

When Do Envy-Free Allocations Exist?

Pasin Manurangsi
Department of EECS
UC Berkeley

Warut Suksompong
Department of Computer Science
University of Oxford

Abstract

We consider a fair division setting in which m indivisible items are to be allocated among n agents, where the agents have additive utilities and the agents' utilities for individual items are independently sampled from a distribution. Previous work has shown that an envy-free allocation is likely to exist when $m = \Omega(n \log n)$ but not when $m = n + o(n)$, and left open the question of determining where the phase transition from non-existence to existence occurs. We show that, surprisingly, there is in fact no universal point of transition—instead, the transition is governed by the divisibility relation between m and n . On the one hand, if m is divisible by n , an envy-free allocation exists with high probability as long as $m \geq 2n$. On the other hand, if m is not “almost” divisible by n , an envy-free allocation is unlikely to exist even when $m = \Theta(n \log n / \log \log n)$.

1 Introduction

Resource allocation is a fundamental task that occurs in a great number of everyday situations, from allocating school supplies to children and course slots in universities to students, to allocating machine processing time to users and kidneys to kidney transplant patients. One of the principal concerns when allocating resources to interested agents is *fairness*: we want all agents to feel that they receive a fair share of the resources. There is a rich and beautiful theory of *fair division* that goes back several decades and has been studied in mathematics, economics, and more recently in computer science (Brams and Taylor, 1996; Moulin, 2003).

In order to reason about fairness, we must define when an allocation is considered to be “fair”. One of the most prominent fairness notions is *envy-freeness*, which means that every agent likes her allocated portion at least as much as that of any other agent (Foley, 1967; Varian, 1974). While an envy-free allocation can always be obtained when we allocate *divisible* goods such as land or machine processing time (Stromquist, 1980), this is not the case when it comes to allocating *indivisible* goods like jewelry and artworks. Indeed, if a single bracelet or painting is to be divided between two agents, then no matter how the division is performed, the agent who does not receive the item will be left envying the other agent.

Given that the existence of envy-free allocations cannot be guaranteed in general for indivisible goods, an important question is therefore *when* such allocations exist. Dickerson et al. (2014) investigated this question under a simple model where the agents have additive utilities and their utilities for individual items are drawn at random from probability distributions. If the number of items, m , is less than the number of agents, n , no envy-free allocation exists since any allocation necessarily leaves some agent empty-handed and envious. Dickerson et al. showed that even when the number of items slightly exceeds the number of agents— $m = n + o(n)$ —an envy-free allocation is still unlikely to exist. However, as soon as the number of items is larger than the number of agents by a logarithmic factor— $m = \Omega(n \log n)$ —an envy-free allocation exists with high probability, and can furthermore be obtained by simply giving each item to the agent with the highest utility for it. Dickerson et al. also found the phase transition from non-existence to existence to be quite sharp in computer experiments, and left open the question of determining where this transition occurs. Is the logarithmic factor in the upper bound necessary, or do we already have existence when, say, $m = 1.001n$?

In this paper, we show that, surprisingly, there is in fact *no* universal point of transition between non-existence and existence. Instead, the transition is governed by the divisibility relation between m and n . On the one hand, if m is divisible by n , we show that an envy-free allocation exists with high probability as long as $m \geq 2n$ (Theorem 3.1). Our result improves upon the aforementioned $m = \Omega(n \log n)$ upper bound and moreover completely closes the gap for the case of divisibility, since Dickerson et al.'s lower bound already implies that the same result does not hold when $m = n$.¹ On the other hand, if m is not “almost” divisible by n , in the sense that the remainder of the division is between n^ϵ and $n - n^\epsilon$ for some constant $\epsilon \in (0, 1)$, we show that an envy-free allocation is unlikely to exist as long as $m = O(n \log n / \log \log n)$ (Theorem 4.1). This comes to within a $\Theta(\log \log n)$ factor of matching their upper bound. Both our existence and non-existence results rely on several new key ideas. In particular, for the existence result we need a com-

¹We do note, however, that Dickerson et al.'s upper bound holds under a weaker assumption on the distributions. For example, it does not assume that the utilities are drawn from the same distribution for all agents and items.

pletely different algorithm, since the welfare-maximizing algorithm used to achieve existence for $m = \Omega(n \log n)$ cannot yield any improvement of this bound (Proposition 3.2).

1.1 Related Work

Besides the work of Dickerson et al. (2014) that we mentioned, several other works have investigated the asymptotic existence and non-existence of fair allocations for various fairness notions. Suksompong (2016) considered *proportional* allocations—allocations in which every agent receives at least $1/n$ of her value for the whole set of items—and showed that such allocations exist with high probability if m is a multiple of n or $m = \omega(n)$. Kurokawa, Procaccia, and Wang (2016) showed that an allocation that satisfies the *maximin share criterion* is likely to exist as long as either m or n goes to infinity.² As in our work, both Kurokawa, Procaccia, and Wang (2016) and Suksompong (2016) used techniques from the theory of matchings in random graphs to establish the existence of fair allocations. Amanatidis et al. (2017) also addressed the existence of allocations satisfying the maximin share criterion. Finally, Manurangsi and Suksompong (2017) considered the setting where goods are allocated to groups of agents and generalized Dickerson et al. (2014)’s results on envy-freeness to that setting.

Since envy-free allocations cannot always be obtained even in the simplest setting with two agents and one item, a recent line of work has focused on relaxations of envy-freeness with the goal of recovering the guaranteed existence. These relaxations include *envy-freeness up to one good*—any envy that an agent has towards another agent can be eliminated by removing *some* item from the latter agent’s bundle—and *envy-freeness up to any good*—any such envy can be eliminated by removing *any* item from the latter agent’s bundle. It has been shown that these relaxations do provide existence guarantees in a number of settings (Lipton et al., 2004; Caragiannis et al., 2016; Conitzer, Freeman, and Shah, 2017; Amanatidis, Birmpas, and Markakis, 2018; Barman, Krisnamurthy, and Vaish, 2018; Biswas and Barman, 2018; Plaut and Roughgarden, 2018).

2 Preliminaries

A set $M = [m]$ of indivisible items is to be allocated to a set $N = [n]$ of agents, where we use $[k]$ to denote the set $\{1, 2, \dots, k\}$. Each agent i has a nonnegative utility $u_i(j)$ for item j . We assume that the utility $u_i(j)$ lies in $[0, 1]$; this does not introduce a loss of generality since we can scale down all utilities by their maximum. The utilities of the agents are *additive*, i.e., $u_i(M') = \sum_{j \in M'} u_i(j)$ for any $M' \subseteq M$. The additivity assumption is made in several works on fair division and, in particular, in all of the works on the asymptotic existence of fair allocations mentioned in Section 1.1.

A *bundle* refers to a subset of M . An *allocation* is a partition of M into n bundles (M_1, M_2, \dots, M_n) , where bundle

²We refer to their paper for the definition, but remark here that both proportionality and the maximin share criterion are weaker than envy-freeness when utilities are additive.

M_i is allocated to agent i . An allocation is said to be *envy-free for agent i* if $u_i(M_i) \geq u_i(M_j)$ for any $j \in N$, and *envy-free* if it is envy-free for every agent $i \in N$.

For agents $i \in N$ and items $j \in M$, the utilities $u_i(j)$ are drawn independently from a distribution \mathcal{U} . A distribution is said to be *non-atomic* if it does not put positive probability on any single point. The condition that we will impose on \mathcal{U} for our results is that it “behaves like a polynomial close to 1” in the sense that the function $g(\alpha) = \Pr_{u \sim \mathcal{U}}[u \geq 1 - \alpha]$ is bounded above and below by a polynomial. This is formalized in the following definition.

Definition 2.1. Let θ, q be any positive real numbers. A probability distribution \mathcal{U} on $[0, 1]$ is said to be (θ, q) -*polynomially bounded below* (resp. *above*) *at 1* if for every $\alpha \in (0, 1]$, we have $\Pr_{u \sim \mathcal{U}}[u > 1 - \alpha] \geq \theta \cdot \alpha^q$ (resp. $\Pr_{u \sim \mathcal{U}}[u > 1 - \alpha] \leq \theta \cdot \alpha^q$).

A probability distribution \mathcal{U} is said to be *polynomially bounded at 1* if there exist constants $\underline{\theta}, \bar{\theta}, q > 0$ such that \mathcal{U} is $(\underline{\theta}, q)$ -polynomially bounded below at 1 and $(\bar{\theta}, q)$ -polynomially bounded above at 1.

We assume in Section 3 that \mathcal{U} is polynomially bounded at 1 and in Section 4 that \mathcal{U} is polynomially bounded below at 1. To illustrate the generality of this definition, consider any non-atomic continuous distribution \mathcal{U} whose probability density function $f_{\mathcal{U}}$ is bounded below (resp. above) around 1, i.e., there exist $\epsilon, \beta > 0$ such that $f_{\mathcal{U}}(x) \geq \beta$ (resp. $f_{\mathcal{U}}(x) \leq \beta$) for all $x \geq 1 - \epsilon$. One can check that \mathcal{U} is polynomially bounded below (resp. above) at 1 with parameters $\theta = \epsilon \cdot \beta$ (resp. $\theta = \max\{\beta, 1/\epsilon\}$) and $q = 1$. This immediately implies that the uniform distribution on $[0, 1]$ and a normal distribution (with any mean and variance) truncated at 0 and 1 are polynomially bounded at 1, as both have probability density functions that are bounded both above and below in $[0, 1]$.

For completeness, let us also provide examples of distributions that are not polynomially bounded at 1. The first example is when $\Pr_{u \sim \mathcal{U}}[u = 1] > 0$. In this case, clearly the distribution is not polynomially bounded above at 1. Another example is if we take any \mathcal{U} such that $\Pr_{u \sim \mathcal{U}}[u \geq 1 - 1/2^i] = 1/2^{i^2}$ for all integer $i \geq 0$. It is not hard to see that this distribution is not polynomially bounded below at 1.

Indeed, for any fixed $q > 0$, we have $\lim_{i \rightarrow \infty} \frac{(1/2^{i^2})}{(1/2^i)^q} = 0$, which means that there is no $\theta > 0$ such that $\Pr_{u \sim \mathcal{U}}[u \geq 1 - \alpha] \geq \theta \cdot \alpha^q$ for all $\alpha \in (0, 1]$.

Finally, a statement is said to hold *with high probability* if the probability that it holds approaches 1 as $n \rightarrow \infty$.

3 Existence

In this section, we investigate the existence front of envy-free allocations. We first show that the welfare-maximizing algorithm of Dickerson et al. (2014) cannot yield any improvement of the $m = \Omega(n \log n)$ bound. We then prove the main existence result of this paper, which holds for any $m \geq 2n$ that is a multiple of n :

Theorem 3.1. *Let $r \geq 2$ be an integer, and suppose that $m = rn$. Assume that \mathcal{U} is polynomially bounded at 1. With*

high probability, there exists an envy-free allocation. Moreover, there is a polynomial-time algorithm that computes such an allocation.

A bonus of our algorithm is that it returns a *balanced* allocation, i.e., one that gives every agent the same number of items. This may be desirable in situations where capacity constraints are involved, for example if we divide artworks between museums or players between sports teams.

3.1 The Limit of the Welfare-Maximizing Algorithm

Recall the main existence result of Dickerson et al. (2014): when $m = \Omega(n \log n)$, the welfare-maximizing algorithm, which allocates each item to the agent who values it most, is likely to produce an envy-free allocation. We observe that this bound is tight up to a constant factor—for $m = n \log n - \omega(n)$ items, the welfare-maximizing allocation is unlikely to be envy-free. An implication of this observation is that the welfare-maximizing algorithm fails to be envy-free in the case where $m = rn$, for any positive integer $r \leq \log n - \omega(1)$. By contrast, the algorithm that we will present finds an envy-free allocation with high probability for any integer $r \geq 2$.

Proposition 3.2. *Let $m = n \log n - \omega(n)$, and suppose that \mathcal{U} is non-atomic. Then, with high probability, the welfare-maximizing allocation is not envy-free.*

Proof. The proposition follows from a classical result on the coupon collector’s problem. In this problem, there is an urn of n coupons. Each turn, a coupon is drawn uniformly at random from the urn and immediately returned to the urn. Erdős and Rényi (1961) proved that with high probability, after $n \log n - \omega(n)$ turns, some coupon has not been drawn.

The connection between the coupon collector’s problem and our setting is fairly simple. First, the non-atomic assumption on the distribution implies that, almost surely, all items yield positive utility to every agent, and every item has only one agent who values it most. As a result, the welfare-maximizing allocation assigns each item to each agent with probability $1/n$. If we view each agent as a coupon in the coupon collector’s problem, Erdős and Rényi’s result implies that with high probability, some agent does not receive any item in this allocation. From the positive utility observation, the allocation cannot be envy-free. \square

3.2 Warm-Up: A Simplified Algorithm for $r \geq 3$

The remainder of Section 3 is devoted to proving Theorem 3.1; we assume throughout that $m = rn$ for some integer $r \geq 2$. As Dickerson et al. (2014) already showed that the theorem holds for $r = \Omega(\log n)$, it suffices for us to establish the statement for $r = O(\log n)$. Nevertheless, we will prove the statement for $r \leq e^{n^{0.1}}$, which is much stronger; while this is not necessary, we do so to demonstrate that our algorithm and its analysis are robust and apply even when the number of items is significantly larger than the number of agents.

Before we proceed to the actual algorithm, let us provide the intuition behind the algorithm by describing a simpler

algorithm that works in all cases except when $r = 2$. For the sake of exposition, we shall restrict ourselves to the case where the distribution \mathcal{U} is the uniform distribution on $[0, 1]$ and $r > 2$ is a constant (i.e., does not grow with n). We shall also sometimes be informal here; all proofs will be formalized in the rest of Section 3.

The simplified algorithm tries to find an allocation that satisfies the following two properties: (i) each agent receives exactly r items, and (ii) each agent has utility at least $\tau := 1 - 2 \log n/n$ for every item that she receives. If at least one such allocation exists, the algorithm outputs any of them. Else, it outputs NULL. Note that determining whether such an allocation exists and finding one if it exists can be done in polynomial time by reducing to matching: we create a bipartite graph $(N \times [r], M, E)$, where $((i, \ell), j) \in E$ if and only if $u_i(j) \geq \tau$. A desired allocation corresponds to a perfect matching in this graph.³

For the sake of convenience, we introduce the notion of *r-matching*, which allows us to focus on the graph with vertex set N instead of $N \times [r]$. In an *r-matching*, each left vertex can be matched to as many as r right vertices, whereas each right vertex is still allowed to be matched to at most one left vertex.

Definition 3.3. An *r-matching* of a bipartite graph G is a subgraph of G such that every left vertex has degree at most r and every right vertex has degree at most 1. An *r-matching* is said to be *perfect* if every left vertex has degree exactly r and every right vertex has degree exactly 1.

As with normal matchings, a perfect *r-matching* can be computed in polynomial time by creating r copies of each left vertex and finding a perfect matching. With this definition, our simplified algorithm can be described as follows.

Algorithm 1 Simplified Algorithm for $r \geq 3$

```

1: procedure THRESHOLDMATCHING $_{\tau}(N, M, \{u_i\}_{i \in [n]})$ 
2:   for  $i = 1, 2, \dots, n$  do
3:      $M_{\geq \tau}(i) \leftarrow \{j \in M \mid u_i(j) \geq \tau\}$ 
4:   Let  $G_{\geq \tau} = (N, M, E_{\geq \tau})$  denote the graph where
      $(i, j) \in E_{\geq \tau}$  iff  $j \in M_{\geq \tau}(i)$ .
5:   if  $G_{\geq \tau}$  contains a perfect r-matching then
6:     return any perfect r-matching of  $G_{\geq \tau}$ 
7:   else
8:     return NULL

```

We now sketch the proof of correctness of Algorithm 1, which consists of two parts. Firstly, we argue that with high probability, the algorithm returns a perfect *r-matching* in $G_{\geq \tau}$ (i.e., does not output NULL). Secondly, we show that the output allocation is envy-free with high probability.

Existence of a perfect *r-matching* in $G_{\geq \tau}$. For the first part, we evoke a classical result regarding the existence of a perfect matching in bipartite random graphs. Recall that

³We write (U, V, E) to denote a bipartite graph with the set of vertices U and V in the partition, which we refer to as the set of *left vertices* and *right vertices* respectively, and the set of edges E .

for any positive integers a, b and any $p \in [0, 1]$, a bipartite graph sampled from the *Erdős-Rényi random bipartite graph distribution* $\mathcal{G}(a, b, p)$ consists of left and right vertex sets A and B of size a and b respectively, and for any pair of vertices $a \in A$ and $b \in B$, the edge (a, b) occurs with probability p independently of other pairs of vertices.

Proposition 3.4 (Erdős and Rényi (1964)). *Let G be a bipartite graph sampled from the Erdős-Rényi random bipartite graph distribution $\mathcal{G}(n, n, p)$, where $p = (\log n + \omega(1))/n$. Then, with high probability, G contains a perfect matching.*

To show that a perfect r -matching is likely to exist in $G_{\geq \tau}$, we arbitrarily partition the item set M into r parts $M^{(1)}, \dots, M^{(r)}$, each of size n . We also create a bipartite graph $H^{(a)}$ for $a = 1, 2, \dots, r$ where the left vertex set is N , the right vertex set is $M^{(a)}$, and each (i, j) is an edge iff $u_i(j) \geq \tau$. Now, since $\tau = 1 - 2 \log n/n$, for each a the graph $H^{(a)}$ is distributed according to the Erdős-Rényi random bipartite graph distribution $\mathcal{G}(n, n, 2 \log n/n)$. As a result, Proposition 3.4 implies that $H^{(a)}$ contains a perfect matching with high probability. By taking the union of the perfect matchings in $H^{(1)}, \dots, H^{(r)}$, we conclude that $G_{\geq \tau}$ contains a perfect r -matching with high probability. This completes the first part of the proof sketch.

Envy-freeness of output allocation. Next, we argue that with high probability, any allocation output by Algorithm 1 is envy-free. Consider any such allocation. Since every agent receives r items, each of which yields utility at least τ to her, her total utility is at least $r \cdot \tau = r - 2r \log n/n$. It therefore suffices to show that with high probability, for every $i' \neq i$, the utility of agent i' for agent i 's bundle M_i is at most $r - 2r \log n/n$. We will show that with high probability, for every $i' \neq i$, agent i' values at most $r - 1$ items in M_i more than $1 - 2r \log n/n$. This is sufficient because these $r - 1$ items can each contribute utility at most 1 to agent i' , whereas the remaining item contributes utility at most $1 - 2r \log n/n$ to her. It follows that the utility of agent i' for M_i does not exceed $(r - 1) + (1 - 2r \log n/n) = r - 2r \log n/n$.

Fix two distinct $i, i' \in [n]$. Let $E_{i, i'}$ denote the “bad” event that there exist r items j_1, \dots, j_r for which $u_i(j_k) \geq \tau$ and $u_{i'}(j_k) \geq 1 - 2r \log n/n$ for $k = 1, 2, \dots, r$. Consider any item $j \in M$. Since we assume that $u_i(j)$ and $u_{i'}(j)$ are drawn independently from the uniform distribution on $[0, 1]$, the probability that item j satisfies the two inequalities above for i and i' is at most $\frac{2 \log n}{n} \cdot \frac{2r \log n}{n} = \frac{4r \log^2 n}{n^2}$. Using the union bound over all subsets of r items, we have

$$\begin{aligned} \Pr[E_{i, i'}] &\leq \binom{m}{r} \cdot \left(\frac{4r \log^2 n}{n^2} \right)^r \\ &\leq \left(\frac{4r^2 \log^2 n}{n} \right)^r = o(n^{-2}), \end{aligned}$$

where we use the inequality $\binom{m}{r} \leq m^r$ and the assumption that $r \geq 3$ is constant. Applying the union bound again over all i, i' , the probability that at least one bad event occurs is $o(1)$. This concludes our proof sketch for the simplified algorithm.

3.3 The Algorithm

Having described the simplified algorithm, we now proceed to the actual algorithm. Before we do so, let us note that Algorithm 1 does *not* work for $r = 2$. This is because when $r = 2$, there is a constant probability that some pair of agents have the same two most valued items. In this case, Algorithm 1 could output an allocation that assigns both items to one of the two agents, which would mean that this agent is envied by the other agent.

To make Algorithm 1 work for $r = 2$, recall that the algorithm could fail if it is possible to find r items in the candidate item set of agent i (i.e., the set of items for which agent i has utility at least τ) that another agent i' values more than $r \cdot \tau$ in total. The modification to the algorithm is simple: remove any such problematic items from the candidate set of i before we try to find a perfect r -matching in the graph.

There are multiple ways to implement this removal step. The way we use, which we feel is quite natural, is to continue removing from the candidate set of agent i an item which agent i' values the most, until the r items in the candidate set of i that are most highly valued by i' are not valued more than $r \cdot \tau$ in total. The pseudo-code of the algorithm is presented below as Algorithm 2; here we use $\text{sum-top}_r(S)$ to denote the sum of the r largest elements of S for any *multiset* of real numbers S (or the sum of all elements if S contains less than r elements). The set in line 5 of the algorithm is considered as a multiset. The appropriate value of τ depends on the distribution \mathcal{U} and will be specified later.

Algorithm 2 Algorithm for any $r \geq 2$

```

1: procedure THRESHOLDMATCHINGWITHREMOVAL
   ( $N, M, \{u_i\}_{i \in [n]}$ )
2: for  $i = 1, 2, \dots, n$  do
3:    $M_{\geq \tau}^*(i) \leftarrow \{j \in M \mid u_i(j) \geq \tau\}$ 
4:   for  $i' \in [n] \setminus \{i\}$  do
5:     while  $\text{sum-top}_r(\{u_{i'}(j) \mid j \in M_{\geq \tau}^*(i)\}) > r \cdot \tau$  do
6:        $M_{\geq \tau}^*(i) \leftarrow M_{\geq \tau}^*(i) \setminus \arg \max_{j \in M_{\geq \tau}^*(i)} u_{i'}(j)$ 
7: Let  $G_{\geq \tau}^* = (N, M, E_{\geq \tau}^*)$  denote the graph where
    $(i, j) \in E_{\geq \tau}^*$  iff  $j \in M_{\geq \tau}^*(i)$ .
8: if  $G_{\geq \tau}^*$  contains a perfect  $r$ -matching then
9:   return any perfect  $r$ -matching of  $G_{\geq \tau}^*$ 
10: else
11: return NULL

```

The above modification ensures that if Algorithm 2 returns an allocation, it must be envy-free. Indeed, each agent i has utility at least τ for every item assigned to her in the r -matching, so her total utility is at least $r \cdot \tau$. On the other hand, by construction of the graph $G_{\geq \tau}^*$, each agent i' values the r items assigned to agent i at most $r \cdot \tau$. Thus, the output allocation must be envy-free.

In order to establish Theorem 3.1, it therefore remains to show that with an appropriate choice of τ , a perfect r -matching in $G_{\geq \tau}^*$ exists with high probability. Recall our assumption that the distribution \mathcal{U} from which the utilities are drawn is polynomially bounded at 1. Let $\underline{\theta}, \bar{\theta}, q > 0$ be

the associated parameters. It suffices to prove the following lemma:

Lemma 3.5. *Set $\tau := 1 - \left(\frac{64 \log m}{\theta n}\right)^{1/q}$ in Algorithm 2, and let $2 \leq r \leq e^{n^{0.1}}$. Then, with high probability, the graph $G_{\geq \tau}^*$ contains a perfect r -matching.*

Note that the condition $r \leq e^{n^{0.1}}$ implies that $\tau > 0$ for large enough n .

The proof of Lemma 3.5 consists of two parts. First, in Section 3.4, we show that only few edges are removed in line 6 of Algorithm 2; in particular, we show that with high probability, at most two edges adjacent to any particular vertex are removed. Then, in Section 3.5, we show that the existence of a perfect r -matching is *locally resilient* (see, e.g., (Sudakov and Vu, 2008)) in the following sense: even if we remove a low-degree subgraph from a random graph sampled from the Erdős-Rényi random bipartite graph distribution with sufficiently large probability, then the remaining graph still contains a perfect r -matching with high probability. Putting these two parts together yields Lemma 3.5; this is done in Section 3.6.

Before we proceed to proving Lemma 3.5, we perform some preliminary calculations. Picking $\alpha = 1 - \tau$ in Definition 2.1, we have $\Pr_{u \sim \mathcal{U}}[u > \tau] \geq \frac{64 \log m}{n}$. On the other hand, writing $\tau' := 3\tau - 2$ and letting $\alpha = 1 - \tau'$ in Definition 2.1 yields $\Pr_{u \sim \mathcal{U}}[u > \tau'] \leq \theta(3(1 - \tau))^q = \frac{C \log m}{n}$ for the constant $C := 3^q \cdot \frac{64\theta}{\theta}$.

3.4 Bounding the Number of Edges Removed

Let $E_{\geq \tau}$ and $E_{\geq \tau}^*$ denote the set of edges as defined in Algorithm 1 and 2, respectively. The main result of this subsection is the following lemma:

Lemma 3.6. *With high probability, the graph $(N, M, E_{\geq \tau} \setminus E_{\geq \tau}^*)$ has maximum degree at most 2.*

In addition to the graphs $G_{\geq \tau}$ and $G_{\geq \tau}^*$ (as defined in Algorithm 1 and 2 respectively), we consider the graph $G_{> \tau'} = (N, M, E_{> \tau'})$ which can be defined analogously. That is, the neighbor set of $i \in N$ in $G_{> \tau'}$ is $M_{> \tau'}(i) := \{j \in M \mid u_i(j) > \tau'\}$.

The next proposition states that, for any edge (i, j) that is removed in line 6 of Algorithm 2, the edge (i, j) must be part of a complete bipartite subgraph $K_{2, \lfloor 2r/3 \rfloor + 1}$ of the graph $G_{> \tau'}$.⁴ We note that this is similar to the argument for the simplified algorithm in Section 3.2, which, in the new language, states that any edge (i, j) that would be removed in line 6 of Algorithm 2 must be part of a complete bipartite subgraph $K_{2, r}$ of the graph $G_{> \tau'}$, where the threshold τ' there is chosen (differently) to be $1 - 2r \log n/n$.

Proposition 3.7. *If $(i, j) \in E_{\geq \tau} \setminus E_{\geq \tau}^*$, then there exists i' such that $|M_{> \tau'}(i) \cap M_{> \tau'}(i')| > 2r/3$ and $j \in M_{> \tau'}(i) \cap M_{> \tau'}(i')$.*

Proof. First, let us argue that $j \in M_{> \tau'}(i) \cap M_{> \tau'}(i')$. The assumption that $(i, j) \in E_{\geq \tau}$ immediately implies

⁴The notation $K_{a,b}$ refers to a complete bipartite graph with left and right vertex sets of size a and b , respectively.

that $j \in M_{> \tau'}(i)$ since $\tau' = 3\tau - 2 < \tau$. Let $i' \in N$ be the vertex in line 5 of Algorithm 2 that causes the removal of the edge (i, j) . At this line, we have $\text{sum-top}_r(\{u_{i'}(j') \mid j' \in M_{\geq \tau}^*(i)\}) > r \cdot \tau$ and $j \in \arg \max_{j' \in M_{\geq \tau}^*(i)} u_{i'}(j')$. Hence, $u_{i'}(j) \geq (r \cdot \tau)/r = \tau > \tau'$, and therefore $j \in M_{> \tau'}(i')$. We have thus shown that $j \in M_{> \tau'}(i) \cap M_{> \tau'}(i')$.

Next, let $y = \min\{r, |M_{> \tau'}(i) \cap M_{> \tau'}(i')|\}$. The condition $\text{sum-top}_r(\{u_{i'}(j') \mid j' \in M_{\geq \tau}^*(i)\}) > r \cdot \tau$ at the time when the edge (i, j) is removed implies that

$$\begin{aligned} r \cdot \tau &< \text{sum-top}_r(\{u_{i'}(j') \mid j' \in M_{\geq \tau}^*(i)\}) \\ &\leq \text{sum-top}_r(\{u_{i'}(j') \mid j' \in M_{> \tau'}(i)\}) \\ &\leq y \cdot 1 + (r - y) \cdot \tau' \\ &= y + (r - y)(3\tau - 2) \\ &= (3 - 3\tau)y + (3\tau - 2)r, \end{aligned}$$

which implies that $y > 2r/3$, as claimed. \square

We now use Proposition 3.7 to prove Lemma 3.6. We consider two cases: $r \geq 3$ and $r = 2$.

The Case $r \geq 3$. The proof for the case $r \geq 3$ is similar to that of the simplified algorithm (Section 3.2). In particular, we show that with high probability, no edge is removed in Algorithm 2. This also means that Algorithms 1 and 2 are equivalent with high probability.

Proposition 3.8. *Let $3 \leq r \leq e^{n^{0.1}}$. Then, with high probability, $E_{\geq \tau} = E_{\geq \tau}^*$.*

Proof. For convenience, let $p_{> \tau'} := \Pr_{u \sim \mathcal{U}}[u > \tau']$. Note that $p_{> \tau'} \leq \frac{C \log m}{n} \leq C_0/n^{0.9}$ for $C_0 := 2C$, where the last inequality holds for sufficiently large n . We will argue that with high probability, there are no distinct $i, i' \in N$ such that $|M_{> \tau'}(i) \cap M_{> \tau'}(i')| > 2r/3$. Together with Proposition 3.7, this implies that no edge is removed and therefore $E_{\geq \tau} = E_{\geq \tau}^*$.

To show this, we use the standard first moment method. Fix distinct $i, i' \in N$ and a subset $S \subseteq M$ of size $x := \lfloor 2r/3 \rfloor + 1$. The probability that $S \subseteq M_{> \tau'}(i) \cap M_{> \tau'}(i')$ is exactly $(p_{> \tau'})^{2x}$. Hence, by taking the union bound over all choices of i, i' and S , the probability that $|M_{> \tau'}(i) \cap M_{> \tau'}(i')| > 2r/3$ for some i, i' is at most

$$\begin{aligned} n^2 \binom{m}{x} (p_{> \tau'})^{2x} &\leq n^2 \left(\frac{em}{x}\right)^x (p_{> \tau'})^{2x} \\ &(\text{since } x \geq 3) \leq (n(p_{> \tau'})^{1.2})^2 \left(\frac{em(p_{> \tau'})^{1.2}}{x}\right)^x \\ &(\text{since } x > 2r/3) < (n(p_{> \tau'})^{1.2})^2 (1.5en(p_{> \tau'})^{1.2})^x \\ &(\text{since } p_{> \tau'} \leq \frac{C_0}{n^{0.9}}) \leq (C_0^{1.2} n^{-0.08})^2 (1.5eC_0^{1.2} n^{-0.08})^x \\ &= o(1), \end{aligned}$$

which concludes the proof. \square

The Case $r = 2$. As argued earlier, in the case $r = 2$, some edges must be removed in order to guarantee that the output allocation is envy-free. The following proposition ensures that with high probability, for any vertex, at most two edges adjacent to it are removed in Algorithm 2.

Proposition 3.9. *Let $r = 2$. Then, with high probability, the graph $(N, M, E_{\geq \tau} \setminus E_{\geq \tau}^*)$ has maximum degree at most 2.*

Proof. Observe that, for $r = 2$, Proposition 3.7 can be restated as follows: if $(i, j) \in E_{\geq \tau} \setminus E_{\geq \tau}^*$, then there exist $i' \in N$ and $j' \in M$ such that i, i', j, j' form a complete bipartite graph $K_{2,2}$ in the graph $G_{> \tau'}$.

Now, suppose that some vertex $u \in N \cup M$ appears in at least three edges in $E_{\geq \tau} \setminus E_{\geq \tau}^*$. The previous paragraph implies each such edge must be contained in a copy of $K_{2,2}$ in the graph $G_{> \tau'}$. Since the three edges from u are distinct, not all three of these copies can be identical. As a result, u must be contained in two different copies of $K_{2,2}$, which means that at least one of the graphs shown in Figure 1 must appear as a subgraph of $G_{> \tau'}$. Notice that by the union bound, for any graph $H = (V_H, E_H)$, the probability that it appears as a subgraph of $G_{> \tau'}$ is at most $(n + m)^{|V_H|} (p_{> \tau'})^{|E_H|} \leq (3n)^{|V_H|} \left(\frac{C \log m}{n}\right)^{|E_H|}$. However, all graphs H in Figure 1 satisfy $|E_H| \geq |V_H| + 1$. Hence, the probability that each of them appears as a subgraph is at most $\frac{(C \log n)^{O(1)}}{n} = o(1)$. Using the union bound, the probability that at least one of these graphs appears as a subgraph of $G_{> \tau'}$ is also $o(1)$. This implies that the probability that at least one of the vertices is adjacent to more than two edges in $E_{\geq \tau} \setminus E_{\geq \tau}^*$ is $o(1)$, as desired. \square

Finally, we note that Propositions 3.8 and 3.9 together imply Lemma 3.6.

3.5 Local Resilience of Perfect r -Matching

In this subsection, we show that in a random bipartite graph sampled from the Erdős-Rényi random bipartite graph distribution $\mathcal{G}(n, rn, p)$ with sufficiently large p , not only does a perfect r -matching exist, but the existence is also robust in the following sense: even if we remove edges from the graph, as long as not too many edges adjacent to each vertex are removed, a perfect r -matching still exists. Such “robustness” is known in the literature as *local resilience*. In particular, the local resilience of perfect matchings was shown by Sudakov and Vu (2008). We will extend their proof to the case of perfect r -matchings. However, we note that our bound will be slightly weaker than theirs, since our main goal is to derive a bound that is sufficient for the algorithm to work and not to find the best possible parameters.

A typical method for establishing the existence of a perfect matching, which was used both by Sudakov and Vu (2008) and by Erdős and Rényi (1964), is to show that the graph satisfies the condition of Hall’s Marriage Theorem. For any graph G and any set S of vertices in G , denote by $N_G(S)$ the set of vertices adjacent to at least one vertex in S .

Proposition 3.10 (Hall’s Marriage Theorem). *Let $G = (A, B, E)$ be any bipartite graph such that $|A| = |B|$. If*

$|N_G(S)| \geq |S|$ for all subsets $S \subseteq A$, then G has a perfect matching.

Recall that $G = (A, B, E)$ has a perfect r -matching if and only if the graph $(A \times [r], B, E')$, where $((a, \ell), b) \in E'$ iff $(a, b) \in E$, has a perfect matching. Hence, Hall’s Marriage Theorem immediately extends to r -matchings:

Proposition 3.11. *Let $G = (A, B, E)$ be any bipartite graph such that $|B| = r|A|$. If $|N_G(S)| \geq r|S|$ for all subsets $S \subseteq A$, then G has a perfect r -matching.*

One way to show that the condition of Hall’s Marriage Theorem is satisfied is to show that there is at least one edge between any sets $S \subseteq A$ and $T \subseteq B$ of appropriate sizes. To ensure that the existence of a perfect r -matching is locally resilient, we need to show not only that one edge exists, but also that many edges exist. This can be done via standard concentration bounds. For any graph G and any sets S, T of vertices in G , denote by $E_G(S, T)$ the set of edges connecting a vertex in S to a vertex in T .

Lemma 3.12. *Let $G = (A, B, E)$ be a graph sampled from the Erdős-Rényi random bipartite graph distribution $\mathcal{G}(n, m, p)$ with $p \geq \frac{64 \log m}{n}$. Then, with high probability, the following holds for all subsets $S \subseteq A$ and $T \subseteq B$ such that $|T| = m - r|S| + 1$:*

$$|E_G(S, T)| > (16 \log m) \cdot \min\{|S|, |T|\}. \quad (1)$$

The proof of Lemma 3.12 can be found in the full version of this paper (Manurangsi and Suksompong, 2018).

With Lemma 3.12 ready, we now establish the local resilience of the existence of perfect r -matchings in random graphs.

Lemma 3.13. *Let $G = (A, B, E)$ be a graph sampled from the Erdős-Rényi random bipartite graph distribution $\mathcal{G}(n, m, p)$ with $p \geq \frac{64 \log m}{n}$. Then, with high probability, for any subgraph $H = (A, B, E')$ of G with maximum degree at most $16 \log m$, the graph $G - H = (A, B, E \setminus E')$ contains a perfect r -matching.*

Proof. From Lemma 3.12, with high probability, (1) holds for all $S \subseteq A, T \subseteq B$ with $|T| = m - r|S| + 1$. We claim that this implies that $G - H = (A, B, E \setminus E')$ contains a perfect r -matching. Suppose for the sake of contradiction that $G - H$ does not contain an r -perfect matching. Proposition 3.11 implies that there exists a set $S \subseteq A$ such that $|N_{G-H}(S)| \leq r|S| - 1$. Let T be any subset of $B \setminus N_{G-H}(S)$ of size $m - r|S| + 1$. Since $T \cap N_{G-H}(S) = \emptyset$, we have $E_{G-H}(S, T) = \emptyset$. However, by (1), $|E_G(S, T)| > (16 \log m) \cdot \min\{|S|, |T|\}$. This means that at least one vertex in $S \cup T$ has degree more than $16 \log m$ in H , which is a contradiction. \square

3.6 Putting Things Together

With Lemmas 3.6 and 3.13 in hand, we can (finally) prove Lemma 3.5.

Proof of Lemma 3.5. First, Lemma 3.6 ensures that with high probability, for each vertex, at most two edges adjacent to it are removed from $G_{\geq \tau}$ in Algorithm 2, where $G_{\geq \tau}$ is




Figure 1: All possible unions of two distinct complete bipartite graphs $K_{2,2}$ which share at least one vertex u , up to isomorphism. The shaded vertices constitute one copy of $K_{2,2}$ whereas the thickened vertices constitute another.

defined as in Algorithm 1. Recall also that $G_{\geq \tau}$ is distributed according to the Erdős-Rényi random bipartite graph distribution $\mathcal{G}(n, m, p)$ with $p = \Pr_{u \sim \mathcal{U}}[u \geq \tau] \geq \frac{64 \log m}{n}$. It therefore follows from Lemma 3.13 that a perfect r -matching exists in $G_{\geq \tau}^*$ with high probability. \square

4 Non-Existence

Our main non-existence result states that envy-free allocations are unlikely to exist if $m = O(n \log n / \log \log n)$ is not “close to” being a multiple of n . This improves upon the $m = n + o(n)$ lower bound of Dickerson et al. (2014) and comes to within a $\Theta(\log \log n)$ factor of matching their upper bound.

Theorem 4.1. *For any real numbers $\theta > 0$, $\epsilon \in (0, 1)$, and $q \geq 1$, there exists $c > 0$ depending only on θ, ϵ, q such that the following holds: For any positive integer $r \leq \frac{c \log n}{\log \log n}$, if $m \in [rn + n^\epsilon, (r + 1)n - n^\epsilon]$ and \mathcal{U} is (θ, q) -polynomially bounded below at 1, then, with high probability, there is no envy-free allocation.*

We remark that since we only require the distribution to be polynomially bounded below, the assumption $q \geq 1$ does not introduce a loss of generality. Next, we give an overview of the proof of Theorem 4.1; the full proof can be found in the full version of this paper (Manurangsi and Suksompong, 2018).

The proof is based on the first moment method; the key is to show that for any fixed allocation, the probability (over the random utilities drawn) that it is envy-free is $\ll 1/n^m$. Since there are n^m possible allocations, the union bound implies that with high probability, no envy-free allocation exists.

To give an intuition for this bound, let us consider a simplified setting where $m = (r + 0.5)n$ and the distribution \mathcal{U} is uniform on $[0, 1]$. Intuitively, the “more balanced” the allocation is, the harder it is to bound the probability that the allocation is envy-free. Following this intuition, let us consider the “most balanced” allocation where $0.5n$ agents receive $r + 1$ items and the remaining agents receive r items. The key observation is that, for the allocation to be envy-free for every agent in the latter group, any such agent must have utility at most r for the $r + 1$ items in the bundle of any agent in the first group. For a fixed agent in the second group and a fixed agent in the first group, this happens with probability

at most $1 - 1/(r + 1)^{r+1}$. Indeed, if each of the $r + 1$ items yields utility at least $r/(r + 1)$ to the agent, the requirement is not satisfied. Now, since there are $0.25n^2$ such pairs of agents, the probability that this fixed allocation is envy-free

is at most $\left(1 - \frac{1}{(r+1)^{r+1}}\right)^{0.25n^2} = \exp\left(\Theta\left(\frac{-n^2}{(r+1)^{r+1}}\right)\right)$.

Hence, as long as $r \ll \log n / \log \log n$, this term is at most, say, $\exp(-n^{1.9})$, which is indeed much smaller than n^{-m} .

The full proof proceeds along the lines of the argument above, but we need to be more careful as we must also deal with other “less balanced” allocations.

5 Discussion

In this paper, we study the existence and non-existence of envy-free allocations and essentially close the gap left open by Dickerson et al. (2014) with regard to the transition between the two phases. On the positive side, we show that if the number of items is a multiple of the number of agents, an envy-free allocation is likely to exist as long as the former quantity is at least twice the latter. On the negative side, we show that if the number of items is not “close to” being a multiple of the number of agents, an envy-free allocation is unlikely to exist even when the former quantity exceeds the latter by almost a logarithmic factor. Both of our results make use of several new ideas that may be useful for other problems in fair division.

As we mentioned earlier, all of the works on the asymptotic existence of fair allocations thus far have assumed that agents are endowed with additive utilities. While additivity provides a reasonable trade-off between simplicity and expressiveness, it would be interesting to establish analogous results that hold for more general classes of utilities. Going beyond additivity introduces several complications; for example, the welfare-maximizing allocation is no longer simply the one that assigns every item to the agent who values it most, and giving an agent several goods that she values highly does not guarantee that the agent will also have a correspondingly high value for the whole bundle. Nevertheless, a starting point may be to prove results for specific distributions over utilities from a well-structured class such as that of submodular valuations.

Another possible avenue for future work is to consider the setting where instead of allocating items to individual agents, we divide them among *groups* of agents (Segal-

Halevi and Nitzan, 2016; Segal-Halevi and Suksompong, 2018; Suksompong, 2018a,b). The agents in each group share the same set of items but may have different preferences. This is the case, for example, when dividing household goods among families or resources between departments in a university. Manurangsi and Suksompong (2017) generalized the results of Dickerson et al. (2014) to the group setting and left a logarithmic gap between existence and non-existence. We are hopeful that the techniques we introduce in the present work will help towards closing this gap as well.

Acknowledgments

This work was partially supported by NSF Awards CCF-1655215, CCF-1813188, CCF-1815434, and by a Stanford Graduate Fellowship. We would like to thank the anonymous reviewers for their helpful comments.

References

- Amanatidis, G.; Markakis, E.; Nikzad, A.; and Saberi, A. 2017. Approximation algorithms for computing maximin share allocations. *ACM Transactions on Algorithms* 13(4):52.
- Amanatidis, G.; Birmpas, G.; and Markakis, E. 2018. Comparing approximate relaxations of envy-freeness. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 42–48.
- Barman, S.; Krisnamurthy, S. K.; and Vaish, R. 2018. Finding fair and efficient allocations. In *Proceedings of the 19th ACM Conference on Economics and Computation*, 557–574.
- Biswas, A., and Barman, S. 2018. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 91–97.
- Brams, S. J., and Taylor, A. D. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2016. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation*, 305–322.
- Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation*, 629–646.
- Dickerson, J. P.; Goldman, J.; Karp, J.; Procaccia, A. D.; and Sandholm, T. 2014. The computational rise and fall of fairness. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, 1405–1411.
- Erdős, P., and Rényi, A. 1961. On a classical problem of probability theory. *Magyar Tudományos Akadémia Matematikai Kutató Intézetének Közleményei* 6:215–220.
- Erdős, P., and Rényi, A. 1964. On random matrices. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences* 8:455–461.
- Foley, D. K. 1967. Resource allocation and the public sector. *Yale Economics Essays* 7(1):45–98.
- Kurokawa, D.; Procaccia, A. D.; and Wang, J. 2016. When can the maximin share guarantee be guaranteed? In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, 523–529.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, 125–131.
- Manurangsi, P., and Suksompong, W. 2017. Asymptotic existence of fair divisions for groups. *Mathematical Social Sciences* 89:100–108.
- Manurangsi, P., and Suksompong, W. 2018. When do envy-free allocations exist? *CoRR* abs/1811.01630.
- Moulin, H. 2003. *Fair Division and Collective Welfare*. MIT Press.
- Plaut, B., and Roughgarden, T. 2018. Almost envy-freeness with general valuations. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2584–2603.
- Segal-Halevi, E., and Nitzan, S. 2016. Envy-free cake-cutting among families. *CoRR* abs/1607.01517.
- Segal-Halevi, E., and Suksompong, W. 2018. Democratic fair allocation of indivisible goods. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 482–488.
- Stromquist, W. 1980. How to cut a cake fairly. *The American Mathematical Monthly* 87(8):640–644.
- Sudakov, B., and Vu, V. H. 2008. Local resilience of graphs. *Random Structures and Algorithms* 33(4):409–433.
- Suksompong, W. 2016. Asymptotic existence of proportionally fair allocations. *Mathematical Social Sciences* 81:62–65.
- Suksompong, W. 2018a. Approximate maximin shares for groups of agents. *Mathematical Social Sciences* 92:40–47.
- Suksompong, W. 2018b. *Resource Allocation and Decision Making for Groups*. Ph.D. Dissertation, Stanford University.
- Varian, H. R. 1974. Equity, envy, and efficiency. *Journal of Economic Theory* 9:63–91.