ (resp. σ_j^r).

Inspired by envy-freeness up to one good (EF1) from classical fair division [Budish, 2011; Lipton *et al.*, 2004], we define the following fairness guarantee in our setting.

Definition 1 (Double Envy-Freeness Up To c Matches (DEF c)). We say that matching M is envy-free up to c matches (EF c) over N^ℓ if for each pair of agents $i, i' \in N^\ell$, there exists $S^\ell \subseteq M_{i'}^\ell$ with $|S^\ell| \leq c$ such that $u_i^\ell(M_i^\ell) \geq u_i^\ell(M_{i'}^\ell \setminus S^\ell)$. Similarly, we say that it is EF c over N^r if, for each pair of agents $j, j' \in N^r$, there exists $S^r \subseteq M_{j'}^r$ with $|S^r| \leq c$ such that $u_j^r(M_j^r) \geq u_j^r(M_{j'}^r \setminus S^r)$. We say that M is DEF c if it is EF c over both N^ℓ and N^r .

When an algorithm takes as input only the preference rankings, it must ensure that the matching it returns is DEF c for all possible valuation functions which could have induced the rankings. It is easy to observe that this is equivalent to satisfying the following stronger guarantee which uses the stochastic dominance (SD) relation. This is akin to the SD-EF1 strengthening of EF1 [Freeman *et al.*, 2020; Aziz, 2020].

Definition 2 (SD Double Envy-Freeness Up To c Matches (SD-DEF c)). We say that matching M is SD-envy-free up to c matches (SD-EF c) over N^ℓ if, for every $t \in [n^\ell]$,

$$\sum_{p=0}^t M(i, \sigma_i^\ell(p)) \geq \sum_{p=0}^t M(i', \sigma_i^\ell(p)) - c, \forall i, i' \in N^\ell,$$

and is SD-EF c over N^r if, for every $t \in [n^r]$,

$$\sum_{p=0}^t M(\sigma_j^r(p), j) \geq \sum_{p=0}^t M(\sigma_j^r(p), j') - c, \forall j, j' \in N^r.$$

M is called SD-DEF c if it is SD-EF c over both N^ℓ and N^r .

Finally, we extend a different fairness notion from classical fair division called the maximin share guarantee (MMS).

Definition 3 (α -Double Maximin Share Guarantee (α -DMMS)). Let \mathcal{M} denote the set of valid matchings. The *maximin share value* of agent $i \in N^\ell$ is defined as

$$MMS_i^\ell = \max_{M \in \mathcal{M}} \min_{i' \in N^\ell} u_i^\ell(M_{i'}^\ell),$$

and the maximin share value of agent $j \in N^r$ is defined as

$$MMS_j^r = \max_{M \in \mathcal{M}} \min_{j' \in N^r} u_j^r(M_{j'}^r).$$

Given $\alpha \in [0, 1]$, matching M is called α -maximin share fair (α -MMS) over N^ℓ if $u_i^\ell(M_i^\ell) \geq \alpha \cdot MMS_i^\ell$ for every $i \in N^\ell$, and α -MMS over N^r if $u_j^r(M_j^r) \geq \alpha \cdot MMS_j^r$ for every $j \in N^r$. It is called α -DMMS if it is α -MMS for both N^ℓ and N^r . When $\alpha = 1$, we write DMMS instead of 1-DMMS.

⁴Ties among agents with equal utility are broken arbitrarily.

The notions of (SD-)DEF1 and DMMS are incomparable to the traditional notions of stability and justified envy-freeness, as the following example shows.

Example 1. Suppose $N^\ell = N^r = \{0, 1, 2, 3\}$, the common degree requirement is 2, and each side has identical ordinal preference $0 \succ 1 \succ 2 \succ 3$ over the other side. The only matching that is stable and eliminates justified envy is the one that matches each $i \in \{0, 1\}$ on the left with every $j \in \{0, 1\}$ on the right, and each $i \in \{2, 3\}$ on the left with every $j \in \{2, 3\}$ on the right. Indeed, if some $i \in \{0, 1\}$ on the left is not matched to some $j \in \{0, 1\}$ on the right, then j must be matched to some $i' \in \{2, 3\}$ on the left, which would make (i, j) a blocking pair, and i would (justifiably) envy i' for her match with j . However, this matching violates DEF1 when, for example, agent 2 on the left has more value for agent 1 on the right than for agents 2 and 3 on the right combined (as this would leave her envious of agents 0 and 1 on the left, even after ignoring their match to agent 0 on the right). Note that this matching also violates DMMS, since each agent on the left could partition those on the right into bundles $\{0, 3\}$, $\{0, 3\}$, $\{1, 2\}$, $\{1, 2\}$, guaranteeing themselves a better bundle than the $\{2, 3\}$ that agents 2 and 3 receive. On the other hand, any one-to-one matching satisfies (SD-)DEF1 and DMMS, but many one-to-one matchings are not stable or free of justified envy.

3 Double Envy-Freeness Up To One Match

In this section, we focus on double envy-freeness up to one match, more specifically, its strengthening SD-DEF1. We begin in Section 3.1 by presenting an impossibility result that holds even under quite restrictive conditions. Then, in Section 3.2, we present an algorithm that efficiently computes an SD-DEF1 matching whenever both groups of agents have identical ordinal preferences. In the full version of the paper, we present an additional positive result for the case that all agents have maximum degree constraint equal to two, and one side has identical preferences.

3.1 SD-DEF1 Matchings May Not Exist

Our first main result says that a complete SD-DEF1 matching may not exist. Observe that without the completeness condition an empty matching is trivially SD-DEF1.

Theorem 1. *A complete SD-DEF1 matching is not guaranteed to exist.*

The proof of Theorem 1, along with all other omitted proofs, can be found in the full version of the paper. It uses a counterexample in which both sides have the same number of agents (that is, $N^\ell = N^r = n$), all agents have the same degree constraint ($d^\ell = d^r = d$), and one group of agents have identical preferences ($u_i^\ell = u_{i'}^\ell$, for all $i, i' \in N^\ell$). Thus, Theorem 1 holds even in this restricted case, and continues to hold for more general settings.

A natural question, that we leave open, is whether the impossibility continues to hold when we relax SD-DEF1 to DEF1. We have found no counterexamples (even for SD-DEF1) via simulation; the counterexample for SD-DEF1 is carefully crafted but relies on the strength of SD-DEF1.

3.2 Identical Ordinal Preferences on Both Sides

Theorem 1 says that a complete SD-DEF1 matching is not guaranteed to exist even under quite restrictive conditions. It is natural to ask whether there exist any settings in which a complete SD-DEF1 matching can be guaranteed. In the full version of the paper, we establish the existence of a complete SD-DEF1 matching by restricting the degree bound (with still only one side having identical preferences). In this section, we consider a different restriction: when agents on both sides have identical preferences, i.e., $\sigma_i^\ell = \sigma_i^r$, for all $i, i' \in N^\ell$ and $\sigma_j^\ell = \sigma_j^r$, for all $j, j' \in N^r$.

Theorem 2. *When $n^\ell d^\ell = n^r d^r$ and both groups of agents have identical ordinal preferences, a complete SD-DEF1 matching always exists and can be computed efficiently.*

The proof of Theorem 2 follows from a series of lemmas. In the main text we focus on the simple case for which $N^\ell = N^r = n$ and $d^\ell = d^r = d$. Theorem 2 follows from progressively reducing the general case to this simple case.

We denote by σ^ℓ and σ^r the ordinal preferences of the agents in N^ℓ and N^r , respectively. Without loss of generality, assume that $\sigma^\ell = \sigma^r = 0 \succ \dots \succ n-1$. We want to find an SD-DEF1 matching under which each agent is matched to exactly d agents on the opposite side. A natural idea is to let agents on one side pick agents on the other side in a round-robin fashion. That is, we construct an ordering R over agents on one side, and these agents take turns according to R in a cyclic fashion with each agent, in her turn, making one match to her most preferred agent (i.e. lowest indexed agent) on the opposite side who has less than d matches so far. A standard argument from classical fair division shows that regardless of the ordering R , the resulting matching will be SD-EF1 over the side that does the picking.⁵ However, as the example below shows, not all orderings R lead to a matching that also satisfies SD-EF1 over the other side.

Example 2. Consider the case where $n = 5$ and $d = 2$. Suppose the ordering R has agents on the left choose in the order 0, 1, 2, 3, 4. Then, agent 0 on the right will be matched to agents 0 and 1 on the left, while agent 1 on the right will be matched to agents 2 and 3 on the left. SD-EF1 is violated as agent 1 significantly envies agent 0 on the right side.

We now show that when R is carefully designed, SD-EF1 can also be satisfied over the other side, resulting in SD-DEF1. Algorithm 1 takes as input parameters $a \in [n]$ and $x \in \{d, n-d\}$, and for any choices of these parameters, constructs an ordering R over the agents on (say) the left side. Algorithm 2 then uses this ordering to run the round-robin procedure while respecting the degree constraints. Example 3 demonstrates these algorithms.

Example 3. Consider the same instance as Example 2, with $n = 5$ and $d = 2$. Suppose we choose $a = 3$ and $x = d = 2$. Then the round robin ordering returned by Algorithm 1 is $R(0) = 3 + 0 = 3$ (setting $i = 0$), $R(1) = 3 + 3 = 1$ ($i = 3$), $R(2) = 3 + 1 = 4$ ($i = 1$), $R(3) = 3 + 4 = 2$ ($i = 4$),

⁵As observed by Biswas and Barman [2018], the standard round robin algorithm is not EF1 when agents have cardinality constraints, but EF1 is retained provided that agents have identical preferences.

Algorithm 1 Round-Robin-Ordering(n, a, x)

```

1: //  $x$  and  $n$  coprime, so  $x^{-1} \pmod{n}$  exists
2: for  $i \in [n]$  do
3:    $R(i) = a + px^{-1} \pmod{n}$ 
4: end for
5: return  $R$ 

```

Algorithm 2 Restricted-Round-Robin-Coprime(n, d)

```

1: Choose  $a \in \{0, \dots, n-1\}$  and  $x \in \{d, n-d\}$ 
2:  $R = \text{Round-Robin-Ordering}(n, a, x)$ 
3: // Round-robin with ordering  $R$  over agents on the left
4:  $M(i, j) = 0, \forall i, j \in [n]$ 
5: for  $j \in [n], t \in [d]$  do
6:    $M(R(j \cdot d + t \pmod{n}), j) = 1$ 
7: end for
8: return  $M$ 

```

and $R(4) = 3 + 2 = 0$ ($i = 2$), with all addition performed mod n . That is, agents on the left choose in order 3, 1, 4, 2, 0. This results in the matching $M_3^\ell = \{0, 2\}$, $M_1^\ell = \{0, 3\}$, $M_4^\ell = \{1, 3\}$, $M_2^\ell = \{1, 4\}$, and $M_0^\ell = \{2, 4\}$ (equivalent to the formula provided directly in Line 6 of Algorithm 2). The fact that this is SD-EF1 over N^ℓ is easy to check. Examining the matching, note that $M_0^r = \{1, 3\} = M_4^\ell$, $M_1^r = \{2, 4\} = M_0^\ell$, $M_2^r = \{0, 3\} = M_4^\ell$, $M_3^r = \{1, 4\} = M_2^\ell$, and $M_4^r = \{0, 2\} = M_3^\ell$. That is, the matchings received by agents on the right are the same as those received by agents on the left, up to a cyclic shift. For this matching, SD-EF1 over N^ℓ immediately implies SD-EF1 over N^r .

The next result shows that for any choices of the parameters, the resulting matching is SD-DEF1. The idea of the proof is to show that the structure in Example 3 holds in general: for any allowed choice of (a, x) , the set of bundles received by agents on the right is the same as the set of bundles received by agents on the left, thus inheriting SD-EF1 from the fact that the matching is constructed by agents on the left choosing in round robin sequence.

Lemma 1. *When $n^\ell = n^r = n$ and $d^\ell = d^r = d$ are coprime and both groups of agents have identical ordinal preferences, Algorithm 2 efficiently computes a complete SD-DEF1 matching.*

Proof. To avoid the \pmod{n} notation in this proof, we will treat integers as belonging to the ring $\mathbb{Z}/n\mathbb{Z}$ of integers modulo n . Thus, addition, multiplication, and multiplicative inverses will be modulo n . Note that x^{-1} exists because $x \in \{d, n-d\} = \{d, -d\}$ is coprime with n .

We claim that the ordering R constructed in Algorithm 1 is a valid ordering over the agents in N^ℓ . Notice that because $x \in \{d, -d\}$ is coprime with n , $(p \cdot x^{-1})_{i \in [n]} = [n]$. Thus, each position in the ordering R is mapped to exactly one agent. Because agents on the left take d turns in a cyclic fashion, it is convenient to think of an *extended* ordering R which is the original R concatenated with itself d times: one can check that this still obeys $R(p) = a + px^{-1}$ for all $p \in [nd]$.

Next, we argue that the matching returned is a valid complete matching. Notice that during the round-robin, d agents

on the left that are consecutive in the ordering pick a given agent on the right before moving on to the next lowest indexed agent on the right. Further, each agent on the left gets d turns. Hence, it is easy to see that every agent is matched to exactly d agents on the opposite side.

As mentioned earlier, the fact that the returned matching M is SD-EF1 over N^ℓ follows directly from the standard round-robin argument in classical fair division: given any pair of agents $i, i' \in N^\ell$, if we ignored the first turn taken by i' , then in each round agent i would get a turn before agent i' does, and hence, would not envy agent i' in the SD sense. It remains to show that M is also SD-EF1 over N^r . We show that for each agent $j \in N^r$, there exists an agent $i \in N^\ell$ such that $M_j^r = M_i^\ell$. SD-EF1 over N^r will then follow from SD-EF1 over N^ℓ given that $\sigma^\ell = \sigma^r$.

Let us focus on agent $j \in N^r$. Because agents on the right are picked from lowest-indexed to highest-indexed, agent j is picked by the d agents from N^ℓ who appear consecutively in the (extended) ordering R at indices $jd + t$ for $t \in [d]$. Given that $R(p) = a + px^{-1}$ for all $p \in [nd]$, we immediately have $M_j^r = \{a + (jd + t)x^{-1} : t = [d]\}$.

Next, let us focus on agent $i \in N^\ell$. If she is matched to some agent $j \in N^r$ in a particular turn, then from the observation above, it must be that $i = a + (jd + t)x^{-1}$ for some $t \in [d]$. Solving this for j , we get that $j = ((i - a)x - t)d^{-1}$. Varying $t \in [d]$ in this equation gets us the d agents on the right that agent i is matched to: $M_i^\ell = \{((i - a)x - t)d^{-1} : t = [d]\}$.

To show that for each $j \in N^r$, there exists $i \in N^\ell$ with $M_i^\ell = M_j^r$, we take two cases.

If $x = n - d = -d$, then $x^{-1} = (-d)^{-1} = -d^{-1}$. In this case, it is easy to check that taking $i = j$ suffices as $M_j^r = M_j^\ell = \{a - j - td^{-1} : t = [d]\}$.

If $x = d$, then $M_j^r = \{j + a + td^{-1} : t = [d]\}$, while

$$\begin{aligned} M_i^\ell &= \{i - a - td^{-1} : t = [d]\} \\ &= \{i - a - (d - 1 - t)d^{-1} : t = [d]\}. \end{aligned}$$

Notice that M_j^r coincides with $M_{j+2a-d^{-1}+1}^\ell$. \square

Algorithm 2 executes round-robin with the left side taking turns, and allows freely choosing $a \in [n]$ and $x \in \{d, n - d\}$ to decide their ordering. Note that if the right side takes turns instead, the algorithm still produces a complete SD-DEF1 matching. However, this extension does not find any new matchings. When $x = n - d$, the matching produced is symmetric ($M_i^\ell = M_i^r$ for all $i \in [n]$), and thus the same regardless of which side takes turns. When $x = d$, the allocations on one side are cyclic shift of the allocations on the other side. Hence, any matching produced by the right side taking turns can also be produced by the left side taking turns with appropriately chosen (a, x) .

What about allowing choices of x other than just d and $n - d$? At least for $n = 7$, $d = 3$, and $a = 0$, it is easy to check by hand that no other choices of x produce an SD-DEF1 matching. On the other hand, could it be that some of the $2n$ choices of (a, x) are redundant and lead to the same matching as other choices? The following result shows that in every instance, all $2n$ choices lead to different matchings.

Proposition 1. For any inputs n and d to Algorithm 2, the $2n$ possible choices of (a, x) result in distinct matchings.

Given Proposition 1, one may be tempted to conjecture that these $2n$ choices generate all complete SD-DEF1 matchings. However, in the full version of the paper, we show that this is not the case, leaving open the question of characterizing the set of all complete SD-DEF1 matchings.

The proof of Theorem 2 continues by reducing the case where n and d are not coprime to the coprime case. Letting $g = \gcd(n, d)$, we divide both sides into g sub-groups of $n' = n/g$ agents each. Then, we run Algorithm 2 a total of g^2 times to match agents from each sub-group on the left to $d' = d/g$ agents from each sub-group on the right. This matches each agent with exactly d agents from the opposite side. Note that we allow each of the g^2 calls to Algorithm 2 to use arbitrary choices of a and x . Nonetheless, we show that the resulting complete matching must be SD-DEF1.

Lemma 2. When $n^\ell = n^r = n$, $d^\ell = d^r = d$, and both groups of agents have identical ordinal preferences, a complete SD-DEF1 matching always exists and can be computed efficiently.

Finally, we turn our attention to the general case in which we drop the constraints $n^\ell = n^r$ and $d^\ell = d^r = d$. We do however require that $n^\ell \cdot d^\ell = n^r \cdot d^r$. The proof of Theorem 2, which appears in the full version of the paper, uses a trick of adding dummy agents to the side with fewer agents, computing an SD-DEF1 matching as per Lemma 2, and then removing the dummy agents. The key is to show that the removal of dummy agents reduces the degrees of the agents on the opposite side exactly as intended and SD-DEF1 is preserved.

We note that it is possible to extend our constructive result slightly beyond the case of $n^\ell \cdot d^\ell = n^r \cdot d^r$. Without loss of generality, assume that $n^\ell \cdot d^\ell < n^r \cdot d^r$. First, note that in this case, no matching is complete. We can still make the degree of each agent on the left equal to d^ℓ , but the best we can hope for is that the degrees of agents on the right differ by at most 1, i.e., they are either $\lfloor n^\ell \cdot d^\ell / n^r \rfloor$ or $\lceil n^\ell \cdot d^\ell / n^r \rceil$.⁶ In this case, the trick outlined in Theorem 2 only works when the dummy agents are added to the left side, i.e., if $n^\ell \leq n^r$. We conjecture that such an SD-DEF1 matching always exists even when $n^\ell > n^r$, but leave it as an open question.

4 Double Maximin Share Guarantee

In this section, we focus first on the existence of DMMS matchings, and second on the existence of matchings that are DMMS and SD-DEF1 concurrently.

We begin by considering the case where agents on both sides have identical preferences, i.e., $u_i^\ell(j) = u_i^r(j)$, for any pair of agents $i, i' \in N^\ell$, and any $j \in N^r$, and similarly $u_j^r(i) = u_j^\ell(i)$, for any pair of agents $j, j' \in N^r$ and any $i \in N^\ell$. We show the following negative result, which stands in contrast to the one-sided fair division setting in which an MMS allocation is guaranteed to exist when agents have identical preferences.

⁶In case that $n^\ell \cdot d^\ell / n^r$ is an integer, we can set this to be d^r and achieve exactly equal degrees on the right side too.

Theorem 3. *A 0.89-DMMS matching may not exist, even when agents on both sides have identical preferences.*

Proof. We denote by u^ℓ and u^r the cardinal preferences of the agents in N^ℓ and N^r respectively. As the utilities are the same across the agents in the same group, we can define $MMS^\ell = MMS_i^\ell$ for all $i \in [n^\ell]$, and $MMS^r = MMS_j^r$ for all $j \in [n^r]$.

Consider the instance with $n = n^\ell = n^r = 7$ and $d = d^\ell = d^r = 3$, $u^\ell(j) = n - j - 1$ for all $j \in [n]$, and $u^r(i) = n - i - 1$ for all $i \in [n]$. Thus, for any complete matching, $\sum_{i \in N^\ell} u^\ell(M_i^\ell) = \sum_{j \in N^r} u^r(M_j^r) = 63$. This means that $MMS^\ell = MMS^r \leq 9$, because if all agents receive equal utility then they each get utility 9. Next, we construct a matching M such that $u^\ell(M_i^\ell) = 9$ for all $i \in [n]$.

Without loss of generality, assume that M_0^ℓ , M_1^ℓ , and M_2^ℓ all contain agent 0. Then, we know that agents 1 and 2 cannot be contained in these bundles, because then they would have value larger than 9, implying that some other agent receives utility less than 9. Without loss of generality, we assume that bundles M_3^ℓ , M_4^ℓ , and M_5^ℓ contain agent 1. Now, we observe that $M_6^\ell = \{2, 3, 4\}$, as there is no other way to have $u^\ell(M_6^\ell) = 9$. As 0 and 2 can not belong to the same bundle (such a bundle would be valued at least 10), we may assume without loss of generality that agent 2 is contained in M_3^ℓ , and M_4^ℓ . Then, the constraint that $u^\ell(M_3^\ell) = u^\ell(M_4^\ell) = 9$ dictates that $M_3^\ell = M_4^\ell = \{1, 2, 6\}$. With these bundles fixed, it is easy to check that the only M_5^ℓ that yields $u^\ell(M_5^\ell) = 9$ is $M_5^\ell = \{1, 3, 5\}$. Lastly, without loss of generality, we may assume that $M_0^\ell = M_1^\ell = \{0, 4, 5\}$, and $M_2^\ell = \{0, 3, 6\}$. Hence, we conclude that the following matching is the only one (subject to permutations of N^ℓ) that satisfies MMS for agents on the left.

- $M_0^\ell = M_1^\ell = \{0, 4, 5\}$
- $M_2^\ell = \{0, 3, 6\}$
- $M_3^\ell = M_4^\ell = \{1, 2, 6\}$
- $M_5^\ell = \{1, 3, 5\}$
- $M_6^\ell = \{2, 3, 4\}$

Now, consider agents $0 \in N^r$ and $4 \in N^r$. Both are matched to agents $0 \in N^\ell$ and $1 \in N^\ell$, but agent $0 \in N^r$ is matched to agent $2 \in N^\ell$ while agent $4 \in N^r$ is matched to agent $6 \in N^\ell$. Therefore, $u^r(M_0^r) \neq u^r(M_4^r)$ (and this difference persists regardless of permutations of N^ℓ). It is therefore not the case that every agent on the right receives utility 9; in particular, one agent receives utility 8 or less, producing the approximation ratio $\alpha = 8/9 < 0.89$. \square

While a DMMS matching may not exist, even when preferences are identical, we can exploit the algorithms presented in Section 3 to obtain an approximation to DMMS.

Theorem 4. *When $n^\ell d^\ell = n^r d^r$ and both groups of agents have identical utilities, every complete SD-DEF1 matching M is also $\frac{1}{d^\ell}$ -MMS over N^ℓ , and $\frac{1}{d^r}$ -MMS over N^r .*

We next show an almost-matching upper bound that can be achieved by any SD-DEF1 matching, to complement Theorem 4. In fact, we show a more general result that trades off the approximation to DMMS with the approximation to double envy-freeness.

Theorem 5. *There exists an instance with $n^\ell = n^r = n$ and $d^\ell = d^r = d$ in which no matching is simultaneously $\frac{c+2}{d}$ -DMMS and SD-DEF c for any $c \in [d]$.*

Finally, we show that a strong impossibility persists even if we only require SD-EF1 on one side and MMS on the other.

Theorem 6. *A matching that satisfies SD-EF1 over N^ℓ and MMS over N^r is not guaranteed to exist, even when agents on both sides have identical preferences.*

5 Discussion

We have introduced a model that bridges two-sided matching and fair division by requiring fairness on both sides of a matching market. We have shown that SD-EF1 can be achieved for agents on both sides, when all agents on the left side and all agents on the right share a common ordinal preference ranking over agents on the other side. When this condition is not satisfied, there may exist no matching that satisfies SD-DEF1. We have also shown that there may not exist a doubly MMS matching even when agents have identical preferences. While we do not rule out a good approximation to DMMS, we show that it is essentially impossible to obtain a good approximation to DMMS if one also requires SD-DEF1.

It is interesting to note that the proofs of Theorems 4 and 5 do not rely on the constraints that an agent in N^ℓ can have up to d matches, and can be matched with an agent in N^r at most once. Therefore, these theorems also hold in a one-sided fair division problem where there are n agents and n/d items with d copies each, and all the agents have identical preferences.

Many interesting avenues for future research remain. For example, one can hope to derive weaker positive results in the case where one side has identical preferences. When the side with identical preferences does the picking, Algorithm 2 remains EF1 for that side. In the full version of the paper, we conduct empirical simulations and observe that Algorithm 2 remains EF1 for some of the agents on the other side as well (and does better than the classical round-robin algorithm in this aspect). It would also be interesting to compare the two-sided fair division setting with its one-sided counterpart (where only one side has preferences and we seek fairness only for this side). In our simulations, we observe that there is a sharp contrast for envy-freeness (one-sided EF is almost always achievable while two-sided DEF almost always isn't). For the maximin share guarantee, however, there is no contrast: both one-sided MMS and two-sided DMMS are almost always achievable. Theoretically analyzing the probability of satisfiability of these notions in random instances would be an interesting direction for the future. One can also consider two-sided versions of other fairness notions, including those that remain interesting when the degree constraint is 1, which could yield further interesting results in the one-to-one or many-to-one settings. Finally, it would also be interesting to derive positive results when each agent can have a different degree constraint.

References

- [Abdulkadiroğlu and Sönmez, 2003] A. Abdulkadiroğlu and T. Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.
- [Abdulkadiroğlu *et al.*, 2005] A. Abdulkadiroğlu, P.A. Pathak, A.E. Roth, and T. Sönmez. The boston public school match. *American Economic Review*, 95(2):368–371, 2005.
- [Aigner-Horev and Segal-Halevi, 2019] E. Aigner-Horev and E. Segal-Halevi. Envy-free matchings in bipartite graphs and their applications to fair division. *arXiv preprint arXiv:1901.09527*, 2019.
- [Aziz, 2020] H. Aziz. Simultaneously achieving ex-ante and ex-post fairness. In *Proc. of 16th WINE*, pages 341–355, 2020.
- [Beynier *et al.*, 2019] A. Beynier, Y. Chevaleyre, L. Gourvès, A. Harutyunyan, J. Lesca, N. Maudet, and A. Wilczynski. Local envy-freeness in house allocation problems. *Autonomous Agents and Multi-Agents Systems*, 33(5):591–627, 2019.
- [Biswas and Barman, 2018] A. Biswas and S. Barman. Fair division under cardinality constraints. In *Proc. of 27th IJCAI*, pages 91–97, 2018.
- [Budish and Cantillon, 2012] E. Budish and E. Cantillon. The multi-unit assignment problem: Theory and evidence from course allocation at harvard. *American Economic Review*, 102(5):2237–71, 2012.
- [Budish, 2011] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Caragiannis *et al.*, 2019] I. Caragiannis, D. Kurokawa, H. Moulin, A.D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–32, 2019.
- [Echenique and Oviedo, 2006] F. Echenique and J. Oviedo. A theory of stability in many-to-many matching markets. *Theoretical Economics*, 1(2):233–273, 2006.
- [Foley, 1967] D. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.
- [Freeman *et al.*, 2020] R. Freeman, N. Shah, and R. Vaish. Best of both worlds: Ex-ante and ex-post fairness in resource allocation. In *Proc. of 21st EC*, pages 21–22, 2020.
- [Gale and Shapley, 1962] D. Gale and L.S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1):9–15, 1962.
- [Gan *et al.*, 2019] J. Gan, W. Suksompong, and A.A. Voudouris. Envy-freeness in house allocation problems. *Mathematical Social Sciences*, 101:104–106, 2019.
- [Gollapudi *et al.*, 2020] S. Gollapudi, K. Kollias, and B. Plaut. Almost envy-free repeated matching in two-sided markets. In *Proc. of 16th WINE*, pages 3–16, 2020.
- [Hatfield *et al.*, 2011] J.W. Hatfield, F. Kojima, Y. Narita, et al. Promoting school competition through school choice: A market design approach. *Stanford: Stanford University Capital and Economic Opportunity Working Group Working Paper*, 18:2011, 2011.
- [Kurokawa *et al.*, 2016] D. Kurokawa, A.D. Procaccia, and J. Wang. When can the maximin share guarantee be guaranteed? In *Proc. of 30th AAAI*, pages 523–529, 2016.
- [Lipton *et al.*, 2004] R.J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proc. of 6th EC*, pages 125–131, 2004.
- [Othman *et al.*, 2010] A. Othman, T. Sandholm, and E. Budish. Finding approximate competitive equilibria: Efficient and fair course allocation. In *Proc. of 9th AAMAS*, pages 873–880, 2010.
- [Patro *et al.*, 2020] G.K. Patro, A. Biswas, N. Ganguly, K.P. Gummadi, and A. Chakraborty. Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. In *Proc. of 29th WWW*, pages 1194–1204, 2020.
- [Procaccia and Wang, 2014] A.D. Procaccia and J. Wang. Fair enough: Guaranteeing approximate maximin shares. In *Proc. of 14th EC*, pages 675–692, 2014.
- [Roth and Sotomayor, 1992] A.E. Roth and M. Sotomayor. Two-sided matching. *Handbook of game theory with economic applications*, 1:485–541, 1992.
- [Roth, 1984] A.E. Roth. Stability and polarization of interests in job matching. *Econometrica: Journal of the Econometric Society*, pages 47–57, 1984.
- [Saidman *et al.*, 2006] S.L. Saidman, A.E. Roth, T. Sönmez, M.U. Ünver, and F.L. Delmonico. Increasing the opportunity of live kidney donation by matching for two and three way exchanges. *Transplantation*, 81:773–782, 2006.
- [Sotomayor, 1999] M. Sotomayor. Three remarks on the many-to-many stable matching problem. *Mathematical social sciences*, 38(1):55–70, 1999.
- [Sühr *et al.*, 2019] T. Sühr, A.J. Biega, M. Zehlike, K.P. Gummadi, and A. Chakraborty. Two-sided fairness for repeated matchings in two-sided markets: A case study of a ride-hailing platform. In *Proc. of the 25th SIGKDD*, pages 3082–3092, 2019.
- [Tadenuma, 2011] K. Tadenuma. Partnership, solidarity, and minimal envy in matching problems. In *Social Ethics and Normative Economics*, pages 155–167. Springer, 2011.
- [Wu and Roth, 2018] Q. Wu and A.E. Roth. The lattice of envy-free matchings. *Games and Economic Behavior*, 109:201–211, 2018.
- [Yokoi, 2020] Y. Yokoi. Envy-free matchings with lower quotas. *Algorithmica*, 82(2):188–211, 2020.