

On the Existence of Envy-Free Allocations Beyond Additive Valuations

Gerdus Benadè

Questrom School of Business, Boston University, benade.at.bu.edu

Daniel Halpern

School of Engineering and Applied Sciences, Harvard University, dhalpern@g.harvard.edu

Alexandros Psomas

Department of Computer Science, Purdue University, apsomas@purdue.edu

Paritosh Verma

Department of Computer Science, Purdue University, verma136@purdue.edu

We study the problem of fairly allocating m indivisible items among n agents. Envy-free allocations, wherein each agent prefers her bundle to the bundle of every other agent, need not exist in the worst case. However, when agents have additive preferences and each agent's value for each item is drawn independently from a distribution, envy-free allocations exist with high probability when $m \in \Omega(n \log n / \log \log n)$. In this paper, we study the existence of envy-free allocations under stochastic valuations far beyond the additive setting. We introduce a new stochastic model where each agent's valuation is sampled by first fixing a worst-case function, and then drawing a uniformly random renaming of the items, independently for each agent; this generalizes various known models. We prove that this random renaming is sufficient to ensure that envy-free allocations exist with high probability in very general settings. For entirely arbitrary valuations and two agents, we prove that envy-free allocations exist with high probability. When valuations are non-negative and "order-consistent," a class that generalizes additive, budget-additive, unit-demand, and single-minded agents, we show SD-envy-free allocations exist for $m \in \omega(n^2)$ when n divides m , and SD-EFX allocations exist for all $m \in \omega(n^2)$.

Key words: Fair Division, Envy-freeness, Beyond Additive Valuations, Hypergraph Isomorphism

1. Introduction

The fair division of constrained resources has been extensively studied by several fields including economics, computer science, and operations research. In the classical *cake cutting* model (Steinhaus 1948), a continuum must be divided fairly between agents with heterogeneous valuation functions; since then, many variations of the problem have been studied in settings differing in divisibility, dynamics, side constraints and fairness notions, to name just a few. Applications of fair allocation problems include dividing assets in an inheritance or divorce settlement (Pratt and Zeckhauser 1990, Brams and Taylor 1996), organ transplantation (Bertsimas et al. 2013), allocating food to individuals or food pantries (Sinclair et al. 2023, Mertzanidis et al. 2024, Alptekinoglu and

Benadè 2023), resource allocation in shared environments like data centers (Butler and Williams 2002, Armony and Ward 2010) and redistricting (Sherstyuk 1998, Benadé et al. 2026).

The operations research literature offers a rich toolkit for incorporating fairness into optimization models. Explicit side constraints can enforce minimum guarantees or balanced service levels; objective-based criteria such as max-min or leximax prioritize the worst-off agents; inequality-averse welfare functions penalize dispersion across outcomes; and a growing body of work characterizes the efficiency-fairness frontier in various settings. Different application contexts demand different approaches; we refer the interested reader to Chen and Hooker (2023) for a more thorough discussion. In some settings, fairness is best understood as requiring equal treatment or ensuring that no group falls below an acceptable threshold. In others, it requires richer comparisons across agents with heterogeneous preferences. In such cases, *envy-freeness* is arguably the gold standard for fairness: it requires that every agent weakly prefers their own bundle over the bundle of any other agent.

We consider the problem of fairly dividing a set $\mathcal{M}$ of m indivisible items among a set $\mathcal{N}$ of n agents. Each agent i has a valuation function $v_i : 2^{\mathcal{M}} \mapsto \mathbb{R}$, which maps each subset of items $S \subseteq \mathcal{M}$ to a value for S . Our primary goal is to find an *envy-free* allocation. An allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ is envy-free if $v_i(A_i) \geq v_i(A_j)$ for all $i, j \in \mathcal{N}$. Because this is our objective, we require that all items be allocated; without this constraint, leaving all items unallocated would trivially achieve envy-freeness. It is easy to see that envy-free allocations do not exist in the worst case: consider a single item and two agents valuing it positively. Motivated to circumvent this simple non-existence result, a line of research in fair division studies the existence of envy-free allocations under stochastic valuations.

To date, this work has focused on additive valuation functions. Recall that a valuation function $v_i : 2^{\mathcal{M}} \mapsto \mathbb{R}$ is additive if $v_i(S \cup T) = v_i(S) + v_i(T)$ for all $S, T \subseteq \mathcal{M}$ with $S \cap T = \emptyset$; an additive function v_i can be succinctly represented with a value $v_{i,j}$ for each item $j \in \mathcal{M}$, such that $v_i(S) = \sum_{j \in S} v_{i,j}$ for all $S \subseteq \mathcal{M}$ and $i \in \mathcal{N}$. Dickerson et al. (2014) show that when agents have additive and non-negative valuation functions, and all item values are drawn independently from a distribution $\mathcal{D}$, allocations that simultaneously satisfy envy-freeness and Pareto efficiency exist with high probability for $m \in \Omega(n \log n)$. On the other hand, envy-free allocations do not exist with constant probability for $m \in n + o(n)$. In the same setting, Manurangsi and Suksompong (2020) show that an envy-free allocation exists with high probability as long as $m \geq 2n$ and n divides m . When m is not “almost divisible” by n , an envy-free allocation is unlikely to exist for $m \in O(n \log n / \log \log n)$.

Manurangsi and Suksompong (2021) close this gap by proving that envy-free allocations exist with high probability when $m \in \Omega(n \log n / \log \log n)$. More recently, Bai and Gözl (2022) extend these bounds to the additive non-i.i.d. case where item values are drawn independently from agent-specific distributions $\mathcal{D}_i$. However, despite all the interest in envy-free allocations under stochastic valuations, little is known beyond the additive model.

The goal of this paper is to study the existence of envy-free allocations under stochastic valuations beyond the additive case. In order to even pose this question, it is necessary to first specify a stochastic model for generating such valuations. The aforementioned additive model is very natural: for every agent, we simply need to sample a value for each item, then the value of any subset of items follows readily from additivity. But how should we go about sampling, for example, a submodular valuation function?

One may hope to draw inspiration from adjacent fields. The Bayesian setting is the dominant paradigm in mechanism design, where the worst-case lens fails to provide useful insights or meaningful guarantees for various fundamental problems including revenue maximization. One standard “beyond additive” model in Bayesian mechanism design, e.g. in the literature on prophet inequalities (see Lucier (2017) for a survey), assumes that valuation functions are drawn from distributions over families of functions (e.g. submodular or XOS functions). Another standard model, e.g. in auctions, is the “ C over independent items” model for a condition C (e.g. “subadditive over independent items”), introduced by Rubinstein and Weinberg (2018), where, informally, an agent’s valuation function is parameterized by a vector of types, drawn from a product distribution; see Rubinstein and Weinberg (2018), Rubinstein and Singla (2017), Cai and Zhao (2017), Chawla and Miller (2016) for more details about this model. Unfortunately, in both these models it is easy to pick a valuation function (or a distribution over valuation functions) such that one item has more value than all other items combined, in which case envy-free allocations trivially do not exist.

In order to study more general stochastic valuations, we need an entirely new approach.

1.1. Our Contributions

A model for general stochastic valuations. Our approach centers on a form of neutrality shared by existing stochastic models. When item values are drawn identically and independently from identical distributions (i.i.d.), there is no a priori “discrimination” between items. It is virtually impossible to make a single item significantly more valuable than all remaining items across all agents — the scenario above that trivially precludes envy-free allocations. This suggests some form of neutrality would be very useful when trying to establish a high-probability existence guarantee.

Motivated by this, we propose a framework that isolates this property by considering the *worst-case over neutral distributions*. Formally, we say a distribution $\mathcal{D}_i$ over valuation functions is neutral if it is invariant under item permutations. That is, for any permutation $\pi : \mathcal{M} \rightarrow \mathcal{M}$, a valuation $v_i \sim \mathcal{D}_i$ is identically distributed to v_i^π , where $v_i^\pi(S) = v_i(\pi^{-1}(S))$ for all $S \subseteq \mathcal{M}$. While we do not claim that such neutral distributions necessarily reflect realistic human preferences, this invariance holds implicitly in nearly all previous models in stochastic fair division (with the exception of (Farhadi et al. 2019), discussed later). By explicitly isolating it, we seek to understand if this structural property is the true driving force behind the existence of fair allocations. We ask: is this property alone sufficient? Do envy-free allocations exist with high probability when each agent i 's valuation is drawn independently from a neutral distribution $\mathcal{D}_i$, even in the worst case over the choice of these distributions?

Observe that any neutral distribution can be equivalently viewed as a mixture over the permutation orbits of fixed valuation functions. Therefore, taking the worst case over all neutral distributions is mathematically equivalent to the following generative process: first, an adversary fixes an arbitrary, worst-case valuation function $v_i : 2^{\mathcal{M}} \rightarrow \mathbb{R}$ for each agent i ; then, the items are renamed uniformly at random and independently across agents. Slightly more formally, after the worst-case valuations v_i are fixed, a uniformly random permutation $\pi_i : \mathcal{M} \rightarrow \mathcal{M}$ is sampled for each agent, and their resulting valuation function becomes $v_i^{\pi_i}$.

Unfortunately, it is not hard to see that neutral distributions alone cannot guarantee the existence of envy-free allocations. Consider, for example, an identical additive function, i.e., $v_{i,j} = 1$ for all $i \in \mathcal{N}, j \in \mathcal{M}$. Random renaming has no power and leaves the values unchanged. For this function, when n does not divide m , envy-free allocations do not exist: It is impossible to perfectly equalize the bundle sizes, so, in any allocation, agents with fewer items will envy those with more. Thus, there certainly must be some requirements depending on divisibility, but one might expect that several other additional conditions must be met for envy-free allocations to exist with high probability.

EF allocations under arbitrary valuations. For $n = 2$, we show this intuition is entirely false. In Theorem 1, we provide a tight characterization demonstrating that, strikingly, the identical additive instance is essentially the *only* valuation for which envy-free allocations fail to exist. When m is even, envy-free allocations exist with high probability; when m is odd (and the indivisibility of items prevents an equal partition), envy-free allocations fail to exist with high probability only under a mild generalization of uniform identical values, roughly, that agents prefer bundles of the same specific size. This failure mode is quite specific; for any other *arbitrarily general* valuation

functions, envy-free allocations reliably exist with probability at least $1 - \frac{1}{m/2+1}$. We impose no constraints on these valuations, they can be superadditive, non-monotone, or mix positive and negative bundle values. We match this with a simple instance where existence fails with probability $1/m$, proving the bound is tight up to a constant.

The proof of Theorem 1 relies on the insight that a valuation function can be represented as a hypergraph on m vertices, where a hyperedge indicates that the corresponding subset of items is preferred to its complement. In this representation, a random permutation of items' names corresponds to picking a random, isomorphic hypergraph. Whenever the hypergraphs corresponding to the two agents' valuation functions are not identical, an envy-free allocation exists: allocate i the bundle corresponding to the hyperedge present in i 's hypergraph and not j 's. Using the orbit-stabilizer theorem, we reduce our question about the existence of envy-free allocations to a question about the number of automorphisms of k -uniform¹ hypergraphs. In Theorem 3, which may be of independent interest, we establish the main technical ingredient required for Theorem 1 and prove that the number of automorphisms of a k -uniform hypergraph on m vertices is at most $\frac{m!}{m-k+1}$. This is generalized to arbitrary hypergraphs under a mild condition in Corollary 1.

EF allocations under order-consistent valuations. In Section 4 we consider a structured valuation class that strictly generalizes additive valuations. Given an order over the items π , we say that a subset of items A stochastically dominates a subset of items B , denoted as $A \succeq_{\pi}^{\text{sd}} B$, when the best item in A (according to π) is weakly better than the best item in B , the second best item in A is weakly better than the second best item in B , and so on. We say that a valuation function v is *order-consistent* with respect to π if, for all bundles $A, B \subseteq \mathcal{M}$, $A \succeq_{\pi}^{\text{sd}} B$ implies that $v(A) \geq v(B)$. By picking π to be the items sorted in order of decreasing value, it is clear that additive valuations are order-consistent. Order-consistent valuations include budget additive valuations (a valuation is budget additive with budget B if each item j has value v_j and the value of $S \subseteq \mathcal{M}$ is $v(S) = \min\{B, \sum_{j \in S} v_j\}$), single-minded valuations (a valuation is single-minded if there is a subset S^* , such that $v(S) = v(S^*) > 0$ for all $S \supseteq S^*$, and $v(S) = 0$ otherwise), and unit-demand valuations (a valuation is unit-demand if there is a value v_j for each item j , and the value for a subset of items S is equal to $\max_{j \in S} v_j$). The class of order-consistent valuation functions, also studied in Bouveret et al. (2010), is incomparable with commonly studied large valuation classes; for example, there exist order-consistent valuations that are not subadditive.

¹ A hypergraph $H = (\mathcal{M}, E)$ is k -uniform if all of its edges are of size k , i.e., for all $e \in E$, we have $|e| = k$.

We prove that, given n agents with arbitrary order-consistent valuation functions over m items, where m is divisible by n , the probability that an envy-free allocation exists after random renaming is at least $1 - O\left(\frac{n^2}{m} + \frac{n \log m}{m^{\frac{n-1}{n}}}\right)$ (Theorem 4). In fact, we prove the existence of a (much) stronger notion of fairness, SD-envy-freeness (Bogomolnaia and Moulin 2001): an allocation $\mathcal{A}$ is SD-envy-free if, for all agents $i, j \in \mathcal{N}$, $A_i \succeq_{\pi_i}^{\text{sd}} A_j$. Note that, if an allocation is SD-envy-free, then it is envy-free for all additive utility functions consistent with the agents' preferences (Aziz et al. 2015). As a corollary to Theorem 4, we get that for $m \in \omega(n^2)$, SD-envy-free allocations exist with high probability. Our proof of Theorem 4 is constructive; we show that a simple Round-Robin process (agents take turns picking the best, according to their order, available item) produces such an allocation.

Theorem 4 is tight both in terms of the divisibility assumption and the bound on m . Standard birthday paradox arguments imply (even for unit-demand valuations) that envy-free allocations may not exist with constant probability for $m \in O(n^2)$. Moreover, m being divisible by n is a necessary condition for SD-envy-freeness because items may have equal value and SD-EF requires equal bundle sizes. When n does not divide m , we prove that SD-EFX allocations exist with probability at least $1 - O\left(\frac{n^2}{m} + n \log m / m^{\frac{n-1}{n}}\right)$.

Though the relation between m and n is asymptotically tight in Theorem 4, we show that it is possible to get better probability bounds for the important case when n is small, e.g., when n is a constant. For example, when $n = 2$ Round-Robin will not find an SD-EF allocation with probability $\Omega(1/\sqrt{m})$, while a careful analysis would say that such allocations do not exist with probability $\Omega(1/m)$. We propose the “Give-Away Round-Robin” algorithm, which first has every agent give one of their least desired remaining items to every other agent, and then proceeds with standard Round-Robin. In Theorem 5, we show that, when m is divisible by n , Give-Away Round-Robin outputs an SD-EF allocation with probability at least $1 - \tilde{O}_n\left(\frac{1}{m}\right)$.

1.2. Related Work

Dickerson et al. (2014) initiated the study of asymptotic fair division and showed that the welfare maximizing algorithm (allocate each good to the agent with the highest value for it) is envy-free with high probability for $m \in \Omega(n \log n)$ for the case of additive, non-negative valuations. This holds when items' values are drawn i.i.d. from a common prior and when agents' valuations are correlated, as long as, for every item j : (1) the probability that agent i has the highest value for j is exactly $1/n$, and (2) the expected value of $v_{i,j}$ conditioned on i having the highest value for j , is greater by constant than the expected value of $v_{i,j}$ conditioned on some other agent having the highest value for j . Manurangsi and Suksompong (2020, 2021) establish tight bounds for the

existence of envy-free allocations in the i.i.d. model: $m \in \Omega(n \log n / \log \log n)$ is a necessary and sufficient condition; similar to our Theorem 4, this bound is achieved by the classic Round-Robin algorithm. Bai and Gözl (2022) extend these results to the case of independent but non-identical additive agents. Kurokawa et al. (2016), Amanatidis et al. (2017), and Farhadi et al. (2019) show that weaker notions of fairness, namely maximin share fairness, also exist with high probability. Farhadi et al. (2019) also study a “stochastic items” model, where every item j has a probability distribution $\mathcal{D}_j$ (and $v_{i,j}$ is drawn from $\mathcal{D}_j$); interestingly, this model is not neutral. Suksompong (2016) shows that proportional allocations exist with high probability, while Manurangsi and Suksompong (2017) study envy-freeness when items are allocated to groups rather than to individuals. For non-additive stochastic valuations, Manurangsi and Suksompong (2021) and Gan et al. (2019) study the so-called *house allocation* problem, where there are m houses, n agents, with $v_{i,j} \sim D_i$ for each agent i and house j , and each agent must be allocated a single house; equivalently, one can think of unit-demand valuations and independently picking a uniformly random ranking of the items for each agent. The former paper shows that an envy-free assignment is likely to exist if $m/n > e + \epsilon$, for any constant $\epsilon > 0$, but unlikely to exist if $m/n < e - \epsilon$.

The existence of fair allocations, and their compatibility with efficiency, has also been studied in dynamic settings with additive valuations (Benadè et al. 2018). The algorithms of Dickerson et al. (2014) and Bai and Gözl (2022) can be readily applied even when items arrive in an online fashion, implying that Pareto efficiency ex-post and envy-freeness with high probability are compatible even in an online setting. The same result was recently shown to be (approximately) true even when the designer does not have access to the exact values of the agents (Benadè et al. 2022), or even when agents’ valuations are correlated (Zeng and Psomas 2020, Benadè et al. 2024).

Bouveret et al. (2010) study preferences represented by SCI-nets, which take the form of a partial order over bundles induced by a linear ordering over items. Order-consistent valuations are exactly those where pairwise dominance of bundles implies envy-freeness. As opposed to our interest here, Bouveret et al. (2010) are interested in computational aspects of the problem (e.g., the complexity of computing an envy-free allocation in the worst-case).

A related research direction is smoothed analysis. Traditionally, in smoothed analysis, an instance is generated by starting from a worst-case instance and adding a small amount of noise. This model has been widely successful in circumventing computational hardness (Spielman and Teng 2004, Blum and Spencer 1995, Feige and Kilian 1998, Kolla et al. 2011). More recently, this model has been used to bypass impossibility results in social choice (Xia 2020, 2021, Flanigan et al. 2023, Xia

2021), mechanism design (Psomas et al. 2019, Blum and Gözl 2021, Blum et al. 2017), and, closer to this paper, fair division (Bai et al. 2022). Specifically, Bai et al. (2022) start from a worst-case instance for additive agents and add an independent boost to the value of each item for each agent; the authors give sufficient conditions (on the boosting) for envy-free allocations to exist with high probability. One could interpret our model as a smoothed/semi-random model that starts from a worst-case instance and adds noise in the form of random renaming. In contrast with existing models, we permute the items and leave the valuation functions untouched.

2. Preliminaries

We consider the problem of dividing a set $\mathcal{M}$ of m indivisible items among a set $\mathcal{N}$ of $n \geq 2$ agents. Throughout, we assume that $\mathcal{N} = \{1, 2, \dots, n\}$ and the items are indexed by $1, 2, \dots, m$. An allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ is an n -partition of $\mathcal{M}$, where A_i denotes the set of items allocated to agent $i \in \mathcal{N}$. Formally, in any allocation $\mathcal{A}$, all items are allocated, $\bigcup_{i \in \mathcal{N}} A_i = \mathcal{M}$ (without this assumption, the allocating no items is trivially envy free) and each item is allocated to only one agent, $A_i \cap A_j = \emptyset$ for all distinct agents $i, j \in \mathcal{N}$. We sometimes refer to a subset of items as a *bundle*.

Each agent $i \in \mathcal{N}$ has a valuation function $v_i : 2^{\mathcal{M}} \mapsto \mathbb{R}$, which maps each subset of items $S \subseteq \mathcal{M}$ to $v_i(S)$, the agent’s value for S . We will assume that the valuations are normalized, i.e., $v_i(\emptyset) = 0$ for all $i \in \mathcal{N}$. We are interested in proving the existence of allocations that are fair with respect to the agent valuations. Our primary notion of fairness is envy-freeness.

DEFINITION 1 (ENVY-FREENESS). An allocation $\mathcal{A} = (A_1, \dots, A_n)$ is *envy-free (EF)* iff each agent prefers her own bundle over the bundle of any other agent, i.e., for all agents $i, j \in \mathcal{N}$, we have $v_i(A_i) \geq v_i(A_j)$.

For the case of indivisible items, envy-freeness can be too stringent of a requirement. A common relaxation is called envy-freeness up to any good, or EFX. EFX is considered the “best fairness analog of envy-freeness” in discrete fair division (Caragiannis et al. 2019). EFX allocations may not exist for monotone valuations, as recently shown by (Akrami et al. 2026), though their existence for additive valuations remains an elusive open problem.

DEFINITION 2 (ENVY-FREENESS UP TO ANY GOOD (EFX)). An allocation $\mathcal{A}$ is *envy-free up to any good (EFX)* iff for all agents $i, j \in \mathcal{N}$ where $A_j \neq \emptyset$, we have $v(A_i) \geq v(A_j \setminus \{g\})$ for all items $g \in A_j$.

When we claim that a property holds with high probability, it means that, for all n and $\delta > 0$, there exists a $m_0 = m_0(n, \delta)$ such that, for all $m > m_0$, the probability is at least $1 - \delta$.

In Section 4, we focus on a class of valuation functions that are consistent with an underlying preference order over the items. Given an ordering of the items $\pi = g_1 > g_2 > \dots > g_m$, we can define a partial order over bundles $\succeq_\pi^{\text{sd}}$ as follows. Let $A, \hat{A} \subseteq \mathcal{M}$ be any two bundles such that $A = \{g_{\sigma_1}, g_{\sigma_2}, \dots, g_{\sigma_{|A|}}\}$ and $\hat{A} = \{g_{\hat{\sigma}_1}, g_{\hat{\sigma}_2}, \dots, g_{\hat{\sigma}_{|\hat{A}|}}\}$ where sequences σ and $\hat{\sigma}$ are sorted in an increasing order. We say that $A \succeq_\pi^{\text{sd}} \hat{A}$ iff $|A| \geq |\hat{A}|$ and for all $k \in \{1, 2, \dots, |\hat{A}|\}$, we have $\sigma_k \leq \hat{\sigma}_k$. Additionally, $A \succ_\pi^{\text{sd}} \hat{A}$ iff $A \succeq_\pi^{\text{sd}} \hat{A}$ and $A \neq \hat{A}$.

DEFINITION 3 (ORDER-CONSISTENCY). A valuation function $v : 2^{\mathcal{M}} \mapsto \mathbb{R}$ is called *order-consistent with respect to an ordering π* iff for all bundles $A, B \subseteq \mathcal{M}$, $A \succeq_\pi^{\text{sd}} B$ implies that $v(A) \geq v(B)$. If v is order consistent with respect to some order π , we will simply call v order-consistent.

DEFINITION 4 (SD-EF). Suppose that each agent $i \in \mathcal{N}$ has a valuation function v_i which is order-consistent with respect to π_i . An allocation $\mathcal{A}$ is *sd-envy-free (SD-EF)* iff for all $i, j \in \mathcal{N}$ we have $A_i \succeq_{\pi_i}^{\text{sd}} A_j$.

SD-EF allocations are also EF. In fact, an allocation $\mathcal{A}$ which is SD-EF for agents with valuation functions $v_1, v_2, \dots, v_n$ that are order-consistent with respect to $\pi_1, \pi_2, \dots, \pi_n$, respectively, is EF for agents with any valuation functions $\hat{v}_1, \hat{v}_2, \dots, \hat{v}_n$ that are order-consistent with respect to $\pi_1, \pi_2, \dots, \pi_n$ respectively. When an allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ is SD-EF, it follows that $|A_i| = |A_j|$ for all $i, j \in \mathcal{N}$, i.e., the number of items is a multiple of the number of agents, $m = qn$ for some integer $q \in \mathbb{N}$. When $m \neq qn$, we consider the analogous strengthening of EFX.

DEFINITION 5 (SD-EFX). Suppose each agent $i \in \mathcal{N}$ has a valuation function v_i that is order-consistent with respect to π_i . An allocation $\mathcal{A}$ is *sd-envy-free up to any item (SD-EFX)* iff for all agents $i, j \in \mathcal{N}$, where $A_j \neq \emptyset$, we have $A_i \succeq_{\pi_i}^{\text{sd}} A_j \setminus \{g\}$ for all items $g \in A_j$.

As with SD-EF and EF, an allocation $\mathcal{A}$ which is SD-EFX is also EFX. While the existence of EFX allocations remains unknown, we give an instance with additive valuations for which SD-EFX allocations do not exist in Appendix C.

Finally, if v is a valuation function and π is a permutation over items, we use v^π to represent the permuted valuation function after renaming, where $v^\pi(S) := v(\pi^{-1}(S))$, and $\pi(S) := \{\pi(g) : g \in S\}$. Order-consistency is maintained under renaming.

3. Two Agents and Arbitrary Valuations

With two agents, an allocation is completely determined by selecting a single bundle, as the second agent receives the remainder. This structural simplicity fundamentally distinguishes the two-agent case from the many-agent setting. In this section, we isolate the specific conditions under which valuations are so rigid that a random renaming of items cannot yield an envy-free split.

To this end, we define a specific class of valuation functions. Intuitively, a valuation is *cardinality-consistent* if an agent's preference between a bundle and its complement depends entirely on the size of the bundle.

DEFINITION 6. A set function $f : 2^{\mathcal{M}} \rightarrow \mathbb{R}$ is *cardinality-consistent* if, for any bundle $S \subseteq \mathcal{M}$, the value $\text{sgn}(f(S) - f(\mathcal{M} \setminus S))$ is non-zero and depends only on $|S|$. Equivalently, for all $S, T \subseteq \mathcal{M}$ with $|S| = |T|$, it holds that:

$$\text{sgn}(f(S) - f(\mathcal{M} \setminus S)) = \text{sgn}(f(T) - f(\mathcal{M} \setminus T)) \in \{-1, 1\}.$$

Since these preferences are invariant for equal-sized bundles, we can succinctly represent a cardinality-consistent valuation using a one-dimensional signature.

DEFINITION 7. Let $f : 2^{\mathcal{M}} \rightarrow \mathbb{R}$ be a cardinality-consistent function. We define its *cardinality signature* as the function $p^f : [m] \rightarrow \{-1, 1\}$ given by $p^f(x) := \text{sgn}(f(S) - f(\mathcal{M} \setminus S))$, where $S \subseteq \mathcal{M}$ is any set with $|S| = x$.

The signature p^f is well-defined because f is cardinality-consistent; the output depends solely on the size of the set. Consequently, this structure inherently restricts the total number of items.

Observation 1 *If a function $f : 2^{\mathcal{M}} \mapsto \mathbb{R}$ is cardinality-consistent, then $|\mathcal{M}|$ must be odd.*

Proof Towards a contradiction, assume that $|\mathcal{M}|$ is even. Consider any set $S \subseteq \mathcal{M}$ such that $|S| = |\mathcal{M}|/2$ and let $T = \mathcal{M} \setminus S$. Since $|S| = |T|$, and f is cardinality-consistent, $\text{sgn}(f(S) - f(T)) = \text{sgn}(f(T) - f(S)) \in \{-1, 1\}$. However, this contradicts the fact that $\text{sgn}(f(S) - f(T)) = -\text{sgn}(f(T) - f(S))$. $\square$

Together, Definitions 6 and 7 generalize the simple counterexample of two agents with identical uniform valuations over an odd number of items — clearly random renaming is ineffective in such cases. Our main result in this section establishes envy-free allocations exist with high probability for arbitrary evaluations *unless* both agents are cardinality-consistent and share the exact same signature.

THEOREM 1. *For $n = 2$ agents with arbitrary valuation functions v_1 and v_2 , the following holds:*

- (a) *If both v_1 and v_2 are cardinality-consistent and have the same cardinality signatures, i.e., $p^{v_1} = p^{v_2}$, then the probability that an envy-free allocation exists is zero.*
- (b) *If both v_1 and v_2 are cardinality-consistent and have different cardinality signatures, i.e., $p^{v_1} \neq p^{v_2}$, then an envy-free allocation exists with probability one.*

(c) If at least one of v_1 and v_2 is not cardinality-consistent, an envy-free allocation exists with probability at least $1 - \frac{1}{m/2+1}$.

For even m , Theorem 1 together with Observation 1 implies that envy-free allocation exists with probability at least $1 - \frac{1}{m/2+1}$.

Theorem 1 highlights an interesting phase transition. If v_1 and v_2 are cardinality-consistent and $p^{v_1} = p^{v_2}$, then the probability that an envy-free allocation exists is zero; otherwise, the probability is at least $1 - \frac{1}{m/2+1}$. This bound is tight up to constants, as the following instance shows: consider m items and $n = 2$ agents where each agent positively values item 1 and has a zero value for all other items. In this instance, envy-free allocations exist with probability exactly $1 - 1/m$ (since the probability of the uniquely valued items coinciding after random renaming is $1/m$), matching the bound in Theorem 1 up to a factor of two.

To prove Theorem 1, we will use hypergraphs to represent the preferences of agents.

3.1. Representing preferences as hypergraphs

The proof of Theorem 1 is based on two key insights. First, the preferences of agents comparing the value of a bundle and its complement can be conveniently represented using hypergraphs. Second, given this hypergraph formulation, the probability of the existence of an envy-free allocation can be directly reduced to counting hypergraph automorphisms.

Before explaining the hypergraph representation, we make an observation.

Observation 2 *If any agent has a valuation function $v : 2^{\mathcal{M}} \rightarrow \mathbb{R}_{\geq 0}$ such that, for some $S \subseteq \mathcal{M}$, we have $v(S) = v(\mathcal{M} \setminus S)$, then an EF allocation exists with probability 1.*

Proof Without loss of generality, let v be the valuation function of agent 1. For any permutation π_1 , agent 1 is indifferent between the bundles $\pi_1(S)$ and $\pi_1(\mathcal{M} \setminus S)$, therefore, either the allocation $(\pi_1(S), \pi_1(\mathcal{M} \setminus S))$ or $(\pi_1(\mathcal{M} \setminus S), \pi_1(S))$ is envy-free. $\square$

Given Observation 2, we can focus on the case where the valuation function v_i of each agent i is such that, for every subset $S \subseteq \mathcal{M}$, we have $v_i(S) \neq v_i(\mathcal{M} \setminus S)$. We use the following hypergraph representation of arbitrary set function preferences. Given a valuation function v such that $v(S) \neq v(\mathcal{M} \setminus S)$ for all $S \subseteq \mathcal{M}$, let $H^v = (\mathcal{M}, E_v)$ be the hypergraph where $E_v := \{S \subseteq \mathcal{M} : v(S) > v(\mathcal{M} \setminus S)\}$. H^v is a hypergraph with m vertices, one for each item, and $2^m/2$ edges, one edge for every pair $\{S, \mathcal{M} \setminus S\}$.

Given a hypergraph $H = (\mathcal{M}, E)$ and a permutation $\pi : \mathcal{M} \rightarrow \mathcal{M}$, we will use $\pi(H) = (\mathcal{M}, E_\pi)$ to denote the hypergraph obtained by permuting the set of vertices by π , i.e., $E_\pi = \{\pi(S) : S \in E\}$.

As an example, consider the valuations and corresponding hypergraphs shown in Figure 1. Each preferred bundle of size 2 corresponds to an edge in the hypergraph, representing the corresponding agent's preferences. Since $H^1 \neq H^2$ there exists an envy-free allocation, in this case $(\{1, 3\}, \{2, 4\})$.




Figure 1 Example with $m = 4$ items and $n = 2$ agents where the hypergraphs corresponding to the valuations of agent 1, 2 are H^1, H^2 respectively; to simplify the figure, we do not list, show all the edges of the hypergraphs. The leftmost table shows all bundles of size $m/2 = 2$ and their complements; the preferred bundles of agent 1 (2) are boxed (underlined, respectively). The rightmost figure shows the hypergraph $\pi(H^1)$ obtained by renaming H^1 using permutation $\pi = (2, 1, 3, 4)$.

The next lemma formalizes the relationship between envy-free allocations and the hypergraph representation of valuations.

LEMMA 1. *Let $v_1, v_2 : 2^{\mathcal{M}} \rightarrow \mathbb{R}_{\geq 0}$ be two valuation functions with corresponding hypergraphs $H^{v_1} = (\mathcal{M}, E_{v_1})$ and $H^{v_2} = (\mathcal{M}, E_{v_2})$, respectively. If the hypergraphs are not identical, i.e. $H^{v_1} \neq H^{v_2}$, then an envy-free allocation exists.*

Proof. Since $H^{v_1} \neq H^{v_2}$, there must exist an edge $S \subset \mathcal{M}$, present in one of the hypergraphs that is not present in both. If $S \in E_{v_1}$ and $S \notin E_{v_2}$, then the allocation $\mathcal{A} = (S, \mathcal{M} \setminus S)$ is envy-free, since $v_1(S) > v_1(\mathcal{M} \setminus S)$ and $v_2(S) < v_2(\mathcal{M} \setminus S)$. Similarly, if $S \notin E_{v_1}$ and $S \in E_{v_2}$, then the allocation $\mathcal{B} = (\mathcal{M} \setminus S, S)$ is envy-free. Thus, envy-free allocations exist when $H^{v_1} \neq H^{v_2}$. $\square$

3.2. Relating envy-freeness to hypergraph isomorphism

When agent i with valuation function v_i and corresponding hypergraph H^{v_i} renames the items by applying permutation π , then the hypergraph $H^{v_i^\pi}$ corresponding to the new valuation function v_i^π is $H^{v_i^\pi} = \pi(H^{v_i})$. That is, permuting the names of the items amounts to permuting the vertices of the hypergraph. This follows directly from the following sequence of equivalences: subset S is present as an edge in $H^{v_i^\pi}$ iff $v_i^\pi(S) > v_i^\pi(\mathcal{M} \setminus S)$ iff $v_i(\pi^{-1}(S)) > v_i(\pi^{-1}(\mathcal{M} \setminus S))$ (from the definition of v_i^π) iff $\pi^{-1}(S)$ is present as an edge in H^{v_i} iff S is present as an edge in $\pi(H^{v_i})$.

In our setting, agents 1 and 2, with valuation functions v_1 and v_2 , apply permutations π_1 and π_2 to rename the items. As per Lemma 1, an envy-free allocation will exist whenever $\pi_1(H^{v_1}) \neq \pi_2(H^{v_2})$.

Since applying permutations π_1, π_2 does not change the structure of the hypergraph, $\pi_1(H^{v_1}) = \pi_2(H^{v_2})$ is only possible when H^{v_1} and H^{v_2} are isomorphic (and therefore, envy-free allocations trivially exist with probability 1 whenever H^{v_1} and H^{v_2} are non-isomorphic).

Accordingly, we assume that H^{v_1} and H^{v_2} are isomorphic. Without loss of generality, we also assume that π_1 is the identity permutation. To see why, observe that for uniformly random π_1, π_2 , the composition $\pi_1^{-1}\pi_2$ is also uniformly random. We view agent 2's renaming of the items as the permutation group (recall that the elements of the permutation group are all permutations $\pi : \mathcal{M} \rightarrow \mathcal{M}$ and the group operation which combines two group elements is the function composition) acting on the set of all hypergraphs on m vertices, denoted $\mathcal{H}_m = \{H = (V, E) : |V| = m\}$ and recall relevant group theoretic concepts. The *orbit* of a hypergraph H is an equivalence class defined as $O(H) = \{G \in \mathcal{H}_m : \exists \text{ permutation } \pi, G = \pi(H)\}$. The permutation group partitions $\mathcal{H}_m$ into ℓ orbits $O_1, O_2, \dots, O_\ell$, where any two hypergraphs are isomorphic if and only if they are in the same orbit. In Figure 1, H^1 and H^2 are not in the same orbit. Despite H^1 and $\pi(H^1)$ being isomorphic and in the same orbit, these valuations still permit an envy-free allocation.

Let O be the orbit containing both H^{v_1} and H^{v_2} . Starting from $H \in O$ and applying a random permutation π , the distribution of the permuted hypergraph $\pi(H)$ is also uniform over O . To upper bound the probability of the non-existence of EF allocations, it is sufficient to lower bound all orbit sizes. Specifically, $\Pr_{\pi_2}[H^{v_1} = \pi_2(H^{v_2})] = 1/|O|$. The orbit-stabilizer theorem (Dummit and Foote 2004) relates a graph's orbit size to the number of its stabilizers.

THEOREM 2 (Orbit-stabilizer theorem (Dummit and Foote 2004)). *If a finite group Π acts on a set X , then for every $x \in X$, $|\Pi| = |\text{Orb}(x)| \cdot |\text{Stab}(x)|$, where $\text{Stab}(x) = \{\pi \in \Pi : \pi \cdot x = x\}$.*

In our context Π is the permutation group and $X = \mathcal{H}_m$. A permutation π is an automorphism (and also a stabilizer) of hypergraph H iff $\pi(H) = H$. Let $\text{aut}(H) := \{\text{permutations } \pi : \pi(H) = H\}$ be the set of all automorphisms. Now $m! = |O(H)| \cdot |\text{aut}(H)|$ by Theorem 2. We now upper bound the number of automorphisms of a hypergraph H .

THEOREM 3. *For any non-empty and non-complete k -uniform hypergraph H on m vertices, the number of automorphisms of H satisfies $|\text{aut}(H)| \leq \frac{m!}{m-k+1}$.*

Proof. We prove this by induction on k . Our induction hypothesis is as follows: for a given k , the number of automorphisms of a non-empty and non-complete k -uniform hypergraph on m vertices is at most $\frac{m!}{m-k+1}$.

Base Case: $k = 1$. Fix an arbitrary m . A 1-uniform hypergraph $H = (\mathcal{M}, E)$ is simply a selection of some ℓ singleton sets of vertices: for all $S \in E$, we have $|S| = 1$. Since H is non-empty and non-complete, we have $0 < \ell < m$. Note that, any automorphism π of H must map any vertex that forms (resp. does not form) an edge to another vertex that forms (resp. does not form) an edge. Since there are ℓ vertices that form an edge and $m - \ell$ vertices that do not form an edge, we have

$$|\text{aut}(H)| \leq \ell!(m - \ell)! = \frac{m!}{\binom{m}{\ell}} \leq \frac{m!}{m} = \frac{m!}{m - k + 1}.$$

Induction Step: Now, assume that the induction hypothesis is true for k . Fix an arbitrary m and let $H = (\mathcal{M}, E)$ be a non-empty, non-complete $(k + 1)$ -uniform hypergraph on m vertices. For any vertex $i \in \mathcal{M}$, define H_i to be the k -uniform hypergraph induced by the vertex i . That is, $H_i = (V_i, E_i)$ where $V_i = \mathcal{M} \setminus \{i\}$ and $E_i = \{e \setminus \{i\} : e \in E, i \in e\}$. We consider two cases, based on whether the hypergraphs $\{H_i\}_{i \in \mathcal{M}}$ are all isomorphic to each other or not.

First, suppose the hypergraphs $\{H_i\}_{i \in \mathcal{M}}$ are not isomorphic to each other. Define $I_1 = \{j \in \mathcal{M} : \exists \text{ there exists a permutation } \pi \text{ satisfying } \pi(H_j) = H_1\}$ as the set of vertices whose corresponding hypergraph is isomorphic to H_1 and let $\ell = |I_1|$. $1 \in I_1$, therefore $\ell > 0$. Since not all H_i 's are isomorphic to each other we have $I_1 \subsetneq \mathcal{M}$, and thus $\ell < m$. In any automorphism π , every vertex in I_1 must be mapped to a vertex in I_1 . Hence, by the same argument as in the base case, the total number of automorphisms of H is at most

$$|\text{aut}(H)| \leq \ell!(m - \ell)! \leq \frac{m!}{m} \leq \frac{m!}{m - k + 1}.$$

Next, suppose the hypergraphs $\{H_i\}_{i \in \mathcal{M}}$ are isomorphic to each other. Because H is non-empty and non-complete and $k \geq 2$, each H_i cannot be empty or complete. That is, the hypergraphs $\{H_i\}_{i \in \mathcal{M}}$ are non-empty, non-complete, k -uniform, and have $m - 1$ vertices. Therefore, as per the induction hypothesis, we have $|\text{aut}(H_j)| \leq \frac{(m-1)!}{(m-1)-k+1}$ for all $j \in \mathcal{M}$. Finally, to count the total number of automorphisms of H , notice that there are exactly m ways to map vertex 1 to another vertex $j \in \mathcal{M}$. However, once this mapping is fixed, the hypergraph induced by the remaining vertices $\mathcal{M} \setminus \{j\}$, must get mapped to the hypergraph H_1 . Hence, the total number of automorphisms of H is at most

$$|\text{aut}(H)| \leq m \cdot \frac{(m-1)!}{(m-1)-k+1} = \frac{m!}{m - (k+1) + 1},$$

giving us the desired inequality. This completes the induction step. $\square$

Theorem 3 also generalizes to non-uniform graphs, given a certain condition holds. This is stated as the following corollary.

COROLLARY 1. *Let $H = (\mathcal{M}, E)$ be a hypergraph. For each $i \in [m - 1]$, define sub-hypergraphs $H_i = (V, E_i)$ such that $E_i = \{e \in E \mid |e| = i\}$. If any one amongst $H_1, H_2, \dots, H_{m-1}$ is both non-empty and non-complete, then $|\text{aut}(H)| \leq \frac{m!}{m/2+1}$.*

Proof. Note that for any $S \subseteq \mathcal{M}$ and permutation $\pi : \mathcal{M} \mapsto \mathcal{M}$, we have $|S| = |\pi(S)|$, i.e., π permutes sets of the same size among themselves. Hence, any automorphism σ of H must also be an automorphism of $H_1, H_2, \dots, H_{m-1}$; this leads to the inequality $|\text{aut}(H)| \leq |\text{aut}(H_i)|$ for all $i \in [m - 1]$. Let H_j be non-empty and non-complete. We have $|\text{aut}(H_j)| = |\text{aut}(\overline{H_j})|$ where $\overline{H_j} = (\mathcal{M}, \overline{E_j})$, $\overline{E_j} = \{\mathcal{M} \setminus e \mid e \in E_j\}$, and $\overline{H_j}$ is non-empty and non-complete since H_j is non-empty and non-complete. Theorem 3 implies that $|\text{aut}(H_j)| \leq \frac{m!}{m-j+1}$ and $|\text{aut}(\overline{H_j})| \leq \frac{m!}{j+1}$. This leads to the desired bound $|\text{aut}(H)| \leq |\text{aut}(H_j)| = |\text{aut}(\overline{H_j})| \leq \min\{\frac{m!}{m-j+1}, \frac{m!}{j+1}\} \leq \frac{m!}{m/2+1}$. $\square$

We combine the previous observations, with Theorem 3 (Corollary 1) to prove Theorem 1, the main result of this section.

3.3. Proof of Theorem 1

We begin by proving (a). Towards this, consider two cardinality-consistent valuations v_1 and v_2 such that $p^{v_1} = p^{v_2}$. By Observation 1, it suffices to consider the case where m is odd. Thus, for every allocation $(X, \mathcal{M} \setminus X)$, the two bundles have different cardinalities. Since v_1 and v_2 are cardinality-consistent, neither agent is indifferent between X and $\mathcal{M} \setminus X$. Therefore, $v_i(X) - v_i(\mathcal{M} \setminus X) \neq 0$, for $i \in \{1, 2\}$. Moreover, since $p^{v_1} = p^{v_2}$, we have, $\text{sgn}(v_1(X) - v_1(\mathcal{M} \setminus X)) = \text{sgn}(v_2(X) - v_2(\mathcal{M} \setminus X)) \in \{-1, 1\}$, or equivalently,

$$(v_1(X) - v_1(\mathcal{M} \setminus X))(v_2(X) - v_2(\mathcal{M} \setminus X)) > 0. \quad (1)$$

Since v_1 is cardinality-consistent and $|X| = |\pi_1^{-1}(X)|$, we get that $(v_1(X) - v_1(\mathcal{M} \setminus X))(v_1(\pi_1^{-1}(X)) - v_1(\pi_1^{-1}(\mathcal{M} \setminus X))) > 0$, or equivalently, $(v_1(X) - v_1(\mathcal{M} \setminus X))(v_1^{\pi_1}(X) - v_1^{\pi_1}(\mathcal{M} \setminus X)) > 0$. Similarly, since v_2 is also cardinality-consistent, we have $(v_2(X) - v_2(\mathcal{M} \setminus X))(v_2^{\pi_2}(X) - v_2^{\pi_2}(\mathcal{M} \setminus X)) > 0$. Combining these inequalities with Equation (1), we get $(v_1^{\pi_1}(X) - v_1^{\pi_1}(\mathcal{M} \setminus X))(v_2^{\pi_2}(X) - v_2^{\pi_2}(\mathcal{M} \setminus X)) > 0$, i.e., even after item renaming, both agents 1 and 2 strictly prefer the same bundle in the allocation $(X, \mathcal{M} \setminus X)$. That is, every allocation $(X, \mathcal{M} \setminus X)$ is not envy-free, or equivalently, the probability that an envy-free allocation exists is zero.

We proceed to prove (b). In particular, we will show that if both v_1 and v_2 are cardinality-consistent and $p^{v_1} \neq p^{v_2}$, then envy-free allocations exist with probability one. Suppose that v_1 and v_2 are cardinality-consistent and $p^{v_1} \neq p^{v_2}$. The latter inequality implies that there exists an $a \in [m]$

such that $v_1(S) > v_1(\mathcal{M} \setminus S)$ and $v_2(T) < v_2(\mathcal{M} \setminus T)$ for any $S, T \subseteq \mathcal{M}$ with $|S| = a$. Additionally, since v_1 and v_2 are cardinality-consistent, for every allocation $(X, \mathcal{M} \setminus X)$ with $|X| = a$ we have $v_1(X) > v_1(\mathcal{M} \setminus X)$ and $v_2(X) < v_2(\mathcal{M} \setminus X)$. That is, every allocation $(X, \mathcal{M} \setminus X)$ is envy-free, or envy-free allocations exist with probability one.

It remains to prove (c). Without loss of generality, assume that v_1 is not cardinality-consistent. This implies the existence of $S, T \subseteq \mathcal{M}$ with $|S| = |T|$ such that $(v_1(S) - v_1(\mathcal{M} \setminus S))(v_1(T) - v_1(\mathcal{M} \setminus T)) \leq 0$. If $(v_1(S) - v_1(\mathcal{M} \setminus S))(v_1(T) - v_1(\mathcal{M} \setminus T)) = 0$, then Observation 2 implies that an envy-free allocation exists with probability one. Henceforth, we consider the case of strict inequality $(v_1(S) - v_1(\mathcal{M} \setminus S))(v_1(T) - v_1(\mathcal{M} \setminus T)) < 0$. Consider hypergraphs $H^{v_1} = (\mathcal{M}, E^{v_1})$ and $H^{v_2} = (\mathcal{M}, E^{v_2})$ corresponding to v_1 and v_2 , respectively. If H^{v_1} and H^{v_2} are not isomorphic, then the inequality $\pi_1(H^{v_1}) \neq \pi_2(H^{v_2})$ always holds, which implies the existence of an envy-free allocation, since there exists a subset $U \subseteq \mathcal{M}$ with $U \in E^{v_1}$ and $\mathcal{M} \setminus U \in E^{v_2}$, making $(U, \mathcal{M} \setminus U)$ envy-free.

Suppose now H^{v_1} and H^{v_2} are isomorphic, and consider the sub-hypergraph $H_{|S|}^{v_1} = (\mathcal{M}, E_{|S|}^{v_1})$ where $E_{|S|}^{v_1} = \{X \subseteq \mathcal{M} \mid |X| = |S| \text{ and } v_1(X) > v_1(\mathcal{M} \setminus X)\}$. Note that $|E_{|S|}^{v_1} \cap \{S, T\}| = 1$ because $(v_1(S) - v_1(\mathcal{M} \setminus S))(v_1(T) - v_1(\mathcal{M} \setminus T)) < 0$ and $|S| = |T|$. It follows that $H_{|S|}^{v_1}$ is non-empty and non-complete, so Corollary 1 applies. By Lemma 1, envy-free allocations do not exist with probability at most $\Pr_\pi[H^{v_1} = \pi(H^{v_2})]$. Since both hypergraphs belong to some orbit $\mathcal{O}$, i.e., they are isomorphic, we have $\Pr_\pi[H^{v_1} = \pi(H^{v_2})] = \frac{1}{|\mathcal{O}|} = \frac{|\text{aut}(H^{v_1})|}{m!}$, where the last inequality follows from the orbit-stabilizer theorem (Dummit and Foote 2004). Combining with $|\text{aut}(H^{v_1})| \leq \frac{m!}{m/2+1}$ from Corollary 1 yields $\Pr_\pi[H^{v_1} = \pi(H^{v_2})] = \frac{|\text{aut}(H^{v_1})|}{m!} \leq \frac{1}{m/2+1}$. Therefore, envy-free allocations exist with probability at least $1 - \frac{1}{m/2+1}$. $\square$

4. Order-Consistent Valuations

In this section, we study order-consistent valuation functions for n agents and establish that SD-EF allocations exist with high probability. In fact, the well-known Round-Robin algorithm outputs such allocations with high probability in polynomial time.

Round-Robin takes as input an ordering of the items for each agent, and agents take turns taking the best available remaining item according to their personal order. For a valuation v_i order-consistent with respect to π_i this means picking the remaining item with lowest index. As long as Round-Robin has access to these orderings after renaming (that is, no other information about the valuation functions is needed), and as long as m is divisible by n , Round-Robin will result in an

SD-EF allocation with (asymptotically optimal) high probability. Otherwise, if m is not divisible by n and there is thus no hope for an EF allocation, Round-Robin will find an SD-EFX allocation with, again, asymptotically optimal probability.

THEOREM 4. *Let $\pi_1, \dots, \pi_n$ be sampled independently and uniformly at random. When $m = qn$, $q \in \mathbb{N}$, $\Pr[\text{Round-Robin is SD-EF}] \geq 1 - O\left(\frac{n^2}{m} + \frac{n \log m}{m^{\frac{n-1}{n}}}\right)$ and, if $m \geq 2n$ (where m may or may not be a multiple of n) $\Pr[\text{No EF allocation exists}] \in \Omega\left(\min\left(\frac{n^2}{m}, 1\right)\right)$. Further, for any m , $\Pr[\text{Round Robin is SD-EFX}] \geq 1 - O\left(\frac{n^2}{m}\right)$.*

Notice that if $m \in \omega(n^2)$, then, as n grows, SD-EF/SD-EFX allocations exist with high probability, while the lower bound implies that even EF allocations do not exist when $m \in O(n^2)$. In this sense, Round-Robin is asymptotically optimal.

For ease of exposition, we break the proof of Theorem 4 into separate lemmas. We first assume that $m = qn$ for some integer q , and prove the existence of SD-EF allocations.

LEMMA 2 (Round-Robin is SD-EF). *When $m = qn$, $\Pr[\text{Round Robin is SD-EF}] \geq 1 - O\left(\frac{n^2}{m} + \frac{n \log m}{m^{\frac{n-1}{n}}}\right)$.*

Second, we show how to extend our arguments to the case of arbitrary m and show the existence of SD-EFX allocations.

LEMMA 3 (Round-Robin is SD-EFX). *For any m , $\Pr[\text{Round Robin is SD-EFX}] \geq 1 - O\left(\frac{n^2}{m}\right)$.*

Finally, we show the lower bound on the probability that EF or SD-EFX allocations exist.

LEMMA 4. *For any m , if $m \geq 2n$, $\Pr[\text{No EF allocation exists}] \in \Omega\left(\min\left(\frac{n^2}{m}, 1\right)\right)$. The same bound holds for the probability that no SD-EFX exists.*

Combined, these three lemmas imply Theorem 4. We prove Lemma 2 in Section 4.1. Lemma 3 largely follows the same proof; we prove it in Section 4.2. Finally, Lemma 4 is proved in Section 4.3.

In Section 4.4, we observe that, even though Round-Robin is asymptotically optimal as n grows, for the important case of small (i.e., constant) n , it leaves a small gap in our understanding. We introduce a new algorithm, Give-Away Round Robin, which closes this gap.

4.1. Round-Robin is SD-EF: The proof of Lemma 2

In this section, we prove Lemma 2

LEMMA 2 (Round-Robin is SD-EF). *When $m = qn$, $\Pr[\text{Round Robin is SD-EF}] \geq 1 - O\left(\frac{n^2}{m} + \frac{n \log m}{m^{\frac{n-1}{n}}}\right)$.*

Our goal will be to show that the probability that Round-Robin does not output an SD-EF allocation is at most $\frac{33n^2}{m} + \frac{16n \log m}{m^{1-1/n}}$. We restrict our analysis to m and n such that $m \geq 2n$ and $\frac{n \log m}{m^{n-1/n}} \leq 1/16$, since otherwise, our target upper bound holds trivially. Fix m, n , and q such that $m = qn$. In each round of the Round-Robin algorithm, lower-indexed agents pick before higher-indexed ones, and therefore the former do not envy the latter. That is, agent i can only envy agent j 's bundle if $j < i$.

For agents $i, j \in \mathcal{N}$ with $j < i$, let $\mathcal{E}^{ij}$ be the event that $A_i \not\geq_{\pi_i}^{\text{sd}} A_j$, i.e., agent i does not SD-prefer their own bundle over that of agent j . $\mathcal{E}^{ij}$ only happens if there exists some $1 \leq k \leq m/n$ such that i prefers j 's k -th best item (according to i) over i 's k -th best item. At a high level, we'd like to compute this probability for each k , and then take a union bound over k to upper bound $\Pr[\mathcal{E}^{ij}]$. The probability that the final allocation is not SD-EF is the probability that one of these events occurs; formally, $\Pr[\bigcup_{i \in \mathcal{N}} \bigcup_{j: j < i} \mathcal{E}^{ij}]$. For $k = 1$ this probability is already $\Theta(n/m)$, so, in order to afford all the union bounds, we will require that the true probabilities of the bad events are much smaller than n/m , and that our analysis is relatively tight.

We use $g \succ_i g'$ to denote that item g is preferred to item g' under π_i . For a set of items $S \subseteq \mathcal{M}$, we write $g \succ_i S$ when g is preferred to all items in S , i.e., $g \succ_i g'$ for all $g' \in S$. Additionally, we use P_t to denote the pool of available items at “time” t , i.e., after t picks have been made (so, $P_0 = \mathcal{M}$ and $P_m = \emptyset$). Let g_t be the t 'th item picked during the execution of the Round Robin algorithm, that is, $P_{t+1} = P_t \setminus \{g_{t+1}\}$. Since agents pick items sequentially as per their index, g_t is picked by agent j exactly when $t \equiv j \pmod n$, $g_{(\ell-1)n+j}$ is the ℓ 'th item picked by agent j , and $A_j = \{g_{0 \cdot n+j}, \dots, g_{(q-1)n+j}\}$.

Taking the perspective of agent i , let b_{jk}^i denote the k 'th best item according to π_i in the final bundle A_j ; we will write b_{jk} if agent i is clear from the context. For agent i 's picks, $b_{ik} = g_{(k-1)n+i}$, i.e., their k 'th pick is also their k 'th favorite item in their bundle. However, this is not the case for items in others' bundles. Since the final bundles are all of size q , $A_i \geq_{\pi_i}^{\text{sd}} A_j$ exactly when $b_{ik} \succ_i b_{jk}$ for all $k \leq q$. Hence, we decompose the event $\mathcal{E}^{ij}$ into events $\mathcal{E}_k^{ij}$, where $\mathcal{E}_k^{ij}$ is the event that $b_{jk} \succ_i b_{ik}$. Now $\mathcal{E}^{ij} = \bigcup_{k=1}^q \mathcal{E}_k^{ij}$, and (by a union bound) our goal will be to upper bound

$$\sum_{i \in \mathcal{N}} \sum_{k=1}^q \Pr \left[\bigcup_{j: j < i} \mathcal{E}_k^{ij} \right]. \quad (2)$$

Fix an agent i and value k . To get a handle on $\Pr \left[\bigcup_{j \in \mathcal{N}: j < i} \mathcal{E}_k^{ij} \right]$, we condition on agent i 's ordering of the items, π_i . Whenever agent i picks g_t , g_t must be the best item remaining in P_{t-1}

according to π_i . When agent $j \neq i$ picks g_t (so $t \not\equiv i \pmod{n}$), the distribution of g_t over the random draws of $\pi_{-i} = \{\pi_1, \dots, \pi_n\} \setminus \{\pi_i\}$, even conditioned on all previous picks $g_1, \dots, g_{t-1}$, is uniformly random over P_{t-1} . Let R_ℓ be the set of items picked between agent i 's ℓ 'th and $(\ell + 1)$ 'th pick. That is, $R_0 = \{g_1, \dots, g_{i-1}\}$, for $1 \leq \ell \leq q - 1$, $R_\ell = \{g_{n(\ell-1)+i+1}, \dots, g_{\ell n+i-1}\}$, and $R_q = \{g_{n(q-1)+i+1}, \dots, g_m\}$. From the perspective of agent i , each item in R_ℓ is picked sequentially and uniformly at random from the remaining items. This allows us to analyze an equivalent two-step process wherein (1) R_ℓ is first selected uniformly at random from P_t , after which, (2) each item in R_ℓ is matched to a uniformly random agent (different than i) that picked in the corresponding timesteps.

We upper bound the expression in (2) by upper bounding each of the summands. Our analysis splits into cases based on the value of k , as care must be taken when comparing i 's worst item to that of other agents.

For all items $k < q = m/n$, i.e. all items except the last one picked by i , we can directly upper bound the probability that i prefers j 's k -th best item. The analysis leverages the insight that, from i 's perspective, and over the random draws of π_{-i} , items picked by other agents look like (uniformly) random selections from the pool of remaining items; therefore, the distribution of other agents' bundles is identical. For i to prefer j 's k -th best item over their own, j must have selected k items all better than i 's k -th pick, which can only occur if these items have all been picked in rounds 1 through k ; this event is unlikely. Lemma 5 shows the precise bound is $\Theta(1/\binom{m/n}{k})$, which requires carefully accounting for the items remaining at each step of Round-Robin and applying Gautschi's inequality (Gautschi 1959).

LEMMA 5. *For $k < q$, it holds that $\Pr \left[\bigcup_{j:j < i} \mathcal{E}_k^{ij} \right] \leq \frac{8}{\binom{q}{k}}$.*

When $k = q = m/n$, the above arguments fail to yield a sufficiently small probability. Instead, define L as the set of (roughly) $3n \log m$ worst items of agent i . We show neutrality implies that, with high probability, by the time i picks their m/n -th and worst item, all items in L have been picked by others. When this happens, i must like their worst item more than the worst (from i 's perspective) item of every agent who picked an item in L . Moreover, conditioned on all items in L being picked by the time of i 's last pick, it is very likely that every other agent received at least one of these items in L .

LEMMA 6. *For $k = q$, it holds that $\Pr \left[\bigcup_{j:j < i} \mathcal{E}_q^{ij} \right] \leq \frac{n}{m} + 2 \left(\frac{8n \log m}{m^{1-1/n}} \right)^{n+1-i}$.*

We show how to conclude the proof of Lemma 2 given these two lemmas, then proceed to prove them. Using Lemmas 5 and 6, we have

$$\begin{aligned}
& \Pr \left[\bigcup_{i,j} \mathcal{E}^{ij} \right] \\
& \leq \sum_{i \in \mathcal{N}} \left(\sum_{k=1}^{q-1} \Pr \left[\bigcup_j \mathcal{E}_k^{ij} \right] + \Pr \left[\bigcup_j \mathcal{E}_q^{ij} \right] \right) \\
& \leq \sum_{i \in \mathcal{N}} \left(\sum_{k=1}^{q-1} \frac{8}{\binom{q}{k}} + \frac{n}{m} + 2 \left(\frac{8n \log m}{m^{1-1/n}} \right)^{n+1-i} \right) \\
& = 8n \sum_{k=1}^{q-1} \frac{1}{\binom{q}{k}} + \frac{n^2}{m} + 2 \sum_{i=1}^n \left(\frac{8n \log m}{m^{1-1/n}} \right)^{n+1-i}.
\end{aligned}$$

We now bound each of these sums. For the first,

$$\begin{aligned}
\sum_{k=1}^{q-1} \frac{1}{\binom{q}{k}} & \leq \frac{1}{\binom{q}{1}} + \frac{1}{\binom{q}{q-1}} + \sum_{k=2}^{q-2} \frac{1}{\binom{q}{k}} \\
& \leq \frac{2}{q} + \sum_{k=2}^{q-2} \frac{1}{\binom{q}{2}} = \frac{2}{q} + \frac{2(q-3)}{q(q-1)} \leq \frac{4}{q} = \frac{4n}{m}.
\end{aligned} \tag{3}$$

For the second, since $\frac{8n \log m}{m^{1-1/n}} \leq 8/16 = 1/2$,

$$\begin{aligned}
\sum_{i=1}^n \left(\frac{8n \log m}{m^{1-1/n}} \right)^{n+1-i} & = \sum_{i=1}^n \left(\frac{8n \log m}{m^{1-1/n}} \right)^i \\
& \leq \sum_{i=1}^{\infty} \left(\frac{8n \log m}{m^{1-1/n}} \right)^i \\
& = \frac{\frac{8n \log m}{m^{1-1/n}}}{1 - \frac{8n \log m}{m^{1-1/n}}} \\
& \leq \frac{16n \log m}{m^{1-1/n}}.
\end{aligned}$$

We conclude that

$$\begin{aligned}
\Pr \left[\bigcup_{i,j} \mathcal{E}^{ij} \right] & \leq \frac{33n^2}{m} + \frac{16n \log m}{m^{1-1/n}} \\
& \in O \left(\frac{n^2}{m} + \frac{n \log m}{m^{\frac{n-1}{n}}} \right),
\end{aligned}$$

as needed. $\square$

Proof of Lemma 5 We directly upper bound each $\Pr \left[\mathcal{E}_k^{ij} \right]$ and union bound over the (at most) n possible choices of j . The equivalent random process of picking R^ℓ and assigning each item randomly discussed above implies $\Pr \left[\mathcal{E}_k^{ij} \right] = \Pr \left[\mathcal{E}_k^{ij'} \right]$, for each $j, j' < i$. Indeed, the items assigned to j and j' in this process appear in exactly the same R_ℓ sets. Now the distributions of A_j and $A_{j'}$ are identical even when conditioning on A_i . As a result, we only upper bound $\Pr \left[\mathcal{E}_k^{ij} \right]$ for $j = i - 1$; the same bound holds for all agents that i could possibly envy (i.e., for all agents j' with $j' < i$).

Consider the item $g_{(k-1)n+i}$, which is picked by i in the k 'th round. As noted above, $g_{(k-1)n+i} = b_{ik}$, so $\mathcal{E}_k^{ij}$ occurs exactly when $b_{jk} >_i b_{ik} = g_{(k-1)n+i}$. This can only occur if there are k timesteps, $t_1, \dots, t_k$, such that j picks an item g_{t_z} satisfying $g_{t_z} >_i g_{(k-1)n+i}$. For any $t > (k-1)n+i$, $g_{(k-1)n+i} >_i g_t$, since i always picks the best available remaining item. Hence, all k of these (unfortunate) picks must have occurred before i 's k 'th pick. Agent j made exactly k picks before i made their k 'th pick, so for $b_{jk} >_i g_{(k-1)n+i}$ to hold, a necessary and sufficient condition is that all these picks were preferred to $g_{(k-1)n+i}$ (according to π_i), i.e., $g_{(\ell-1)n+j} >_i g_{(k-1)n+i}$ for all $\ell = 1, \dots, k$.

Let $P_t^{(r)}$ be the r 'th best item remaining in P_t according to π_i . For all $t < t'$, $|P_t \setminus P_{t'}| = t' - t$, i.e., there are $t' - t$ fewer items available after $t' - t$ picks. This implies that for all $t < (k-1)n+i$, $g_{(k-1)n+i} \geq_i P_t^{((k-1)n+i-t)}$; this is tight only when all picks $g_t, \dots, g_{(k-1)n+i-1}$ are preferred to $g_{(k-1)n+i}$ in π_i . Furthermore, $g_{(k-1)n+i} = P_{(k-1)n+i-1}^{(1)}$ as i picks the best available item remaining at time $(k-1)n+i-1$. $|P_t \setminus P_{((k-1)n+i-1)}| = (k-1)n+i-t-1$, so at least one of $P_t^{(1)}, \dots, P_t^{((k-1)n+i-t)}$ must be available at time $(k-1)n+i-1$, and the best available item at time $(k-1)n+i-1$ must be at least as good as the worst of these. This implies that if $g_{(\ell-1)n+j} >_i g_{(k-1)n+i}$, then $g_{(\ell-1)n+j} >_i P_{(\ell-1)n+j-1}^{((k-\ell)n+(i-j)+1)}$. Using the fact that $j = i - 1$, this simplifies to $g_{(\ell-1)n+j} \geq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+1)}$ or, in words, agent j 's ℓ 'th pick must be one of the top $(k-\ell)n+1$ available items according to π_i . This necessary condition means that

$$\Pr \left[\mathcal{E}_k^{ij} \right] \leq \Pr \left[\bigcap_{\ell=1}^k \left\{ g_{(\ell-1)n+j} \geq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+1)} \right\} \right].$$

The key observation is that each of the events $g_{(\ell-1)n+j} \geq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+1)}$, over the samples of π_{-i} , are mutually independent. Hence, we get $\Pr \left[\bigcap_{\ell=1}^k \left\{ g_{(\ell-1)n+j} \geq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+1)} \right\} \right] = \prod_{\ell=1}^k \Pr \left[g_{(\ell-1)n+j} \geq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+1)} \right]$. Furthermore, $g_{(\ell-1)n+j}$ is a uniform random sample from $P_{(\ell-1)n+j-1}$, so

$$\Pr \left[g_{(\ell-1)n+j} \geq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+1)} \right] = \frac{(k-\ell)n+1}{|P_{(\ell-1)n+j-1}|}.$$

Noticing that $|P_t| = m - t$, we get that,

$$\begin{aligned}
 \Pr \left[\mathcal{E}_k^{ij} \right] &= \prod_{\ell=1}^k \frac{(k-\ell)n+1}{|P_{(\ell-1)n+j-1}|} \\
 &= \prod_{\ell=1}^k \frac{(k-\ell)n+1}{m - ((\ell-1)n+j-1)} \\
 &= \prod_{\ell=1}^k \frac{(k-\ell)n+1}{(q-\ell)n+n-j+1} \\
 &\leq \prod_{\ell=1}^k \frac{(k-\ell)n+1}{(q-\ell)n+1} \\
 &= \prod_{\ell=1}^k \frac{k-\ell+1/n}{q-\ell+1/n} \\
 &= \frac{\prod_{\ell=0}^{k-1} (1/n+\ell)}{\prod_{\ell=0}^{k-1} (q-k+1/n+\ell)} \\
 &= \frac{(1/n)^{\bar{k}}}{(q-k+1/n)^{\bar{k}}}.
 \end{aligned} \tag{4}$$

The inequality follows because $(n-j) \geq 0$, and the second to last equality follows by a change of variables. The notation $(x)^{\bar{r}}$ here represents the rising factorial $(x)^{\bar{r}} = x(x+1) \cdots (x+r-1)$. We make use of the known equality $(x)^{\bar{r}} = \frac{\Gamma(x+r)}{\Gamma(x)}$ to get that

$$\begin{aligned}
 \Pr \left[\mathcal{E}_k^{ij} \right] &= \frac{\Gamma(k+1/n)}{\Gamma(1/n)} \cdot \frac{\Gamma(q-k+1/n)}{\Gamma(q+1/n)} \\
 &\leq \frac{2}{n} \cdot \frac{\Gamma(k+1/n)\Gamma(q-k+1/n)}{\Gamma(q+1/n)},
 \end{aligned}$$

where we replaced $\Gamma(1/n)$ by using the fact that $\Gamma(1+x) = x\Gamma(x)$ for all x (which implies that $\Gamma(1/n) = n\Gamma(1+1/n)$), and that the Γ function has a global minimum on the positive reals above $1/2$ (Wrench 1968) (and thus, $n\Gamma(1+1/n) \geq \frac{n}{2}$).

Ideally, we could replace the $1/n$ terms inside the gamma function with 1s, which would allow us to convert the Γ 's to the corresponding factorial terms, and would simplify to $1/\binom{q}{k}$. We show that this is approximately correct by making use of Gautschi's inequality (Gautschi 1959), which states

$$x^{1-s} \leq \frac{\Gamma(x+1)}{\Gamma(x+s)} \leq (x+1)^{1-s},$$

for all $x > 0$ and $s \in (0, 1)$. We state our bound as the following claim, the proof of which is deferred to Appendix B.

$$\text{CLAIM 1. } \frac{2}{n} \cdot \frac{\Gamma(k+1/n)\Gamma(q-k+1/n)}{\Gamma(q+1/n)} \leq \frac{8}{n} \cdot \frac{1}{\binom{q}{k}}.$$

Therefore, for all $j < i$ and $k < q$, $\Pr \left[\mathcal{E}_k^{ij} \right] \leq \frac{8}{n} \cdot \frac{1}{\binom{q}{k}}$. Union bounding over the (at most) n possible values of j , we have that

$$\Pr \left[\bigcup_{j:j < i} \mathcal{E}_k^{ij} \right] \leq \frac{8}{\binom{q}{k}}. \quad \square$$

Proof of Lemma 6. Let $L \subseteq \mathcal{M}$ be the set of the bottom $\lceil 3n \log m \rceil + (n-1)$ items according to π_i . Notice that $\lceil 3n \log m \rceil + (n-1) \leq 3n \log m + n \leq 4n \log m$ since $\log(m) \geq \log(2n) \geq \log(4) \geq 1$ (we always interpret $\log$ as the natural $\log$). Recall that b_{iq} is the worst item in agent i 's bundle. We consider the event that $b_{iq} \notin L$. We show that the probability of $b_{iq} \in L$ is small and, conditioned on $b_{iq} \notin L$, it is unlikely that i envies anyone. Formally, notice that $\Pr \left[\bigcup_{j:j < i} \mathcal{E}_q^{ij} \right] = \Pr \left[\bigcup_{j:j < i} \mathcal{E}_q^{ij} \mid b_{iq} \notin L \right] \cdot \Pr [b_{iq} \notin L] + \Pr \left[\bigcup_{j:j < i} \mathcal{E}_q^{ij} \mid b_{iq} \in L \right] \cdot \Pr [b_{iq} \in L] \leq \Pr \left[\bigcup_{j:j < i} \mathcal{E}_q^{ij} \mid b_{iq} \notin L \right] + \Pr [b_{iq} \in L] \leq \sum_{j:j < i} \Pr \left[\mathcal{E}_q^{ij} \mid b_{iq} \notin L \right] + \Pr [b_{iq} \in L]$. We upper bound each of these terms individually.

$$\text{CLAIM 2. For all } j < i, \Pr \left[\mathcal{E}_q^{ij} \mid b_{iq} \notin L \right] \leq 1/m.$$

$$\text{CLAIM 3. } \Pr [b_{iq} \in L] \leq 2 \left(\frac{8n \log m}{m^{1-1/n}} \right)^{n+1-i}.$$

Given these two bounds, we can conclude the proof of the lemma as follows:

$$\begin{aligned} & \Pr \left[\bigcup_{j:j < i} \mathcal{E}_q^{ij} \right] \\ & \leq \sum_{j:j < i} \Pr \left[\mathcal{E}_q^{ij} \mid b_{iq} \notin L \right] + \Pr [b_{iq} \in L] \\ & \leq \frac{n}{m} + 2 \left(\frac{8n \log m}{m^{1-1/n}} \right)^{n+1-i}. \end{aligned} \quad (5)$$

It remains to prove the two claims.

Proof of Claim 2. Notice that $b_{iq} \notin L$ implies $A_i \cap L = \emptyset$, so all the items in L must have been chosen by other agents. If $A_j \cap L \neq \emptyset$ it follows that $b_{jq} \in L$ and, hence, $b_{iq} \succ_i b_{jq}$. Therefore, $\Pr \left[\mathcal{E}_q^{ij} \mid b_{iq} \notin L \right] \leq \Pr [A_j \cap L \neq \emptyset \mid b_{iq} \notin L]$ and it suffices to upper bound the latter.

To analyze the bundle of an agent $j \neq i$, we revert to the second view of the allocation process in which the bundles $R_0, \dots, R_q$ are first sampled and then matched to agents different from i . We condition on any valid $R_0, \dots, R_q$ such that $L \subseteq \bigcup_{\ell} R_{\ell}$, and now consider sampling the actual bundles $A_{j'}$ for $j' \neq i$ from these. Except for R_0 and R_q , each $|R_{\ell}| = n-1$ and A_j contains exactly one item from R_{ℓ} chosen uniformly at random. Therefore, with probability $\frac{|R_{\ell} \cap L|}{n-1}$, $A_j \cap R_{\ell} \cap L \neq \emptyset$.

In words, if there are r items from L in R_ℓ , then j receives one of them with probability $r/(n-1)$. Importantly, these are independent across rounds ℓ . It follows that $\Pr[A_j \cap L = \emptyset \mid b_{iq} \notin L] \leq \Pr\left[\sum_{\ell=1}^{q-1} \mathbb{I}[A_j \cap L \cap R_\ell = \emptyset] = 0\right]$. By linearity of expectation and the choice of L ,

$$\begin{aligned} \mathbb{E}\left[\sum_{\ell=1}^{q-1} \mathbb{I}[A_j \cap L \cap R_\ell]\right] &= \sum_{\ell=1}^{q-1} \frac{|R_\ell \cap L|}{n-1} \\ &\geq \frac{|L| - |R_0| - |R_q|}{n-1} \\ &\geq \frac{|L| - (n-1)}{n-1} \\ &\geq \frac{3n \log m}{n-1} \\ &\geq 3 \log m. \end{aligned} \tag{6}$$

The sum $\sum_{\ell=1}^{q-1} \mathbb{I}[A_j \cap L \cap R_\ell = \emptyset]$ is the sum of independent Bernoulli variables indicating whether j received an item from L in round ℓ . We have $\Pr\left[\sum_{\ell=1}^{q-1} \mathbb{I}[A_j \cap L \cap R_\ell = \emptyset] = 0\right] \leq \Pr\left[\sum_{\ell=1}^{q-1} \mathbb{I}[A_j \cap L \cap R_\ell = \emptyset] < \left(1 - \sqrt{\frac{2}{3}}\right) \mathbb{E}\left[\sum_{\ell=1}^{q-1} \mathbb{I}[A_j \cap L \cap R_\ell]\right]\right]$. Applying a Chernoff bound, we get an upper bound of $\exp\left(-\frac{1}{2} \left(\sqrt{\frac{2}{3}}\right)^2 \mathbb{E}\left[\sum_{\ell=1}^{q-1} \mathbb{I}[A_j \cap L \cap R_\ell]\right]\right) \leq \exp\left(\frac{-\frac{2}{3}(3 \log m)}{2}\right) = \exp(-\log m) = \frac{1}{m}$. Putting this all together, we have that

$$\begin{aligned} \Pr\left[\mathcal{E}_q^{ij} \mid b_{iq} \notin L\right] &\leq \Pr[A_j \cap L = \emptyset \mid b_{iq} \notin L] \\ &\leq 1/m, \end{aligned} \tag{7}$$

which concludes the proof of Claim 2. $\square$

The proof of Claim 3 is deferred to Appendix B. This concludes the proof of Lemma 6. $\square$

4.2. Round-Robin is SD-EFX: The proof of Lemma 3

We extend the analysis of Section 4.1 to get SD-EFX allocations for arbitrary m .

LEMMA 3 (Round-Robin is SD-EFX). *For any m , $\Pr[\text{Round Robin is SD-EFX}] \geq 1 - O\left(\frac{n^2}{m}\right)$.*

Fix an arbitrary m and let $m = qn + r$ for $1 \leq r \leq n$ and $r, q \in \mathbb{N}$. Let $\mathcal{A}^q$ be the allocation after qn steps, and $\mathcal{A}$ be the complete allocation of all $qn + r$ items. Let P_t denote the pool of available items at “time” t , i.e., after t picks have been made (so, $P_0 = \mathcal{M}$ and $P_m = \emptyset$). Notice that $\Pr[\mathcal{A} \text{ is SD-EFX}] \geq \Pr[\mathcal{A} \text{ is SD-EFX and } \mathcal{A}^q \text{ is SD-EF}] = \Pr[\mathcal{A} \text{ is SD-EFX} \mid \mathcal{A}^q \text{ is SD-EF}] \cdot \Pr[\mathcal{A}^q \text{ is SD-EF}]$. We show that $\Pr[\mathcal{A} \text{ is SD-EFX} \mid \mathcal{A}^q \text{ is SD-EF}] = 1$ and $\Pr[\mathcal{A}^q \text{ is SD-EF}] \geq 1 - O\left(\frac{n^2}{m}\right)$.

We begin with the former. Fix arbitrary samples $\pi_1, \dots, \pi_n$ such that $\mathcal{A}^q$ is SD-EF. We want to show that $\mathcal{A}$ is SD-EFX. Consider agents $i, j \in \mathcal{N}$. We show that either $A_i^q \succeq_{\pi_i}^{\text{sd}} A_j$ or, for all $g \in A_j$, $A_i^q \succeq_{\pi_i}^{\text{sd}} A_j \setminus \{g\}$. Since $A_i^q \subseteq A_i$, the same holds for A_i . It must be that either $A_j = A_j^q$ or $A_j = A_j^q \cup \{g\}$ for some $g \in P_{qn}$. If $A_j = A_j^q$ it follows that $A_i^q \succeq A_j^q = A_j$, since $\mathcal{A}^q$ is SD-EF by assumption. Suppose instead that $A_j = A_j^q \cup \{g^*\}$ for some $g^* \in P_{qn}$. The key observation is that $g^* \in P_{qn}$, implying it was available (and not selected) at every time step that i picked an item, so $g' \succ g^*$ for all $g' \in A_i$. Note that, this does not immediately imply the EFX condition, since g^* might not be the worst item (from i 's perspective) that j owns. Let b_{ik} and b_{jk} denote the k 'th best item in A_i^q and A_j^q according to π_i , respectively. By assumption $b_{ik} \succ_i b_{jk}$ for all k . Let b'_{jk} be the k 'th best item in $A_j \setminus \{g\}$. If $g^* = g$ then $A_i^q \succeq_{\pi_i}^{\text{sd}} A_j \setminus \{g^*\} = A_j^q$. Otherwise, if $g^* \neq g$, suppose $g^* = b'_{j\ell}$ for some ℓ . Notice that for all $k < \ell$, $b'_{jk} = b_{jk}$ so $b_{ik} \succ_i b'_{jk}$. For $k \geq \ell$, we have that $b_{ik} \succ_i g^* = b'_{j\ell} \geq_i b'_{jk}$. Hence, $A_i^q \succeq_{\pi_i}^{\text{sd}} A_j \setminus \{g^*\}$, as needed.

Next, we show $\Pr[\mathcal{A}^q \text{ is SD-EF}] \geq 1 - O\left(\frac{n^2}{m}\right)$. Specifically, we will show that the probability that $\mathcal{A}^q$ is not SD-EF is upper bounded by $\frac{2066n^2}{m}$. We assume that $m \geq 2000n^2$ (otherwise the bound is trivial), which implies that $qn \geq m/2$. The analysis is very similar to showing Round Robin is SD-EF when $m = qn$, so we simply describe which changes need to be made. Let b_{jk}^i denote the k 'th best item according to π_i in the bundle A_j^q ; for ease of notation we write b_{jk} when i is clear from the context. Let $\mathcal{E}_k^{ij}$ be the event that $b_{jk} \succ_i b_{ik}$, and let $\mathcal{E}^{ij} = \bigcup_{k=1}^q \mathcal{E}_k^{ij}$ be $A_i^q \not\succeq_{\pi_i}^{\text{sd}} A_j^q$. That is, $\mathcal{E}^{ij}$, b_{jk}^i , and $\mathcal{E}_k^{ij}$ are defined analogously to the proof of Lemma 2, but with respect to $\mathcal{A}^q$ instead of $\mathcal{A}$. Let R_ℓ be the set of items picked between agent i 's ℓ 'th and $(\ell + 1)$ 'th pick, and $R_q = P_{qn}$ be the set of items remaining items after the allocations in $\mathcal{A}^q$.

On a high level, our goal is again to upper bound by getting bounds on $\Pr[\mathcal{E}_k^{ij}]$ and $\sum_{i \in \mathcal{N}} \Pr\left[\bigcup_{j: j < i} \mathcal{E}_q^{ij}\right]$, similarly to Lemmas 5 and 6. Using very similar bounds we have

$$\begin{aligned}
\Pr\left[\bigcup_{i,j} \mathcal{E}^{ij}\right] &\leq \sum_{i \in \mathcal{N}} \sum_{k=1}^q \Pr\left[\bigcup_{j: j < i} \mathcal{E}_k^{ij}\right] \\
&\leq \sum_{i \in \mathcal{N}} \left(\sum_{k=1}^{q-1} \Pr\left[\bigcup_j \mathcal{E}_k^{ij}\right] + \Pr\left[\bigcup_j \mathcal{E}_q^{ij}\right] \right) \\
&\leq n \sum_{k=1}^{q-1} \Pr\left[\bigcup_j \mathcal{E}_k^{ij}\right] + \sum_{i \in \mathcal{N}} \Pr\left[\bigcup_j \mathcal{E}_q^{ij}\right] \\
&\leq 8n \sum_{k=1}^{q-1} \frac{1}{\binom{q}{k}} + \sum_{i \in \mathcal{N}} \Pr\left[\bigcup_j \mathcal{E}_q^{ij}\right].
\end{aligned}$$

Since $qn \geq m/2$ we can, similarly to Inequality (3), bound the first term by $\frac{64n^2}{m}$. For the second term, one complication is that our bound on $\Pr \left[\bigcup_j \mathcal{E}_q^{ij} \right]$ needs to take care of the $n = 2$ case, separately. This is caused because the analog of Lemma 6 is not a bound of $\frac{n}{m} + 2 \left(\frac{8n \log m}{m^{1-1/n}} \right)^{n+1-i}$, but a bound of $\frac{n}{m} + 2 \left(\frac{10n \log m}{m^{1-1/n}} \right)^{n+1-i+r}$, which is too weak for $n = 2$. We explain these differences, and how to address them, in Appendix B.1.

4.3. Round-Robin is asymptotically optimal: The proof of Lemma 4

We now prove Lemma 4.

LEMMA 4. *For any m , if $m \geq 2n$, $\Pr [\text{No EF allocation exists}] \in \Omega \left(\min \left(\frac{n^2}{m}, 1 \right) \right)$. The same bound holds for the probability that no SD-EFX exists.*

Fix n and m with $m \geq 2n$. We prove that EF and SD-EFX allocations do not exist with probability at least $\min \left(\frac{n^2}{8m}, \frac{1}{8} \right)$. For SD-EFX we first show that if two agents have the same favorite item, no SD-EFX allocation exists. For EF, if two agents have unit demand valuations, then clearly having the same favorite item also implies that no EF allocations exist. We then show using standard birthday paradox bounds that the probability two agents have the same favorite item is at least $\min \left(\frac{n^2}{8m}, \frac{1}{8} \right)$.

For the former, suppose two agents have the same favorite item g with $m \geq 2n$. We show that no SD-EFX allocation exists. Fix an allocation $\mathcal{A}$. First note that if $\mathcal{A}$ is not *balanced*, i.e., $||A_i| - |A_j|| \leq 1$, then it cannot be SD-EFX. Indeed, if $|A_i| - |A_j| > 1$, then for any $g \in A_i$, $|A_i \setminus \{g\}| > |A_j|$ so it cannot be the case that $|A_j| \geq_{\pi_j}^{\text{sd}} A_i \setminus \{g\}$. Next, suppose $\mathcal{A}$ is balanced. Since the two agents have the same favorite item, there is some agent i whose favorite good is g and who did not receive g . Further, there is an agent j such that $g \in A_j$, and since $\mathcal{A}$ is balanced, $|A_j| \geq 2$. Therefore, there is some $g' \neq g$ such that $g' \in A_j$. However, since $g \in A_j \setminus \{g'\}$, $A_i \not\geq_{\pi_i}^{\text{sd}} A_j \setminus \{g'\}$, so $\mathcal{A}$ is not SD-EFX.

Next, we show that the probability two agents have the same favorite item is at least $\min \left(\frac{n^2}{8m}, \frac{1}{8} \right)$. Notice that each agent's favorite item is a uniformly selected item, independent of other agents. The probability that two agents have the same first item is known from the collision analysis of hash functions. Indeed, it is known that, as long as $n \leq \sqrt{2m}$, the probability that two agents have the same first item is at least $\frac{n(n-1)}{4m}$ (Katz and Lindell 2020). In other words, as long as $m \geq n^2/2$, two agents have the same first item with probability at least $\frac{n(n-1)}{4m} \geq \frac{n^2}{8m}$. For $m < n^2/2$, notice that the probability of a collision is decreasing in m . Hence, the probability of a collision is at least the probability with n agents and $m' = \lceil n^2/2 \rceil \leq n^2$ items. The previous result says that this probability is then at least $\frac{n^2}{8m'} \geq \frac{1}{8}$.

4.4. Give-Away Round Robin

As we’ve shown, Round-Robin is asymptotically optimal as n grows. However, while it is often reasonable to assume that the number of items is large, it may be less reasonable to assume the same for the number of agents. When n is a constant, no SD-EF allocations exist with probability $\Omega_n(1/m)$. Although SD-EF allocations exist with probability $1 - O_n(1/m)$, when m is divisible by n , Theorem 4 only gives an upper bound of $1 - O_n(\frac{\log m}{m^{1-1/n}})$ on the probability of finding such an allocation. A careful analysis shows the analysis of Round Robin is relatively tight: with probability $\Omega_n(\frac{1}{m^{1-1/n}})$, the last agent will be left with their worst item, which implies the resulting allocation cannot be SD-EF. For example, with $n = 2$, Round-Robin will not find an SD-EF allocation with probability $\Omega(1/\sqrt{m})$, while the lower bound only implies such allocations do not exist with probability $\Omega(1/m)$. This motivates an alternative algorithm that closes this gap.

The probability of Round-Robin failing is greatest at the beginning and end of the process: at the beginning of the process, SD-EF fails if any other agent picks i ’s most preferred item; towards the end, there is a risk that all the remaining items are bad. Give-Away Round-Robin neutralizes the late-round risk: to avoid receiving a truly bad item, agents first *give away* one of their least desired items to each of the other agents, and then proceed with regular Round-Robin on the remaining items. We formalize this in Algorithm 1, and prove it achieves optimal (up to log factors) probability of finding an SD-EF allocation.

THEOREM 5. *Let $\pi_1, \dots, \pi_n$ be sampled independently and uniformly at random. When $m = qn$, $q \in \mathbb{N}$, $\Pr[\text{Give-Away Round Robin is SD-EF}] \geq 1 - \tilde{O}_n\left(\frac{1}{m}\right)$.*

For Round-Robin we could be certain that, other than the last item, items picked by agent i are “good” in the sense that they were the best in some pool of items. For Give-Away Round-Robin, each agent essentially starts with a small, random set of items. This complicates the analysis. For example, it is much harder to argue that SD-EF does not fail because of, say, the first ten items. In our analysis, we consider various cases (failure because of ‘high,’ ‘middle,’ ‘low,’ and last items), which need separate, delicate treatment.

Proof The proof follows a very similar structure to the proof of Theorem 4, though with more intricate analysis in various places. Fix $m = qn$ for an integer q . We will show the probability that Give-Away Round Robin is not SD-EF is upper bounded by $\frac{67584n^3 \log^2 m + 25216n^4 \log m + 2354n^5 + 128n^2 \log m + 24n^3 + 3n^2 + n}{m}$. Throughout, we will assume that m is lower bounded by the numerator, as otherwise, the bound is trivial.

Algorithm 1 Give-Away Round Robin

```

1: Input: Valuations  $v_1, \dots, v_n$  which are order consistent with respect to  $\pi_1, \dots, \pi_n$ 
2: Output: An allocation  $\mathcal{A} = (A_1, A_2, \dots, A_n)$ 
3: _____
4: Set  $A_i \leftarrow \emptyset$  for all agents  $i \in \{1, 2, \dots, n\}$ 
5: Set  $P \leftarrow \mathcal{M}$ 
6: Phase 1: Give-away ( $n(n-1)$  items)
7: for  $i \leftarrow 1$  to  $n$  do
8:   for  $j \leftarrow 1$  to  $n$ , where  $j \neq i$  do
9:     Let  $g \in P$  be the unallocated item with the highest index in  $\pi_i$ 
10:    Set  $A_j \leftarrow A_j \cup \{g\}$ 
11:    Set  $P \leftarrow P \setminus \{g\}$ 
12:   end for
13: end for
14: Phase 2: Round-Robin
15: while  $P \neq \emptyset$  do
16:   for  $i \leftarrow 1$  to  $n$  do
17:     if  $P \neq \emptyset$  then
18:       Let  $g \in P$  be the unallocated item with the lowest index in  $\pi_i$ 
19:       Set  $A_i \leftarrow A_i \cup \{g\}$ 
20:       Set  $P \leftarrow P \setminus \{g\}$ 
21:     end if
22:   end for
23: end while
24: return  $\mathcal{A} = (A_1, A_2, \dots, A_n)$ 

```

For $i \neq j$, we again let $\mathcal{E}^{ij}$ be the event that $A_i \not\prec_{\pi_i}^{\text{sd}} A_j$. Unlike in Round Robin, the give-away phase makes it possible that lower-indexed agents envy higher-indexed agents, so $\mathcal{E}^{ij}$ is defined for all agents $i \neq j$. We again wish to upper bound $\Pr \left[\bigcup_{i \in \mathcal{N}} \bigcup_{j: j \neq i} \mathcal{E}^{ij} \right]$. We reuse the notation $g \succ_i g'$ to denote preferences under π_i , and b_{jk}^i for i 's k 'th favorite item in A_j . We decompose $\mathcal{E}^{ij}$ into $\bigcup_{k=1}^q \mathcal{E}_k^{ij}$ as before, where $\mathcal{E}_k^{ij}$ is the event that $b_{jk} \succ_i b_{ik}$. Using a union bound, we wish to upper

bound

$$\sum_{i \in \mathcal{N}} \Pr \left[\bigcup_{k \leq q} \bigcup_{j: j \neq i} \mathcal{E}_k^{ij} \right].$$

We let γ_{ij} for $i \neq j$ be the item given from i to j in the give-away phase. We use $\tilde{P}_t$ for $0 \leq t \leq n(n-1)$ to denote the pool of available items after t giveaway steps. We then let P_t be the pool after t Round Robin steps (so $\tilde{P}_0 = \mathcal{M}$, $\tilde{P}_{n(n-1)} = P_0$, and $P_{m-n(n-1)} = \emptyset$). We let g_t be the item taken at the t 'th step of Round Robin. Therefore, $A_j = \{\gamma_{ij} \mid i \neq j\} \cup \{g_t \mid t \equiv j \pmod n\}$.

We fix an agent i and condition on π_i . The random process of selecting items is now similar to in Round Robin. For times t when γ_{ij} is given, γ_{ij} is the worst remaining item from $\tilde{P}_{t-1}$, and for $\gamma_{i'j}$ with $i' \neq i$ it is a uniformly selected item from $\tilde{P}_{t-1}$ (from agent i 's perspective). During the Round Robin phase, for g_t with $t \equiv i \pmod n$, g_t is the best available item from P_{t-1} , and when $t \not\equiv i \pmod n$, g_t is a uniformly selected item from P_{t-1} . For the Round Robin phase, we again consider the sets R_ℓ of items picked in the Round Robin phase between i 's ℓ 'th and $(\ell+1)$ 'th pick, so $R_0 = \{g_1, \dots, g_{i-1}\}$, for $1 \leq \ell \leq q-n$, $R_\ell = \{g_{n(\ell-1)+i+1}, \dots, g_{\ell n+i-1}\}$, and $R_q = \{g_{n(q-1)+i+1}, \dots, g_m\}$. As with Round Robin, each item in R_ℓ is picked sequentially and uniformly at random from the remaining items. This allows us to analyze an equivalent two-step process wherein R_ℓ is first selected uniformly at random from P_t , after which each item in R_ℓ is matched to a (non- i) agent that is picked in the corresponding timesteps uniformly at random.

We now break down $\bigcup_{k \leq q}$ via a union bound into different groupings of k .

$$\begin{aligned} & \sum_{i \in \mathcal{N}} \Pr \left[\bigcup_{k \leq q} \bigcup_{j: j \neq i} \mathcal{E}_k^{ij} \right] \\ &= \sum_{i \in \mathcal{N}} \left(\Pr \left[\bigcup_{k: k < 32 \log m + 4n} \bigcup_{j: j \neq i} \mathcal{E}_k^{ij} \right] \right. \\ & \quad + \Pr \left[\bigcup_{k: 32 \log m + 4n \leq k \leq q - 32 \log m - 5n} \bigcup_{j: j \neq i} \mathcal{E}_k^{ij} \right] \\ & \quad + \Pr \left[\bigcup_{k: q - 32 \log m - 5n < k \leq q-1} \bigcup_{j: j \neq i} \mathcal{E}_k^{ij} \right] \\ & \quad \left. + \Pr \left[\bigcup_{k: k = q} \bigcup_{j: j \neq i} \mathcal{E}_k^{ij} \right] \right). \end{aligned}$$

Fix an agent i . The following lemmas handle each of the summands individually.

$$\text{LEMMA 7. } \Pr \left[\bigcup_{k:k < 32 \log m + 4n} \bigcup_{j:j \neq i} \mathcal{E}_k^{ij} \right] \leq \frac{2048n^2 \log^2 m + 640n^3 \log m + 50n^4}{m}.$$

$$\text{LEMMA 8. } \Pr \left[\bigcup_{k:32 \log m + 4n \leq k \leq q - 32 \log m - 5n} \bigcup_{j:j \neq i} \mathcal{E}_k^{ij} \right] \leq \frac{1}{m}.$$

$$\text{LEMMA 9. } \Pr \left[\bigcup_{k:q - 32 \log m - 5n < k \leq q - 1} \bigcup_{j:j \neq i} \mathcal{E}_k^{ij} \right] \leq \frac{65536n^2 \log^2 m + 24576n^3 \log m + 2304n^4 + 128n \log m + 24n^2 + n}{m}.$$

$$\text{LEMMA 10. } \Pr \left[\bigcup_{k:k=q} \bigcup_{j:j \neq i} \mathcal{E}_k^{ij} \right] \leq \frac{2n}{m}.$$

Combining we have $\sum_{i \in \mathcal{N}} \Pr \left[\bigcup_{k \leq q} \bigcup_{j:j \neq i} \mathcal{E}_k^{ij} \right] \leq n \left(\frac{67584n^2 \log^2 m + 25216n^3 \log m + 2354n^4 + 128n \log m + 24n^2 + 3n + 1}{m} \right)$, which is at most $\frac{67584n^3 \log^2 m + 25216n^4 \log m + 2354n^5 + 128n^2 \log m + 24n^3 + 3n^2 + n}{m} \in \tilde{O}_n \left(\frac{1}{m} \right)$. $\square$

We prove Lemmas 7 to 10 in Appendix B.2.

5. Conclusion

In this paper, we study the existence of envy-free allocations beyond the case of additive valuations. We introduce a simple model to study this question: starting from a worst-case valuation function, randomly rename the items. We show that, in this model, if valuations are order-consistent (a valuation class general enough to include additive, unit-demand, budget additive, single-minded, etc.), even SD-envy-free allocations exist with high probability. In fact, a simple Round-Robin process will output such an allocation with high probability. Our bound on the probability is asymptotically tight, but we can improve upon it for the important case of a constant number of agents, using a variation of Round-Robin, that might be of independent interest. For arbitrary valuations, we characterize when envy-free allocations exist for the case of $n = 2$ agents. Our proof reduces the question about the existence of envy-free allocations to a question about the number of automorphisms in a uniform hypergraph.

An important problem we leave open is whether envy-free allocations exist for general valuation functions for $n > 2$ agents. Our approach of viewing valuations as hypergraphs immediately fails since now an allocation is not just a set and its complement. Interestingly, the existence of envy-free allocations in our model does not seem to become any easier, even if one is willing to make fairly common structural assumptions on the valuation functions. For example, one might attempt to prove a version of Theorem 1 for, e.g., monotone submodular functions (where strong concentration results are readily available, for example, in Vondrák (2010)), hoping to get a proof for $n = 2$ that can extend to more agents. However, we show in Appendix A that, at least when n divides m , any valuations $v_1, \dots, v_n$ can be transformed into monotone submodular valuations $v'_1, \dots, v'_n$ whose envy-free allocations are precisely the balanced envy-free allocations of $v_1, \dots, v_n$. Thus, proving

the existence of envy-free allocations for monotone submodular valuations is at least as hard as proving the existence for arbitrary valuations in this case.

Finally, it is worth noting that although neutral distributions generalize many existing models, it is highly unlikely that realistic human preferences satisfy this property. Thus, it would be quite interesting to explore alternate stochastic models for beyond additive valuations, which can capture correlation between the agents' valuations.

Appendix A: Reducing Arbitrary Valuations to Submodular Valuations

We show that when n divides m , the existence of balanced envy-free allocations for arbitrary valuation functions reduces to the existence of envy-free allocations for monotone, submodular valuation functions.

PROPOSITION 1. *Let $m = qn$ for some $q \in \mathbb{N}$. For any valuation functions $v_1, \dots, v_n : 2^M \rightarrow \mathbb{R}$, there exist monotone, submodular valuation functions $v'_1, \dots, v'_n$ such that, for any permutations $\pi_1, \dots, \pi_n$, an allocation $\mathcal{A}$ is envy-free under $((v'_1)^{\pi_1}, \dots, (v'_n)^{\pi_n})$ if and only if $\mathcal{A}$ is a balanced envy-free allocation under $(v_1^{\pi_1}, \dots, v_n^{\pi_n})$.*

Proof. For each $i \in \mathcal{N}$, let $\tilde{v}_i : \binom{M}{q} \rightarrow [0, 1/2]$ be an order-preserving rescaling of v_i on bundles of size q , and construct v'_i as follows:

$$v'_i(S) = \begin{cases} |S| & \text{if } |S| < q, \\ q - 1/2 + \tilde{v}_i(S) & \text{if } |S| = q, \\ q & \text{if } |S| > q. \end{cases}$$

It is straightforward to verify that v'_i is monotone and submodular: the marginal gain of adding an item to S is 1 for $|S| \leq q - 2$, in $[1/2, 1]$ for $|S| = q - 1$, in $[0, 1/2]$ for $|S| = q$, and 0 for $|S| \geq q + 1$. Under any permutations $\pi_1, \dots, \pi_n$, if an allocation $\mathcal{A} = (A_1, \dots, A_n)$ is unbalanced, there exists agents i, j so that $|A_i| < q$ and $|A_j| > q$, so $(v'_i)^{\pi_i}(A_i) \leq q - 1 < q = (v'_i)^{\pi_i}(A_j)$ and $\mathcal{A}$ is not envy-free. On balanced allocations ($|A_1| = \dots = |A_n| = q$), $(v'_i)^{\pi_i}$ and $v_i^{\pi_i}$ orders bundles identically, so $\mathcal{A}$ is envy-free under $((v'_1)^{\pi_1}, \dots, (v'_n)^{\pi_n})$ if and only if it is envy-free under $(v_1^{\pi_1}, \dots, v_n^{\pi_n})$. $\square$

Appendix B: Omitted proofs

Proof of Claim 1 We have $\frac{2}{n} \cdot \frac{\Gamma(k+1/n)\Gamma(q-k+1/n)}{\Gamma(q+1/n)} = \frac{2}{n} \cdot \frac{\Gamma(k+1)\Gamma(q-k+1)}{\Gamma(q+1)} \cdot \frac{\Gamma(k+1/n)}{\Gamma(k+1)} \cdot \frac{\Gamma(q-k+1/n)}{\Gamma(q-k+1)} \cdot \frac{\Gamma(q+1)}{\Gamma(q+1/n)}$. Then, using Gautschi's inequality we can get that

$$\begin{aligned} & \frac{2}{n} \cdot \frac{\Gamma(k+1/n)\Gamma(q-k+1/n)}{\Gamma(q+1/n)} \\ & \leq \frac{2}{n} \cdot \frac{k!(q-k)!}{q!} \frac{1}{k^{1-1/n}} \frac{1}{(q-k)^{1-1/n}} (q+1)^{1-1/n} \\ & = \frac{2}{n} \frac{1}{\binom{q}{k}} \left(\frac{q+1}{k(q-k)} \right)^{1-1/n} \\ & < \frac{2}{n} \frac{1}{\binom{q}{k}} \left(\frac{2q}{q/2} \right)^{1-1/n} \\ & \leq \frac{2}{n} \frac{1}{\binom{q}{k}} 4^{1-1/n} \\ & \leq \frac{8}{n} \cdot \frac{1}{\binom{q}{k}} \end{aligned}$$

where the third inequality holds because $q+1 \leq 2q$, both $k \geq 1$, $(q-k) \geq 1$, and at least one k and $(q-k)$ is $\geq q/2$.

□

Proof of Claim 3 Recall that $b_{iq} = g_{n(q-1)+i}$. Further, $g_{n(q-1)+i} \in L$ requires that $P_{n(q-1)+i-1} \subseteq L$, otherwise an item $g \notin L$, which i prefers to any item in L , is available (and i would have preferred to choose it). Since $|P_{n(q-1)+i-1}| = m - (n(q-1) + i - 1) = n - i + 1$, we have $\Pr[b_{iq} \in L] = \Pr[P_{n(q-1)+i-1} \subseteq L] = \sum_{S \in \binom{L}{n-i+1}} \Pr[P_{n(q-1)+i-1} = S]$. Fixing an arbitrary $S \in \binom{L}{n-i+1}$ we have $\Pr[P_{n(q-1)+i-1} = S] = \Pr\left[\bigcap_{t=1}^{n(q-1)+i-1} \{g_t \notin S\}\right] = \prod_{t=1}^{n(q-1)+i-1} \Pr[g_t \notin S \mid \bigcap_{t'=1}^{t-1} g_{t'} \notin S] = \prod_{t=1}^{n(q-1)+i-1} \Pr[g_t \notin S \mid S \subseteq P_{t-1}]$. For timesteps t where i made a pick, i.e., $t \equiv i \pmod{n}$, $\Pr[g_t \notin S \mid S \subseteq P_{t-1}]$ is hard to compute, as it depends on exactly which items are available. In these cases, we use the trivial upper bound of 1. For other timesteps t , g_t is a uniformly random item from P_{t-1} , so $\Pr[g_t \notin S \mid S \subseteq P_{t-1}] = 1 - \frac{|S|}{|P_{t-1}|} = 1 - \frac{n+1-i}{m+1-t} = \frac{m+1-t-(n+1-i)}{m+1-t}$. It follows that $\Pr[P_{n(q-1)+i-1} = S] \leq \prod_{t=1: t \not\equiv i \pmod{n}}^{n(q-1)+i-1} \frac{m+1-t-(n+1-i)}{m+1-t}$. To ease notation, we perform a change of variables to $t' = n(q-1) + i - t$, i.e., we reverse the product. The condition $t \not\equiv i \pmod{n}$ becomes $t' \not\equiv 0 \pmod{n}$, and we have $t' = n(q-1) + i - t = m+1-t-(n+1-i)$. Therefore, we can simplify to

$$\begin{aligned} & \prod_{t'=1: t' \not\equiv 0 \pmod{n}}^{n(q-1)+i-1} \frac{t'}{t' + (n+1-i)} \\ & = \frac{\prod_{t'=1}^{n(q-1)+i-1} \frac{t'}{t' + (n+1-i)}}{\prod_{t'=1: t' \equiv 0 \pmod{n}}^{n(q-1)+i-1} \frac{t'}{t' + (n+1-i)}} \\ & = \prod_{t'=1}^{n(q-1)+i-1} \frac{t'}{t' + (n+1-i)} \\ & \quad \cdot \prod_{t'=1: t' \equiv 0 \pmod{n}}^{n(q-1)+i-1} \frac{t' + (n+1-i)}{t'}. \end{aligned}$$

Observing that $n(q-1) + i - 1 + (n+1-i) = qn = m$, the first product is exactly equal to

$$\frac{(m - (n+1-i))!(n+1-i)!}{m!} = \frac{1}{\binom{m}{n+1-i}}$$

$$\begin{aligned}
&\leq \frac{1}{(m - (n + 1 - i))^{n+1-i}} \\
&\leq \frac{1}{(m - n)^{n+1-i}} \\
&\leq_{(m \geq 2n)} \left(\frac{2}{m}\right)^{n+1-i}.
\end{aligned}$$

The second product simplifies, as we only need to consider values $t' = \ell n$ for $1 \leq \ell \leq q - 1$. Hence,

$$\begin{aligned}
\prod_{t'=1:t'=0 \bmod n}^{n(q-1)+i-1} \frac{t' + (n + 1 - i)}{t'} &= \prod_{\ell=1}^{q-1} \frac{\ell n + (n + 1 - i)}{\ell n} \\
&= \prod_{\ell=1}^{q-1} \frac{\ell + \frac{n+1-i}{n}}{\ell} \\
&= \frac{(1 + \frac{n+1-i}{n})^{q-1}}{(q-1)!},
\end{aligned}$$

where the notation $(x)^{\bar{r}}$ represents the rising factorial $(x)^{\bar{r}} = x(x+1) \cdots (x+r-1)$. We use the fact that $(x)^{\bar{r}} = \frac{\Gamma(x+r)}{\Gamma(x)}$ to get that the above is equal to,

$$\frac{\Gamma(q + \frac{n+1-i}{n})}{\Gamma(1 + \frac{n+1-i}{n})\Gamma(q)}.$$

Similarly to Claim 1, using the fact that the Gamma function is lower bounded by $1/2$ and Gautschi's inequality Gautschi (1959), along with the assumption that $m \geq 2n$, the above is at most

$$\begin{aligned}
\frac{2\Gamma(q + \frac{n+1-i}{n})}{\Gamma(q)} &\leq 2 \left(q + \frac{n+1-i}{n}\right)^{\frac{n+1-i}{n}} \\
&\leq 2(q+1)^{\frac{n+1-i}{n}} \\
&\leq 2m^{\frac{n+1-i}{n}} \\
&= 2(m^{1/n})^{n+1-i}.
\end{aligned}$$

Combining with the previous bounds, we obtain for any $S \in \binom{L}{n+1-i}$ that, $\Pr[P_{n(q-1)-i+1} = S] \leq \left(\frac{2}{m}\right)^{n+1-i}$. $2(m^{1/n})^{n+1-i} = 2\left(\frac{2}{m^{1-1/n}}\right)^{n+1-i}$. The number of possible choices of S is

$$\binom{|L|}{n+1-i} \leq |L|^{n+1-i} \leq (4n \log m)^{n+1-i}. \quad (8)$$

Union bounding over all of these choices, we get that $\Pr[b_{iq} \in L] = \Pr[P_{n(q-1)-i+1} \subseteq L] \leq 2\left(\frac{8n \log m}{m^{1-1/n}}\right)^{n+1-i}$, which concludes the proof of Claim 3. $\square$

B.1. Missing details from the proof of Lemma 3

Here, we flesh out the missing details from the proof of Lemma 3.

For $k < q$, the proof of Lemma 5 goes through almost identically, to give us $\Pr[\mathcal{E}_k^{ij}] \leq \frac{8}{\binom{q}{k}}$. The only difference is that the equality in $m - ((\ell - 1)n + j - 1) = (q - \ell)n + n - j + 1$ of Equation (4) becomes a weak inequality; the inequality is in the right direction for the same conclusion to hold.

For $k = q$, the proof of Lemma 6, the analysis is very similar, even though more care is needed. To ensure that there are at least $3 \log m$ items remaining when lower bounding the expectation in Equation (6), since $|R_0| + |R_q|$ is only upper

bounded by $(n-1)+r \leq 2n-1$, we slightly change L , which is now defined as the bottom $\lceil 3n \log m \rceil + 2n-1$ items. It follows that $|L| \leq 5n \log m$. A second change is in the proof of Claim 3. $|P_{n(q-1)+i-1}|$, the number of items remaining when i makes their q 'th pick, is now $n-i+1+r$. The analysis continues to go through by replacing occurrences of $n-i+1$ with $n-i+1+r$. In the change of variables, we have $t' = n(q-1)+i-t = m+1-t-(n+1-i+r)$. This means that Inequality (8) becomes

$$\binom{|L|}{n+1-i+r} \leq |L|^{n+1-i+r} \leq (5n \log m)^{n+1-i+r},$$

and Inequality (5) becomes

$$\Pr \left[\bigcup_{j:j < i} \mathcal{E}_q^{ij} \right] \leq \frac{n}{m} + 2 \left(\frac{10n \log m}{m^{1-1/n}} \right)^{n+1-i+r}. \quad (9)$$

For $n=2$, we will need a bound that is even stronger than (9). When there are two agents the only envy event we need to consider is $\mathcal{E}_q^{21}$. We claim that we need only include i 's worst $1+r \leq 3$ items in L to get $\Pr [\mathcal{E}_q^{21} \mid b_{2q} \notin L] = 0$. Indeed, suppose $b_{2q} \notin L$, so b_{2q} is not one of the 3 worst items according to π_2 . Since $|R_q| \leq 2$ and $L \cap A_2^q = \emptyset$, it must be that $A_1^q \cap L \neq \emptyset$. We conclude $b_{1q} \in L$, so $b_{2q} \succ_2 b_{1q}$ and event $\mathcal{E}_q^{21}$ does not occur. Therefore, when $n=2$,

$$\sum_{i \in N} \Pr \left[\bigcup_{j:j < i} \mathcal{E}_q^{ij} \right] = \Pr [\mathcal{E}_q^{21}] \leq 2 \left(\frac{2n}{m^{1/2}} \right)^{1+r} \quad (10)$$

We are now ready to upper bound the probability that $\mathcal{A}^q$ is not SD-EF.

$$\begin{aligned} \Pr \left[\bigcup_{i,j} \mathcal{E}^{ij} \right] &\leq \sum_{i \in N} \sum_{k=1}^q \Pr \left[\bigcup_{j:j < i} \mathcal{E}_k^{ij} \right] \\ &\leq \sum_{i \in N} \left(\sum_{k=1}^{q-1} \Pr \left[\bigcup_j \mathcal{E}_k^{ij} \right] + \Pr \left[\bigcup_j \mathcal{E}_q^{ij} \right] \right) \\ &\leq n \sum_{k=1}^{q-1} \Pr \left[\bigcup_j \mathcal{E}_k^{ij} \right] + \sum_{i \in N} \Pr \left[\bigcup_j \mathcal{E}_q^{ij} \right] \\ &\leq 8n \sum_{k=1}^{q-1} \frac{1}{\binom{q}{k}} + \sum_{i \in N} \Pr \left[\bigcup_j \mathcal{E}_q^{ij} \right]. \end{aligned}$$

Since $qn \geq m/2$ we can refine Inequality (3) to get $\sum_{k=1}^{q-1} \frac{1}{\binom{q}{k}} \leq \frac{4}{q} \leq \frac{8n}{m}$, so the first term is at most $\frac{64n^2}{m}$.

When $n=2$, the second term is bounded by $2 \left(\frac{2n}{m^{1/2}} \right)^{1+r}$ by (10). By our assumption $m \geq 2000n^2$ and $r \geq 1$, so $\frac{2n}{m^{1/2}} \leq 1$ and

$$2 \left(\frac{2n}{m^{1/2}} \right)^{1+r} \leq 2 \left(\frac{2n}{m^{1/2}} \right)^{1+1} = \frac{8n^2}{m},$$

from which it follows that $\Pr [\bigcup_{i,j} \mathcal{E}^{ij}] \leq \frac{64n^2}{m} + \frac{8n^2}{m} = \frac{72n^2}{m} \leq \frac{2066n^2}{m}$.

Next, we show the same bound for $n \geq 3$. Here

$$\begin{aligned} \Pr \left[\bigcup_j \mathcal{E}_q^{ij} \right] &\leq \frac{n^2}{m} + 2 \sum_{i \in N} \left(\frac{10n \log m}{m^{1-1/n}} \right)^{n+1-i+r} \\ &\leq \frac{n^2}{m} + 2 \sum_{i=1}^{\infty} \left(\frac{10n \log m}{m^{1-1/n}} \right)^{i+r} \\ &\leq \frac{n^2}{m} + 2 \sum_{i=1}^{\infty} \left(\frac{10n \log m}{m^{2/3}} \right)^{i+r}. \end{aligned}$$

We claim that $\frac{\log m}{m^{1/6}} \leq \frac{6}{e}$. To see why, observe that the derivative of $\frac{\log m}{m^{1/6}}$ is $\frac{6-\log(m)}{6m^{7/6}}$ the function is maximized over positive reals exactly at $m = e^6$, where it takes value $\frac{6}{e}$. Substituting into the above,

$$\Pr \left[\bigcup_j \mathcal{E}_q^{ij} \right] \leq \frac{n^2}{m} + 2 \sum_{i=1}^{\infty} \left(\frac{60n}{em^{1/2}} \right)^{i+r}.$$

Since $(60/e)^2 \leq 500$ and we are restricted to $m \geq 2000n^2$, $\frac{60n}{em^{1/2}} \leq \sqrt{\frac{500}{2000}} = 1/2$. Using the fact that $r \geq 1$, this infinite series is at most

$$\frac{n^2}{m} + 4 \left(\frac{60n}{em^{1/2}} \right)^2 \leq \frac{2001n^2}{m}.$$

Hence,

$$\Pr \left[\bigcup_j \mathcal{E}_q^{ij} \right] \leq \frac{2066n^2}{m},$$

as needed.

B.2. Missing lemmas from Section 4.4

Proof of Lemma 7 Let T be the top $\lfloor 32 \log m + 4n \rfloor$ items according to π_i . Notice that if $T \subseteq A_i$, $b_{ik} > b_{jk}$ for all $k \leq 32 \log m + 4n$. More formally, this means that,

$$\Pr \left[\bigcup_{k: k < 32 \log m + 4n} \bigcup_{j: j \neq i} \mathcal{E}_k^{ij} \right] \leq \Pr [T \not\subseteq A_i],$$

so we directly upper bound the latter. For $T \not\subseteq A_i$ to occur, it must be the case that some other agent j received an item in T before i made their $(\lfloor 32 \log m + 4n \rfloor)$ 'th pick. Since we are assuming $m \geq 32 \log m + 4n + n^2 \geq |T| + n^2$, i would never give away an item in T during the giveaway phase as a worse item must be available. Hence, the only way that another agent receives an item in T is if they are given it by a non- i agent in the giveaway phase or they pick it in the Round Robin phase before time $n(|T| - 1) + i \leq n|T|$. From i 's perspective all of these items are chosen uniformly at random from the pool. There are at most $(n-1)(|T| + n) \leq n|T| + n^2 \leq 32n \log m + 5n^2$ such items. When they are taken, the pool is always of size at least $m - (32n \log m + 5n^2)$. Hence, each is selected with probability at most $\frac{32n \log m + 5n^2}{m - 32n \log m - 5n^2}$. Since we are assuming $m \geq 64n \log m + 10n^2$, this probability is at most $\frac{64n \log m + 10n^2}{m}$. Union bounding over the at most $32n \log m + 5n^2$ choices of k , this implies that the total probability is at most

$$\begin{aligned} & \Pr \left[\bigcup_{k: k < 32 \log m + 4n} \bigcup_{j: j \neq i} \mathcal{E}_k^{ij} \right] \\ & \leq \frac{(64n \log m + 10n^2)(32n \log m + 5n^2)}{m} \\ & = \frac{2048n^2 \log^2 m + 640n^3 \log m + 50n^4}{m}. \quad \square \end{aligned}$$

Proof of Lemma 8 Here, we upper bound each $\Pr [\mathcal{E}_k^{ij}]$ individually and union bound over the at most q choices of k and n choices of j . Fix agent $j \neq i$ and k such that $32 \log m + 4n \leq k \leq q - 32 \log m - 5n$. We first consider $j < i$ and later show how to extend it to $j > i$. Notice that $b_{ik} \geq_i g_{(k-1)n+i}$, i.e., i 's k 'th favorite item is at least as good as their k 'th pick in the Round Robin phase (though it may be better if i was lucky during the giveaway phase). Additionally, if $b_{jk} >_i b_{ik} \geq_i g_{(k-1)n+i}$, this implies that A_j contains at least k items strictly preferred to $g_{(k-1)n+i}$. None of these

could have been chosen after time $(k-1)n+i$ in the Round Robin phase as $g_{(k-1)n+i}$ was the best available item from $P_{(k-1)n+i-1}$. Additionally, $g_{(k-1)n+i} \succ \gamma_{ij}$ (the item i gave to j) as γ_{ij} was the worst available item at a time when $g_{(k-1)n+i}$ was available. Hence, for $\mathcal{E}_k^{ij}$ to hold, at least k of $\{\gamma_{i'j} \mid i' \neq i, j\} \cup \{g_j, \dots, g_{(k-1)n+j}\}$ are preferred to $g_{(k-1)n+i}$. Since $|\{\gamma_{i'j} \mid i' \neq i, j\}| = n-2$, this implies that $|\{\ell \mid g_{(\ell-1)n+j} \succ_i g_{(k-1)n+i}\}| \geq k - (n-2)$. More formally, we have that

$$\Pr \left[\mathcal{E}_k^{ij} \right] \leq \Pr \left[\sum_{\ell=1}^k \mathbb{I}[g_{(\ell-1)n+j} \succ_i g_{(k-1)n+i}] \geq k - (n-2) \right].$$

As before, a necessary condition for $g_{(\ell-1)n+j} \succ_i g_{(k-1)n+i}$ is $g_{(\ell-1)n+j} \succeq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+i-j)}$, hence,

$$\Pr \left[\sum_{\ell=1}^k \mathbb{I}[g_{(\ell-1)n+j} \succ_i g_{(k-1)n+i}] \geq k - (n-2) \right] \leq \Pr \left[\sum_{\ell=1}^k \mathbb{I}[g_{(\ell-1)n+j} \succeq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+i-j)}] \geq k - (n-2) \right],$$

which is equal to $\Pr \left[\sum_{\ell=1}^k \mathbb{I}[g_{(\ell-1)n+j} \not\succeq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+i-j)}] \leq n-2 \right]$. Now each of these indicators is an independent Bernoulli random variable. This means we can use a Chernoff bound on the sum. To do so, we first bound the expectation of the sum. We have that

$$\begin{aligned} & \mathbb{E} \left[\sum_{\ell=1}^k \mathbb{I}[g_{(\ell-1)n+j} \not\succeq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+i-j)}] \right] \\ &= \sum_{\ell=1}^k \Pr \left[g_{(\ell-1)n+j} \not\succeq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+i-j)} \right] \\ &= \sum_{\ell=1}^k 1 - \frac{(k-\ell)n+i-j}{|P_{(\ell-1)n+j-1}|} \\ &= \sum_{\ell=1}^k 1 - \frac{(k-\ell)n+i-j}{n(q-(n-1)) - ((\ell-1)n+j-1)} \\ &= \sum_{\ell=1}^k 1 - \frac{(k-\ell)n+i-j}{n(q-n-\ell)+1-j} \\ &= \sum_{\ell=1}^k \frac{n(q-n-k)+1-i}{n(q-n-\ell)+1-j} \\ &\geq \sum_{\ell=1}^k \frac{n(q-n-k)+1-i+(i-1)}{n(q-n-\ell)+1-j+(i-1)} \\ &= \sum_{\ell=1}^k \frac{n(q-n-k)}{n(q-n-\ell)+i-j} \\ &\geq \sum_{\ell=1}^k \frac{n(q-n-k)}{n(q-n-\ell)+n} \\ &\geq \sum_{\ell=1}^k \frac{q-n-k}{q-n-\ell+1} \\ &\geq k \cdot \frac{q-n-k}{q-n+1} \\ &\geq \frac{(q-n-k)k}{q-n+1} + \frac{(q-n-k)}{q-n+1} - 1 \\ &= \frac{(q-n-k)(k+1)}{q-n+1} - 1. \end{aligned}$$

The last expression is at least $\frac{\max(q-n-k, k+1) \cdot \min(q-n-k, k+1)}{q-n+1} - 1 \geq \frac{1}{2} \cdot \min(q-n-k, k+1) - 1 \geq 16 \log m + 2n - 1$.

Let $\mu = \mathbb{E} \left[\sum_{\ell=1}^k \mathbb{I}[g_{(\ell-1)n+j} \not\geq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+i-j)}] \right]$. The above implies that $(n-2) \leq \frac{1}{2}\mu$. Hence, a Chernoff bound implies that

$$\begin{aligned} & \Pr \left[\sum_{\ell=1}^k \mathbb{I}[g_{(\ell-1)n+j} \not\geq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+i-j)}] \leq n-2 \right] \\ & \leq \Pr \left[\sum_{\ell=1}^k \mathbb{I}[g_{(\ell-1)n+j} \not\geq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+i-j)}] \leq \frac{1}{2}\mu \right] \\ & \leq \exp \left(-\frac{\left(\frac{1}{2}\right)^2 \mu}{2} \right) \\ & = \exp \left(-\frac{\mu}{8} \right) \\ & \leq \exp(-2 \log m) = \frac{1}{m^2}. \end{aligned}$$

All together, this implies that for $j < i$, $\Pr \left[\mathcal{E}_k^{ij} \right] \leq 1/m^2$.

For $j > i$, the argument can be modified as follows. Notice that j 's k 'th pick will be after i 's k 'th pick, hence, we need only sum over ℓ from 1 to $k-1$. This gives $\Pr \left[\sum_{\ell=1}^{k-1} \mathbb{I}[g_{(\ell-1)n+j} >_i g_{(k-1)n+i}] \geq k - (n-2) \right] \leq \Pr \left[\sum_{\ell=1}^{k-1} \mathbb{I}[g_{(\ell-1)n+j} \geq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+i-j)}] \geq k - (n-2) \right] = \Pr \left[\sum_{\ell=1}^{k-1} \mathbb{I}[g_{(\ell-1)n+j} \not\geq_i P_{(\ell-1)n+j-1}^{((k-\ell)n+i-j)}] \leq n-3 \right]$.

The analysis of the expectation goes through (replacing k with $k-1$) since we never use that $j < i$, at least until getting the expectation is lower bounded by

$$\sum_{\ell=1}^{k-1} \frac{q-n-k}{q-n-\ell+1}.$$

Now, we can continue with

$$\sum_{\ell=1}^{k-1} \frac{q-n-k}{q-n-\ell+1} \geq \sum_{\ell=1}^k \frac{q-n-k}{q-n-\ell+1} - 1,$$

at which point the remainder of the analysis continues to hold (off by one) until $\mu \geq 16 \log m + 2n - 2$. This is still sufficient for $(n-3) \leq \frac{1}{2}\mu$ and $\mu \geq 16 \log m$, which is all that was needed for the Chernoff bound to hold. This means that $\Pr \left[\mathcal{E}_k^{ij} \right] \leq 1/m^2$ for $j > i$ as well.

Union bounding over the q possible choices of k and n possible choices of j , we get that

$$\Pr \left[\bigcup_{k: 32 \log m + 4n \leq k \leq q - 32 \log m - 5n} \bigcup_{j: j \neq i} \mathcal{E}_k^{ij} \right] \leq \frac{qn}{m^2} = \frac{1}{m}. \quad \square$$

Proof of Lemma 9 Let L be the last $\lfloor 64n \log m + 12n^2 \rfloor$ items according to π_i . Just as in Lemma 6, we will decompose $\Pr \left[\bigcup_{k: q-32 \log m-5n < k \leq q-1} \bigcup_{j: j \neq i} \mathcal{E}_k^{ij} \right] \leq \Pr \left[\bigcup_{k: q-32 \log m-5n < k \leq q-1} \bigcup_{j: j \neq i} \mathcal{E}_k^{ij} \mid b_{(q-1)i} \notin L \right] + \Pr \left[b_{(q-1)i} \in L \right] \leq \sum_{j: j \neq i} \Pr \left[\bigcup_{k: q-32 \log m-5n < k \leq q-1} \mathcal{E}_k^{ij} \mid b_{(q-1)i} \notin L \right] + \Pr \left[b_{(q-1)i} \in L \right]$.

We upper-bound each of these summands. We begin with $\Pr \left[b_{(q-1)i} \in L \right]$. We again split decompose this depending on if i received any goods in L during the giveaway phase.

$$\Pr \left[b_{(q-1)i} \in L \right] \leq \Pr \left[b_{(q-1)i} \mid \{ \gamma_{ji} \mid j \neq i \} \cap L = \emptyset \right] + \Pr \left[\{ \gamma_{ji} \mid j \neq i \} \cap L \neq \emptyset \right].$$

For $\Pr[\{\gamma_{ji} \mid j \neq i\} \cap L \neq \emptyset]$, recall that γ_{ji} are just random items from the pool from when the pool is of size at least $m - n^2$, hence, the probability each is an element of L is at most $\frac{|L|}{m-n^2}$. Using the assumption that $m \geq 2n^2$, we have that this is at most $\frac{2|L|}{m} \leq \frac{128n \log m + 20n^2}{m}$.

For $\Pr[b_{(q-1)i} \in L \mid \{\gamma_{ji} \mid j \neq i\} \cap L = \emptyset]$, since we are conditioning on i not being given any items in L during the giveaway phase, the only way that i will receive such items is if they pick them during round robin. For $b_{(q-1)i} \in L$, i must receive at least 2 items in L . This means that even when i made their second to last pick at Round Robin $g_{n(q-n-1)+i}$ the pool was $P_{n(q-n-1)+i-1}$, the only available items were from L , i.e., $P_{n(q-n-1)+i-1} \subseteq L$. Notice that $|P_{n(q-n-1)+i-1}| = m - n(n-1) - (n(q-n-1)+i-1) = 2n - i + 1$. We will union bound over all possible sets $S \in \binom{L}{2n-i+1}$ analyzing the probability that $\Pr[P_{n(q-n-1)+i-1} = S \mid \{\gamma_{ji} \mid j \neq i\} \cap L \neq \emptyset]$.

We can decompose $\Pr[P_{n(q-n-1)+i-1} = S \mid \{\gamma_{ji} \mid j \neq i\} \cap L \neq \emptyset] = \Pr[S \subseteq P_0 \mid \{\gamma_{ji} \mid j \neq i\} \cap L \neq \emptyset] \cdot \prod_{t=1}^{P_{n(q-n-1)+i-1}} \Pr[g_t \notin S \mid S \subseteq P_{t-1} \wedge \{\gamma_{ji} \mid j \neq i\} \cap L \neq \emptyset]$.

For the first term and all terms where $t \equiv i \pmod n$, this probability may be hard to compute, but we can upperbound it by 1. For all other terms with $t \not\equiv i \pmod n$,

$$\Pr[g_t \notin S \mid S \subseteq P_{t-1} \wedge \{\gamma_{ji} \mid j \neq i\} \cap L \neq \emptyset] = 1 - \frac{|S|}{|P_{t-1}|}.$$

Hence, this entire product is upper bounded by

$$\prod_{t=1: t \not\equiv i \pmod n}^{n(q-n-1)+i-1} \frac{n(q - (n-1)) + 1 - t - (2n - i + 1)}{n(q - (n-1)) + 1 - t}.$$

Notice that this is now the same analysis as the proof of Claim 3 with m replaced with $q(n-1)$ and $n+1-i$ replaced with $2n+1-i$. Hence, the above product is upper bounded by

$$2 \left(\frac{2}{(n(q - (n-1)))^{1-1/n}} \right)^{2n+1-i}.$$

Using the fact that $m \geq 2n^2$ so $n(q - (n-1)) \geq m/2$, this is at most

$$2 \left(\frac{4}{m^{1-1/n}} \right)^{2n+1-i}.$$

Union bounding over the at most $\binom{|L|}{2n+1-i} \leq |L|^{2n+1-i} \leq (64n \log m + 10n^2)^{2n+1-i}$ subsets S , we get that

$$\Pr[b_{(q-1)i} \mid \{\gamma_{ji} \mid j \neq i\} \cap L = \emptyset] \leq \left(\frac{256n \log m + 48n^2}{m^{1-1/n}} \right)^{2n+1-i} \leq \left(\frac{256n \log m + 48n^2}{m^{1/2}} \right)^{2n+1-i}.$$

Since we have restricted to $m \geq (256n \log m + 40n^2)^2$, $m^{1/2} \geq 256n \log m + 48n^2$, so the inside is at most 1. Hence, since $2n+1-i \geq n+1 \geq 2$, this is at most

$$\left(\frac{64n \log m + 12n^2}{m^{1/2}} \right)^2 = \frac{65536n^2 \log^2 m + 24576n^3 \log m + 2304n^4}{m}.$$

Hence, we have that

$$\Pr[b_{(q-1)i} \in L] \leq \frac{65536n^2 \log^2 m + 24576n^3 \log m + 2304n^4 + 128n \log m + 24n^2}{m}.$$

Finally, fix some agent $j \neq i$. We consider $\Pr[\bigcup_{k: q-32 \log m - 5n < k \leq q-1} \mathcal{E}_k^{ij} \mid b_{(q-1)i} \notin L]$. Notice that if $b_{(q-1)i} \notin L$, as long as $|A_j \cap L| \geq 32 \log m + 5n$, then for all $k \leq 32 \log m + 5n$, $b_{ik} \geq_i b_{i(q-1)} >_i b_{jk}$, so $\mathcal{E}_k^{ij}$ does not hold. Hence, we wish to upperbound the probability that $|A_j \cap L| < 32 \log m + 5n$. We use the second equivalent process, and consider

sampling all of $\gamma_{i'j'}$ items along with sets $R_0, \dots, R_q$ and goods picked by agent i , g_t with $t \equiv i \pmod n$, but we have not yet sampled which non- i agent receives which items in $R_0, \dots, R_q$. Notice that conditioned on $b_{i(q-1)} \notin L$, $|A_i \cap L| \leq 1$. Since $n(n-1)$ items are in the giveaway phase and $|R_0| + |R_q| = n-1$, this implies that $\left|L \cap \left(\bigcup_{\ell=1}^{q-1} R_\ell\right)\right| \geq |L| - n^2 \geq \lfloor 64n \log m + 11n^2 \rfloor \geq 64n \log m + 10n^2$. Notice that $|A_j \cap L| \geq \sum_{\ell=1}^{q-1} \mathbb{I}[A_j \cap R_\ell \cap L \neq \emptyset]$. Further, this is a sum of independent Bernoulli variables with $\Pr[A_j \cap R_\ell \cap L \neq \emptyset] = \frac{|R_\ell \cap L|}{n-1}$. Hence, $\mathbb{E}[\sum_{\ell=1}^{q-1} \mathbb{I}[A_j \cap R_\ell \cap L \neq \emptyset]] \geq \sum_{\ell=1}^{q-1} \frac{|R_\ell \cap L|}{n-1} = \frac{|L \cap (\bigcup_{\ell=1}^{q-1} R_\ell)|}{n-1} \geq 64 \log m + 10n$. We can then use a Chernoff bound to show that $\Pr[|A_j \cap L| \leq 32 \log m + 5n] \leq \Pr\left[\sum_{\ell=1}^{q-1} \mathbb{I}[A_j \cap R_\ell \cap L \neq \emptyset] \leq 32 \log m + 5n\right] \leq \Pr\left[\sum_{\ell=1}^{q-1} \mathbb{I}[A_j \cap R_\ell \cap L \neq \emptyset] \leq \frac{1}{2} \mathbb{E}\left[\sum_{\ell=1}^{q-1} \mathbb{I}[A_j \cap R_\ell \cap L \neq \emptyset]\right]\right] \leq \exp\left(-\frac{(1/2)^2 \mathbb{E}\left[\sum_{\ell=1}^{q-1} \mathbb{I}[A_j \cap R_\ell \cap L \neq \emptyset]\right]}{2}\right) \leq \exp\left(-\frac{32 \log m}{8}\right) \leq \frac{1}{m^4} \leq \frac{1}{m}$.

Hence, we have that

$$\sum_{j:j \neq i} \Pr\left[\bigcup_{k:q-32 \log m - 5n < k \leq q-1} \mathcal{E}_k^{ij} \mid b_{(q-1)i} \notin L\right] \leq \frac{n}{m}.$$

Putting this together, we have that $\Pr\left[\bigcup_{k:q-32 \log m - 5n < k \leq q-1} \bigcup_{j:j \neq i} \mathcal{E}_k^{ij}\right]$ is at most

$$\frac{65536n^2 \log^2 m + 24576n^3 \log m + 2304n^4 + 128n \log m + 24n^2 + n}{m}. \quad \square$$

Proof of Lemma 10 Let L be the last $n-1$ items according to π_i . A sufficient condition for $\bigcup_{k:k=q} \bigcup_{j:j \neq i} \mathcal{E}_k^{ij} = \bigcup_{j:j \neq i} \mathcal{E}_q^{ij}$ to hold is that $\{\gamma_{ij} \mid i \neq j\} = L$, i.e., i is able to give away their bottom $n-1$ goods. In this case, $b_{iq} \succ_i L$ while $b_{jq} \in L$ for all $j \neq i$. For this to hold, the only necessary condition is that no other agent has given away a good in L prior to i giving them away. There are at most n^2 items given away before i is able to give, and each of these is an element of L with probability at most $\frac{|L|}{m-n^2}$. Since we are assuming that $M \geq 2n^2$, this is at most $\frac{2n}{m}$. Hence,

$$\Pr\left[\bigcup_{k:k=q} \bigcup_{j:j \neq i} \mathcal{E}_k^{ij}\right] \leq \frac{2n}{m}. \quad \square$$

Appendix C: Non-existence of SD-EFX allocations

Complementing Theorem 4, in this section we show that SD-EFX allocations may not exist, even for additive valuations. Additionally, this also contrasts with the fact that the existence of EFX remains an elusive open problem.

We first consider an instance $m = 4$ items and $n = 2$ agents having identical additive valuation function $v : 2^{[4]} \rightarrow \mathbb{R}_{\geq 0}$. The function v is such that $v(\{1\}) = 4$, $v(\{2\}) = 1 + \epsilon$, $v(\{3\}) = 1$, and $v(\{4\}) = 1 - \epsilon$. Indeed, the only EFX allocations in this instance are $\mathcal{A} = (\{1\}, \{2, 3, 4\})$ and $\widehat{\mathcal{A}} = (\{2, 3, 4\}, \{1\})$. Note that an SD-EFX allocation must also be EFX, therefore, if an SD-EFX allocation exists for this instance it must be $\mathcal{A}$ or $\widehat{\mathcal{A}}$.

Furthermore, if an allocation is SD-EFX, then it must also be EFX for any other instance in which both agents have the same preference order over the items. Now consider a second instance wherein both agents instead have the additive function $\widehat{v}$ such that $\widehat{v}(\{1\}) = 4$, $\widehat{v}(\{2\}) = 3$, $\widehat{v}(\{3\}) = 2$, and $\widehat{v}(\{4\}) = 1$. Note that the preference order of items is the same in v and $\widehat{v}$, therefore, if $\mathcal{A}$ (or $\widehat{\mathcal{A}}$) is SD-EFX then it must be EFX for the second instance. It is easy to see that neither $\mathcal{A}$ nor $\widehat{\mathcal{A}}$ is EFX for the second instance. This shows that SD-EFX allocations do not exist, even for additive valuations.

Acknowledgments

Alexandros Psomas and Paritosh Verma are supported in part by an NSF CAREER award, a Google AI for Social Good award, and research awards from Google and Supra.

References

- Akrami H, Mayorov A, Mehlhorn K, Srinivas S, Weidenbach C (2026) A counterexample to EFX; $n \geq 3$ agents, $m \geq n + 5$ items, monotone valuations; via SAT-solving. arXiv:2604.18216.
- Alptekinoglu A, Benadè G (2023) Achieving rawlsian justice in food rescue. *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*.
- Amanatidis G, Markakis E, Nikzad A, Saberi A (2017) Approximation algorithms for computing maximin share allocations. *ACM Transactions on Algorithms* 13(4):1–28.
- Armony M, Ward AR (2010) Fair dynamic routing in large-scale heterogeneous server systems. *Operations Research* 58(3):624–637.
- Aziz H, Gaspers S, Mackenzie S, Walsh T (2015) Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence* 227:71–92.
- Bai Y, Feige U, Gözl P, Procaccia AD (2022) Fair allocations for smoothed utilities. *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*, 436–465.
- Bai Y, Gözl P (2022) Envy-free and pareto-optimal allocations for agents with asymmetric random valuations. *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 53–59.
- Benadè G, Halpern D, Psomas A (2022) Dynamic fair division with partial information. *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 3703–3715.
- Benadè G, Kazachkov AM, Procaccia AD, Psomas A (2018) How to make envy vanish over time. *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, 593–610.
- Benadè G, Kazachkov AM, Procaccia AD, Psomas A, Zeng D (2024) Fair and efficient online allocations. *Operations Research* 72(4):1438–1452.
- Benadè G, Procaccia AD, Tucker-Foltz J (2026) You can have your cake and redistrict it too. *Operations Research* (in press).
- Bertsimas D, Farias VF, Trichakis N (2013) Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research* 61(1):73–87.
- Blum A, Caragiannis I, Haghtalab N, Procaccia AD, Procaccia EB, Vaish R (2017) Opting into optimal matchings. *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2351–2363.
- Blum A, Gözl P (2021) Incentive-compatible kidney exchange in a slightly semi-random model. *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 138–156.
- Blum A, Spencer J (1995) Coloring random and semi-random k -colorable graphs. *Journal of Algorithms* 19(2):204–234.

- Bogomolnaia A, Moulin H (2001) A new solution to the random assignment problem. *Journal of Economic theory* 100(2):295–328.
- Bouveret S, Endriss U, Lang J (2010) Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI)*, 387–392.
- Brams SJ, Taylor AD (1996) *Fair Division: From cake-cutting to dispute resolution* (Cambridge University Press).
- Butler M, Williams HP (2002) Fairness versus efficiency in charging for the use of common facilities. *Journal of the Operational Research Society* 53(12):1324–1329.
- Cai Y, Zhao M (2017) Simple mechanisms for subadditive buyers via duality. *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC)*, 170–183.
- Caragiannis I, Gravin N, Huang X (2019) Envy-freeness up to any item with high nash welfare: The virtue of donating items. *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*, 527–545.
- Chawla S, Miller JB (2016) Mechanism design for subadditive agents via an ex ante relaxation. *Proceedings of the 2016 ACM Conference on Economics and Computation (EC)*, 579–596.
- Chen VX, Hooker JN (2023) A guide to formulating fairness in an optimization model. *Annals of Operations Research* 326(1):581–619.
- Dickerson J, Goldman J, Karp J, Procaccia A, Sandholm T (2014) The computational rise and fall of fairness. *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, volume 28.
- Dummit DS, Foote RM (2004) *Abstract algebra*, volume 3 (Wiley Hoboken).
- Farhadi A, Ghodsi M, Hajiaghayi MT, Lahaie S, Pennock D, Seddighin M, Seddighin S, Yami H (2019) Fair allocation of indivisible goods to asymmetric agents. *Journal of Artificial Intelligence Research* 64:1–20.
- Feige U, Kilian J (1998) Heuristics for finding large independent sets, with applications to coloring semi-random graphs. *Proceedings of the 39th Symposium on Foundations of Computer Science (FOCS)*, 674–683.
- Flanigan B, Halpern D, Psomas A (2023) Smoothed analysis of social choice revisited. *Proceedings of the 19th Conference on Web and Internet Economics (WINE)*, 290–309.
- Gan J, Suksompong W, Voudouris AA (2019) Envy-freeness in house allocation problems. *Mathematical Social Sciences* 101:104–106.
- Gautschi W (1959) Some elementary inequalities relating to the gamma and incomplete gamma function. *Journal of Mathematics and Physics* 38(1):77–81.
- Katz J, Lindell Y (2020) *Introduction to modern cryptography* (CRC press).
- Kolla A, Makarychev K, Makarychev Y (2011) How to play unique games against a semi-random adversary: Study of semi-random models of unique games. *Proceedings of the 52nd Symposium on Foundations of Computer Science (FOCS)*, 443–452.
- Kurokawa D, Procaccia AD, Wang J (2016) When can the maximin share guarantee be guaranteed? *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*.

- Lucier B (2017) An economic view of prophet inequalities. *ACM SIGecom Exchanges* 16(1):24–47.
- Manurangsi P, Suksompong W (2017) Asymptotic existence of fair divisions for groups. *Mathematical Social Sciences* 89:100–108.
- Manurangsi P, Suksompong W (2020) When do envy-free allocations exist? *SIAM Journal on Discrete Mathematics* 34(3):1505–1521.
- Manurangsi P, Suksompong W (2021) Closing gaps in asymptotic fair division. *SIAM Journal on Discrete Mathematics* 35(2):668–706.
- Mertzanidis M, Psomas A, Verma P (2024) Automating food drop: The power of two choices for dynamic and fair food allocation. *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*, 243–243.
- Pratt JW, Zeckhauser RJ (1990) The fair and efficient division of the winsor family silver. *Management Science* 36(11):1293–1301.
- Psomas A, Schwartzman A, Weinberg SM (2019) Smoothed analysis of multi-item auctions with correlated values. *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*, 417–418.
- Rubinstein A, Singla S (2017) Combinatorial prophet inequalities. *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1671–1687.
- Rubinstein A, Weinberg SM (2018) Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. *ACM Transactions on Economics and Computation* 6(3-4):1–25.
- Sherstyuk K (1998) How to gerrymander: A formal analysis. *Public Choice* 95(1):27–49.
- Sinclair SR, Jain G, Banerjee S, Yu CL (2023) Sequential fair allocation: Achieving the optimal envy-efficiency trade-off curve. *Operations Research* 71(5):1689–1705.
- Spielman DA, Teng SH (2004) Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM* 51(3):385–463.
- Steinhaus H (1948) The problem of fair division. *Econometrica* 16:101–104.
- Suksompong W (2016) Asymptotic existence of proportionally fair allocations. *Mathematical Social Sciences* 81:62–65.
- Vondrák J (2010) A note on concentration of submodular functions. arXiv:1005.2791.
- Wrench JW (1968) Concerning two series for the gamma function. *Mathematics of Computation* 22(103):617–626.
- Xia L (2020) The smoothed possibility of social choice. *Proceedings of the 33rd Annual Conference on Neural Information Processing Systems (NeurIPS)*, 11044–11055.
- Xia L (2021) The semi-random satisfaction of voting axioms. *Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 6075–6086.
- Zeng D, Psomas A (2020) Fairness-efficiency tradeoffs in dynamic fair division. *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 911–912.