Envy-Free and Pareto-Optimal Allocations for Agents with Asymmetric Random Valuations

Yushi Bai[1] and **Paul Gölz**[2]
[1]Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China
[2]Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, USA
bys18@mails.tsinghua.edu.cn, pgoelz@cs.cmu.edu

Abstract

We study the problem of allocating m indivisible items to n agents with additive utilities. It is desirable for the allocation to be both fair and efficient, which we formalize through the notions of *envy-freeness* and *Pareto-optimality*. While envy-free and Pareto-optimal allocations may not exist for arbitrary utility profiles, previous work has shown that such allocations exist with high probability assuming that all agents' values for all items are independently drawn from a common distribution. In this paper, we consider a generalization of this model where each agent's utilities are drawn independently from a distribution *specific to the agent*. We show that envy-free and Pareto-optimal allocations are likely to exist in this asymmetric model when $m = \Omega\left(n \log n\right)$, which is tight up to a log log gap that also remains open in the symmetric subsetting. Furthermore, these guarantees can be achieved by a polynomial-time algorithm.

1 Introduction

Imagine that the neighborhood children go trick-or-treating and return successfully, with a large heap of candy between them. They then try to divide the candy amongst themselves, but quickly reach the verge of a fight: Each has their own conception of which sweets are most desirable, and, whenever a child suggests a way of splitting the candy, another child feels unfairly disadvantaged. As a (mathematically inclined) adult in the room, you may wonder: Which allocation of candies should you suggest to keep the peace? And, is it even possible to find such a fair distribution?

In this paper, we study the classic problem of fairly dividing m items among n agents [Bouveret *et al.*, 2016], as exemplified by the scenario above. We assume that the items we seek to divide are goods (i.e., receiving an additional piece of candy never makes a child less happy), that items are indivisible (candy cannot be split or shared), and that the agents have *additive* valuations (roughly: a child's value for a piece of candy does not depend on which other candies they receive).

We will understand an allocation to be fair if it satisfies two axioms: *envy-freeness* (EF) and *Pareto-optimality* (PO). First, fair allocations should be envy-free, which means that no agent should strictly prefer another agent's bundle to their own. After all, if an allocation violates envy-freeness, the former agent has good reason to contest it as unfair. Second, fair allocations should be Pareto-optimal, i.e., there should be no reallocation of items making some agent strictly better off and no agent worse off. Not only does this axiom rule out allocations whose wastefulness is unappealing; it is also arguably necessary to preserve envy-freeness: Indeed, if a chosen allocation is envy-free but not Pareto-optimal, rational agents can be expected to trade items after the fact, which might lead to a final allocation that is not envy-free after all. Unfortunately, even envy-freeness alone is not always attainable. For instance, if two agents like a single item, the agent who does not receive it will always envy the agent who does.

Motivated by the fact that worst-case allocation problems may not have fair allocations, a line of research in fair division studies asymptotic conditions for the existence of such allocations, under the assumption that the agents' utilities are random rather than adversarially chosen [e.g., Dickerson *et al.* 2014; Manurangsi and Suksompong 2019]. Specifically, these papers assume that all agents' utilities for all items are independently drawn from a common distribution $\mathcal{D}$, a model which we call the *symmetric model*. Among the algorithms shown to satisfy envy-freeness in this setting, only one is also Pareto-optimal: the (utilitarian) *welfare-maximizing algorithm*, which assigns each item to the agent who values it the most. This algorithm is Pareto-optimal, and it is also envy-free with high probability as the number of items m grows in $\Omega\left(n \log n\right)$.[1] Since envy-free allocations may exist with only vanishing probability for $m \in \Theta\left(n \log n/\log\log n\right)$ in the symmetric model [Manurangsi and Suksompong, 2019], the above result characterizes almost tightly when envy-free and Pareto-optimal allocations exist in this model.

Zooming out, however, this positive result is unsatisfying in that, outside of this specific random model, the welfare-maximizing algorithm can hardly be called "fair": For example, if an agent A tends to have higher utility for most items than agent B, the welfare-maximizing algorithm will allocate most items to agent A, which can cause large envy for agent B. In short, the welfare-maximizing algorithm leads

[1]In fact, Dickerson *et al.* [2014] prove this result for a somewhat more general model than the one presented above (certain correlatedness between distributions is also allowed), but their model assumes the key symmetry between agents that we discuss below.

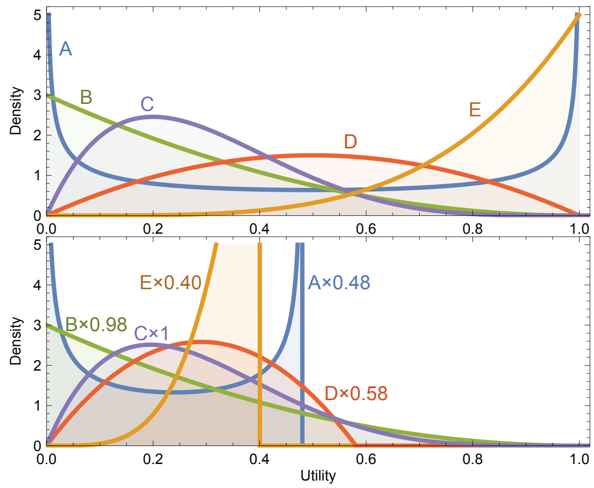

Figure 1: The top panel shows probability density functions of five agents' utility distributions. The bottom panel shows densities after scaling distributions by the given multipliers. When drawing an independent sample from each scaled distribution, each sample is the largest with probability $1/5$.

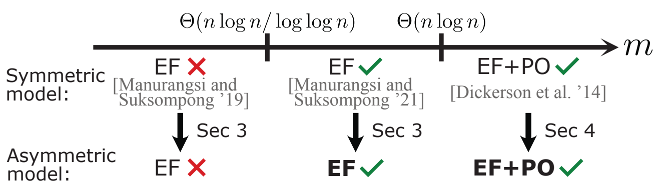

Figure 2: Existing and new results on when EF and EF+PO allocations are guaranteed to exist in both models. Bold results are new.

to fair allocations only because the model assumes each agent to be equally likely to have the largest utility, an assumption that limits the lessons that can be drawn from this model.

Motivated by these limitations of prior work, this paper investigates the existence of fair allocations in a generalization of the symmetric model, which we refer to as the *asymmetric model*. In this model, each agent i is associated with their own distribution $\mathcal{D}_i$, from which their utility for all items is independently drawn. Within this model, we aim to answer the question: *When do envy-free and Pareto-optimal allocations exist for agents with asymmetric valuations?*

1.1 Our Techniques and Results

In Section 3, we study which results in the symmetric model generalize to the asymmetric model. In particular, we apply an analysis by Manurangsi and Suksompong [2021] to the asymmetric model in a black-box manner to prove envy-free allocations exist when $m \in \Omega(n \log n / \log \log n)$, which is tight with existing impossibility results on envy-freeness. However, this approach does not preserve Pareto-optimality.

Using a new approach, we prove in Section 4 that generalizing the random model from symmetric to asymmetric agents does not substantially decrease the frequency of envy-free and Pareto-optimal allocations. The key idea is to find a multiplier $\beta_i > 0$ for each agent i such that, when drawing an independent sample u_i from each utility distribution $\mathcal{D}_i$, each agent i has an equal probability of $\beta_i u_i$ being larger than the $\beta_j u_j$ of all other agents $j \neq i$, which we call the agent's *resulting probability* from these multipliers. Fig. 1 illustrates how five utility distributions can be rescaled in this way. If all resulting probabilities of a set of multipliers equal $1/n$, we call these multipliers *equalizing*.

A set of equalizing multipliers defines what we call its *multiplier allocation*, which assigns each item to the agent i whose utility weighted by β_i is the largest. Put differently, the multiplier allocation simulates the welfare-maximizing algorithm in an instance in which each agent i's distribution is scaled by β_i. Just like the welfare-maximizing allocations, the multiplier allocations are Pareto-optimal by construction, and the similarity between both allocation types allows us to apply proof techniques developed for the welfare-maximizing algorithm and the symmetric setting to show envy-freeness.

The core of our paper is a proof that equalizing multipliers always exist, which we show using Sperner's lemma. Since an algorithm based on this direct proof would have exponential running time, we design a polynomial-time algorithm for computing approximately equalizing multipliers, i.e., multipliers whose resulting probabilities lie within $[1/n - \delta, 1/n + \delta]$ for a $\delta > 0$ given in the input.

Having established the existence of equalizing multipliers, we go on to show that the multiplier allocation is envy-free with high probability. To obtain this result, we demonstrate a constant-size gap between each agent's expected utility for an item conditioned on them receiving the item and the agent's expected utility conditioned on another agent receiving the item, and then use a variant of the argument of Dickerson *et al.* [2014] to show that multiplier allocations are envy-free with high probability when $m \in \Omega(n \log n)$. This guarantee extends to the case where we allocate based on multipliers that are sufficiently close to equalizing, which means that our polynomial-time *approximate multiplier algorithm* is Pareto-optimal and envy-free with high probability.

In Section 5, we empirically evaluate how many items are needed to guarantee envy-free and Pareto-optimal allocations for a fixed collection of agents. We find that the approximate multiplier algorithm needs relatively large numbers of items to ensure envy-freeness; that the round robin algorithm violates Pareto-optimality in almost all instances; and that the Maximum Nash Welfare (MNW) algorithm achieves both axioms already for few items but that its running time limits its applicability. For larger numbers of items, the approximate multiplier algorithm satisfies both axioms and excels by virtue of its running time.

1.2 Related Work

The question of when fair allocations exist for random utilities was first raised by Dickerson *et al.* [2014], whose main result we have already discussed. Our paper also builds on work by Manurangsi and Suksompong; Manurangsi and Suksompong [2019; 2021], who prove the lower bound on the existence of envy-free allocations mentioned in the introduction and that the classic round robin algorithm produces envy-free allocations in the symmetric model for slightly lower m than the welfare-maximizing algorithm. A bit further afield, Suksompong [2016] and Amanatidis *et al.* [2017] study the existence of proportional and maximin-share allocations (two relaxations of envy-freeness) in the symmetric model, and Manurangsi and Suksompong [2017] study envy-freeness when items are allocated to groups rather than to individuals. None of these papers consider Pareto-optimality, perhaps because fair division yields few tools for simultaneously guaranteeing envy-freeness and Pareto-optimality.

The asymmetric model we investigate has been previously used, for example, by Kurokawa *et al.* [2016] to study the existence of maximin-share allocations. While part of their proof applies the results by Dickerson *et al.* to construct envy-free allocations in the asymmetric model, as do we, their allocation algorithm is not Pareto-optimal (see Section 3). Farhadi *et al.* [2019] also consider maximin-share allocations in the asymmetric model, for agents with weighted entitlements. Zeng and Psomas [2020] study allocation problems in the asymmetric model, when items arrive online. While they do consider and achieve Pareto-optimality, they only obtain approximate notions of envy-freeness. Finally, Bai *et al.* [2022] study the existence of envy-free allocations and of both proportional and Pareto-optimal allocations in an expressive utility model based on smoothed analysis.

2 Preliminaries

General Definitions. We consider a set M of m indivisible items being allocated to a group $N = \{1, \ldots, n\}$ of n agents. Each agent $i \in N$ has a *utility* $u_i(\alpha) \geq 0$ for each item $\alpha \in M$, indicating their degree of preference for the item. The collection of agent–item utilities make up a *utility profile*. An *allocation* $\mathcal{A} = \{A_i\}_{i \in N}$ is a partition of the items into n *bundles*: $M = A_1 \cup \cdots \cup A_n$, where agent i gets the items in bundle A_i. Under our assumption that the agents' utilities are *additive*, agent i's *utility* for a subset of items $A \subseteq M$ is $u_i(A) = \sum_{\alpha \in A} u_i(\alpha)$.

An allocation $\mathcal{A} = \{A_i\}_{i \in N}$ is said to be *envy-free* (EF) if $u_i(A_i) \geq u_i(A_j)$ for all $i, j \in N$, i.e., if each agent weakly prefers their own bundle to any other agent's bundle. We say that an allocation $\mathcal{A} = \{A_i\}_{i \in N}$ is *Pareto dominated* by another allocation $\mathcal{A}' = \{A'_i\}_{i \in N}$ if $u_i(A_i) \leq u_i(A'_i)$ for all $i \in N$, with at least one inequality holding strictly. An allocation is *Pareto-optimal* (PO) if it is not Pareto dominated by any other allocation. An allocation is called *fractionally Pareto-optimal* (fPO) if it is not even Pareto dominated by any "fractional" allocation of items. For our purposes, it suffices to note that an allocation is fPO iff there exist multipliers $\{\beta_i > 0\}_{i \in N}$ such that each item α is allocated to an agent i with maximal $\beta_i u_i(\alpha)$ [Negishi, 1960].

Asymmetric Model. In our asymmetric model, each agent i is associated with a *utility distribution* $\mathcal{D}_i$, a nonatomic probability distribution over $[0, 1]$. The model assumes that the utilities $u_i(\alpha)$ for all $\alpha \in M$ are independently drawn from $\mathcal{D}_i$. For simplicity, we just write u_i as a random variable for $u_i(\alpha)$ if we are not talking about a specific item α, where $u_i \sim \mathcal{D}_i$. Let f_i and F_i denote the probability density function (PDF) and cumulative distribution function (CDF) of $\mathcal{D}_i$. For our main result, we make the following assumptions on utility distributions: *(a) Interval support*: The support of each $\mathcal{D}_i$ is an interval $[a_i, b_i]$ for $0 \leq a_i < b_i \leq 1$. *(b) (p, q)-PDF-boundedness*: For constants $0 < p < q$, the density of each $\mathcal{D}_i$ is bounded between p and q within its support. These two assumptions are weaker than those by Manurangsi and Suksompong [2021], who additionally require all distributions to have support $[0, 1]$. A random event occurs *with high probability* if the event's probability converges to 1 as $n \to \infty$.

3 Takeaways From the Symmetric Model

We quickly review results obtained in the symmetric model, and to which degree they carry over to the asymmetric model.

Non-Existence of EF Allocations: Since the symmetric model is a special case of the asymmetric model — in which all $\mathcal{D}_i$ are equal — this negative result immediately applies:

Proposition 1 (Manurangsi and Suksompong 2019).[2] *There exists $c > 0$ such that, if $m = (\lfloor c \log n / \log \log n \rfloor + 1/2)\, n$ and all utility distributions are uniform on $[0, 1]$, then, with high probability, no envy-free allocation exists.*

Existence of EF Allocations: In the symmetric model, Manurangsi and Suksompong [2021] give an allocation algorithm, round robin, that satisfies EF with high probability. An interesting property of this algorithm is that an agent's allocation given a utility profile depends not on the *cardinal* information of the agents' utilities, but only on their *ordinal* preference order over items. Using this property, we prove in the full version that their result generalizes to the asymmetric model since, in a nutshell, an agent i's envy of the other agents is indistinguishable between the asymmetric model and a symmetric model with common distribution $\mathcal{D}_i$.

Proposition 2. *When distributions have interval support and are (p, q)-PDF-bounded, if $m \in \Omega\,(n \log n / \log \log n)$, an envy-free allocation exists with high probability.*

To our knowledge, we are the first to observe that the analysis by Manurangsi and Suksompong generalizes in this way, which improves on the previously best known upper bound of $m \in \Omega\,(n \log n)$ in the asymmetric model due to Kurokawa *et al.* [2016].

Existing Approaches Do Not Provide EF+PO: Generalizing the existence result for EF and PO allocations by Dickerson *et al.* [2014] to the asymmetric model is more challenging than the round robin result above, since cardinal information is crucial for the PO property. In the full version of the paper, we illustrate this point by considering how Kurokawa *et al.* [2016] apply the theorem of Dickerson *et al.* to prove the existence of EF allocations in the asymmetric model; namely, they assign each item to the agent for whom the item is in the highest percentile of their utility distribution. On an example, we show that this approach fundamentally violates PO, and that assigning items based on multipliers is the most natural way to guarantee PO. In the full version, we also give an example showing that normalizing each agent's values to add up to one — perhaps the most obvious way to obtain multipliers — is not sufficient to provide EF.

[2] Here, we present a special case; the original result holds for different choices of distribution and leaves some flexibility in m.

4 Existence of EF+PO Allocations

We now prove our main theorem:

Theorem 3. *Suppose that all utility distributions have interval support and are (p,q)-PDF-bounded for some p,q. If $m \in \Omega(n \log n)$ as $n \to \infty$,[3] then, with high probability, an envy-free and (fractionally) Pareto-optimal allocation exists and can be found in polynomial time.*

In Section 4.1, we prove that we can always find multipliers $\{\beta_i\}_{i\in N}$ that equalize each agent's probability of receiving a random-utility item in the multiplier allocation (which allocates item α to the agent with maximal $\beta_i u_i(\alpha)$ and is trivially fPO). We also discuss how to efficiently find multipliers leading to approximately equalizing probabilities. Next, in Section 4.2, we show that an agent's expected utility for an item allocated to themselves is larger by a constant than their expected utility for an item allocated to another agent. In Section 4.3, we combine these properties to prove envy-freeness.

4.1 Existence of Equalizing Multipliers

For a set of multipliers $\vec{\beta} \in \mathbb{R}^n_{>0}$ and an agent i, we denote i's resulting probability by

$$\begin{aligned} p_i(\vec{\beta}) &:= \mathbb{P}[\beta_i u_i = \max\nolimits_{j\in N} \beta_j u_j] \\ &= \int_0^1 f_i(u) \prod\nolimits_{j\in N\setminus\{i\}} F_j(\beta_i/\beta_j\, u)\, du. \end{aligned} \tag{1}$$

Existence Proof Using Sperner's Lemma

The existence of equalizing multipliers can be established quite easily using Sperner's lemma:

Theorem 4. *For any set of utility distributions, there exists a set of equalizing multipliers.*

Proof sketch. Since scaling all multipliers by the same factor does not change the resulting probabilities, we may restrict our focus to multipliers within the $(n-1)$-dimensional simplex $S = \{\vec{\beta} \in \mathbb{R}^n_{\geq 0} \mid \sum_{i\in N} \beta_i = 1\}$. We define a coloring function $f : S \to N$, which maps each set of multipliers in the simplex to an agent with maximum resulting probability. Clearly, points on a face $\beta_i = 0$ are not colored with color belonging to agent i since some other agent has a positive multiplier and must thus have a greater scaled utility than i.

Now, consider a simplicization of S, i.e., a partition of S into small simplices meeting face to face (generalizing the notion of a triangulation in the 2-D simplex). Sperner's lemma shows the existence of a small simplex that is panchromatic, i.e., whose n vertices are each colored with a different agent. This small simplex constitutes a neighborhood of multipliers such that, for each agent i, there is a set of multipliers $\vec{\beta}$ in this neighborhood such that agent i's resulting probability is larger than that of any other agent, and, as a consequence, such that i has a resulting probability of at least $1/n$.

By successively refining the simplicization, we can make these neighborhoods arbitrarily small. In the full version, we prove the existence of a set of exactly equalizing multipliers, which follows from the Bolzano-Weierstraß Theorem and continuity of the functions p_i on $\mathbb{R}^n_{>0}$. □

An Approximation Algorithm for Equalizing Multipliers

The proof above is succinct, but not particularly helpful in finding equalizing multipliers computationally.[4] Though the application of Sperner's lemma can be turned into an approximation algorithm, using it to find multipliers such that all resulting probabilities lie within $\delta > 0$ of $1/n$ requires $poly(n)\, n!\, \left(\log(2\,q)\, \left(1 + \frac{4\,n\,q}{\delta}\right)\right)^n$ time (see full version). This large runtime complexity points to a more philosophical shortcoming of our proof of Theorem 4, namely, that it does very little to elucidate the structure of how multipliers map to resulting probabilities. Given that the proof barely made use of any properties of the p_i other than continuity, it is natural that the resulting algorithm resembles a complete search over the space of multipliers.

By contrast, our polynomial-time algorithm for finding approximately-equalizing multipliers will be based on three structural properties of the p_i (proved in the full version):

Local monotonicity: If we change the multipliers from $\vec{\beta}$ to $\vec{\beta}'$, and if agent i's multiplier increases by the largest factor ($\beta'_i/\beta_i = \max_{j\in N} \beta'_j/\beta_j$), then i's resulting probability weakly increases.

Bounded probability change: If we change a set of multipliers by increasing i's multiplier by a factor of $(1+\epsilon)$ for some $\epsilon > 0$ and leaving all other multipliers equal, then i's resulting probability increases by at most $2\,q\,\epsilon$.

Bound on multipliers: If i's multiplier is at least $2\,q$ times as large as j's multiplier, then i must have a strictly larger resulting probability than j.

Crucially, we can combine the first two properties to control how the resulting probabilities evolve while changing the multipliers in a specific "step" operation, which is the key building block of our approximation algorithm:

Step guarantee: If we change a set of multipliers by increasing the multipliers of a subset S of agents by a factor of $(1+\epsilon)$ while leaving the other multipliers unchanged, then (a) the resulting probabilities of all $i \in S$ weakly increase, but at most by $2\,q\,\epsilon$, and (b) the resulting probabilties of all $i \notin S$ weakly decrease, also by at most $2\,q\,\epsilon$ (proof in full version).

Algorithm 1 keeps track of a set of multipliers $(1+\epsilon)^{\vec{z}} = ((1+\epsilon)^{z_1}, \ldots, (1+\epsilon)^{z_n})^\mathsf{T}$. In each loop iteration, we use the step operation to increase all resulting probabilities originally below $1/n$ and decrease all resulting probabilities originally above $1/n$, both by a bounded amount so that they cannot overshoot $1/n$ by too much. After polynomially many steps, all resulting probabilities lie within a band around $1/n$, which means that the multipliers are approximately equalizing.

Theorem 5. *In time $\mathcal{O}(n^2\,q\,\log(q)\,\delta^{-1})$, Algorithm 1 computes a vector of multipliers $\vec{\beta}$ such that, for all $i \in N$, $1/n - \delta \leq p_i(\vec{\beta}) \leq 1/n + \delta$.*

[3] Alternatively, we may assume $n \in O(m/\log m)$ as $m \to \infty$ to avoid the assumption that $n \to \infty$, as do Dickerson *et al.* [2014].

[4] When measuring running time, we assume that the algorithm has access to an oracle allowing it to compute the $p_i(\vec{\beta})$ for a given $\vec{\beta}$ in constant time. This choice abstracts away from the distribution-specific cost and accuracy of computing the integral in Eq. (1).

Algorithm 1: Equalizing Multipliers

Input : An oracle to compute resulting probabilities, a constant $0 < \delta \leq 1$, and a PDF upper bound q

Output: A vector of multipliers $\vec{\beta} \in \mathbb{R}^n_{>0}$

$\epsilon \leftarrow \delta/(2\,q)$;

$\vec{z} \leftarrow \vec{0}$;

while $\exists i \in N.\ \left|p_i\big((1+\epsilon)^{\vec{z}}\big) - 1/n\right| > \delta$ **do**

  $S \leftarrow \{i \in N \mid p_i\big((1+\epsilon)^{\vec{z}}\big) \leq 1/n\}$;

  $\vec{z} \leftarrow \vec{z} + \mathbb{1}_S$;

return $(1+\epsilon)^{\vec{z}}$

Proof. At the beginning of an iteration of the loop, partition the agents into three sets Z_ℓ, Z_m, and Z_h depending on whether $p_i\big((1+\epsilon)^{\vec{z}}\big)$ is smaller than $1/n - \delta$, is in $[1/n - \delta, 1/n + \delta]$, or is larger than $1/n + \delta$, respectively. We make two observations: **(a)** Once an agent is in Z_m, they will always stay there since, by the step guarantee, their probability moves by at most $\delta = 2\,q\,\epsilon$ per iteration and moves up whenever it was below $1/n$ and down whenever it was above $1/n$. **(b)** Agents cannot move between Z_ℓ and Z_h within one iteration, since the probabilities belonging to Z_ℓ and Z_h are separated by a gap of size $2\,\delta$, whereas the step guarantee shows that an agents' probability moves by at most δ.

Next, we show that the algorithm terminates; specifically, that it exits the loop after at most $T := \lceil \frac{\log(2\,q)}{\log(1+\epsilon)} \rceil \cdot (n-1)$ iterations. For the sake of contradiction, suppose that at the beginning of the $(T+1)$th iteration of the loop, some agent was not yet in Z_m. For now, say that such an i was in Z_h and let i have maximal $p_i\big((1+\epsilon)^{\vec{z}}\big)$. Then, since i has always been in Z_h, z_i has never been increased and it still holds that $z_i = 0$. At the same time, in each round, the multipliers of some $|S| \geq 1$ other agents get increased, from which it follows that some other agent j must have $z_j \geq \lceil \frac{\log(2\,q)}{\log(1+\epsilon)} \rceil$. Then, $\beta_j/\beta_i = (1+\epsilon)^{z_j - z_i} \geq 2\,q$, which implies that $p_j > p_i$ by the bounds-on-multipliers property, which contradicts our choice of i. The case where $i \in Z_\ell$ is symmetric: This time, choose an i with minimal $p_i\left((1+\epsilon)^{\vec{z}}\right)$. Since $i \in Z_\ell$, z_i must have been increased in every round and equal T. Furthermore, in each previous round, $|S| \leq n-1$, since the algorithm would not have re-entered the loop if all probabilities were $1/n$, suggesting that some agent's probability is larger than $1/n$ and is thus not included in S. Hence, there must be another agent j with $z_j \leq T - \lceil \frac{\log(2\,q)}{\log(1+\epsilon)} \rceil$. This implies that $p_j < p_i$, contradicting our choice of i.

It follows that the loop is executed at most T times. Taking into account that each iteration requires $\mathcal{O}(n)$ oracle queries, the total time complexity is in $T \cdot \mathcal{O}(n) \in \mathcal{O}(n^2\,q\,\log(q)/\delta)$.[5] The bound on the resulting probabilities follows from the fact that $Z_m = N$ when the algorithm exits. □

As another demonstration of the rich structure in the p_i, we show in the full version that the equalizing multipliers are unique, using a strengthened local-monotonicity property.

[5] $T \cdot \mathcal{O}(n) = \lceil \log(2\,q)/\log(1+\delta/(2\,q)) \rceil\,(n-1)\,\mathcal{O}(n) \leq \lceil \log(2\,q)\,((4\,q)/\delta) \rceil\,(n-1)\,\mathcal{O}(n) \in \mathcal{O}(n^2\,q\,\log(q)/\delta)$, where the inequality holds since $\delta/(2\,q) \leq 1$.

4.2 Gap between Expected Utilities

Having established the existence of (approximately) equalizing multipliers $\vec{\beta}$, we will now analyze the corresponding multiplier allocation, which assigns each item α to the agent i with maximal $\beta_i\,u_i(\alpha)$. By definition, this allocation satisfies fPO and thus PO, so it remains to show EF. In our exposition, we will focus on exactly equalizing multipliers, but all observations extend to multipliers that are "sufficiently close" to equalizing, which we make explicitly in Proposition 6.

As sketched in Section 1.1, we now prove that an agent i's expected utility for an item they receive themselves is strictly larger than i's expected utility for an item that another agent receives in the multiplier allocation. In fact, we will prove that there is a *constant* gap between these conditional expectations i.e., a constant $C_{p,q} > 0$ such that, for all $i \neq j \in N$, $\mathbb{E}\left[u_i \,|\, \beta_i u_i = \max_{k \in N} \beta_k u_k\right] \geq C_{p,q} + \mathbb{E}\left[u_i \,|\, \beta_j u_j = \max_{k \in N} \beta_k u_k\right]$. Bounding this gap is the main idea of the proof by Dickerson *et al.* Their proof approach is applicable since, by scaling the utilities by equalizing multipliers, we bring a key property exploited by Dickerson *et al.* to the asymmetric model: as does the welfare-maximizing algorithm in the symmetric setting, the multiplier allocation gives a random item to each agent with equal probability. Thus, by concentration, all agents receive similar numbers of items. A positive gap $C_{p,q}$ furthermore ensures that agents prefer the average item in their own bundle to the average item in another bundle. The last two statements imply that the allocation is likely to be envy-free.

In the full version, we show that some utility distributions do not yield a constant gap, and we motivate our assumptions, interval support and (p,q)-PDF-boundedness, using such distributions. In the full version, we derive the desired gap:

Proposition 6. *For any collection of agents whose utility distributions are (p,q)-PDF-bounded and have interval support, given a set of multipliers $\vec{\beta}$ such that $|p_i - 1/n| < 1/(2\,n)$ for all $i \in N$, it holds that for any $i \neq j \in N$,*

$$\mathbb{E}\left[u_i \middle| \beta_i u_i = \max_{k \in N} \beta_k u_k\right] - \mathbb{E}\left[u_i \middle| \beta_j u_j = \max_{k \in N} \beta_k u_k\right] \geq C_{p,q}$$

for a constant $C_{p,q} \in (0,1]$ that only depends on p and q.

4.3 The Multiplier Allocation Satisfies EF

In the full version, we combine the existence of equalizing multipliers and the positive gap to prove Theorem 3, i.e., that the multiplier allocation is EF with high probability, and that this even holds when assigning based on approximately equalizing multipliers. Here, we sketch the argument: First, we run Algorithm 1 to find approximately equalizing multipliers $\vec{\beta}$ with an accuracy $\delta := C_{p,q}/(4\,n)$, which requires $\mathcal{O}(n^2/\delta) \subseteq \mathcal{O}(n^3)$ time by Theorem 5. Second, we allocate all items based on $\vec{\beta}$, in $\mathcal{O}(m\,n)$ time). This yields the *approximate multiplier allocation*, which is fPO

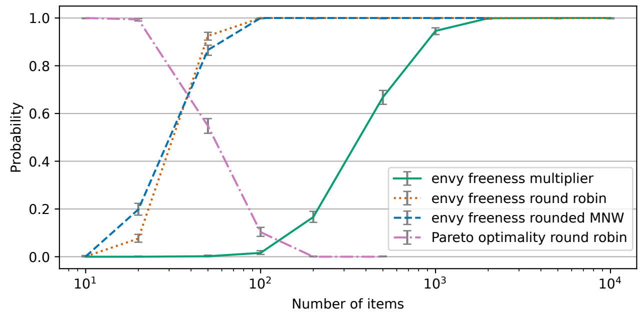

Figure 3: Probability of different algorithms satisfying EF and PO for $n = 10$ distributions. The multiplier and rounded MNW algorithms are always PO and therefore not shown. Each datapoint corresponds to 1 000 random instances, and flies indicate 95% confidence intervals. PO and MNW data not available for large m due to computational cost.

by construction. It remains to show EF: Proposition 6 applies to $\vec{\beta}$ since it satisfies the proposition's precondition: $|p_i - 1/n| \leq \delta \leq \frac{1}{4\,n} < \frac{1}{2\,n}$. By arithmetic, the positive gap guaranteed by the proposition implies that, for any two agents $i \neq j$, $\mathbb{E}[u_i(A_i)] - \mathbb{E}[u_i(A_j)] \geq \frac{m}{2\,n} C_{p,q}$. Assuming that $m \in \Omega(n \log n)$, we prove that, with high probability, all $u_i(A_j)$ stay within a distance of $\frac{m}{4n} C_{p,q}$ from their expectations by concentration. Then,

$$\begin{aligned} u_i(A_i) &\geq \mathbb{E}[u_i(A_i)] - \tfrac{m}{4n} C_{p,q} \\ &\geq \big(\mathbb{E}[u_i(A_j)] + \tfrac{m}{2n} C_{p,q}\big) - \tfrac{m}{4n} C_{p,q} \\ &\geq \mathbb{E}[u_i(A_j)] + \tfrac{m}{4n} C_{p,q} \geq u_i(A_j), \end{aligned}$$

which implies that i does not envy j for any i and j, i.e., the allocation is envy-free (and Pareto optimal) as claimed.

5 Empirical Results

After characterizing the existence of EF and PO allocations from an asymptotic angle, we now empirically investigate allocation problems for a concrete set of agents. We use a set of ten agents with utility distributions from a simple parametric family of $(0.1, 1.9)$-PDF-bounded distributions.[6] We compute multipliers for these ten distributions by implementing a variant of Algorithm 1. Specifically, we repeatedly run the algorithm with exponential decreasing δ, starting each iteration from the last set of multipliers, which allows the algorithm to change multipliers faster in the first rounds and empirically leads to a sublinear running time in δ^{-1}. For a requested accuracy of $\delta = 10^{-5}$, our algorithm runs in 30 seconds on consumer hardware, and we verify analytically that the resulting multipliers indeed lie within this tolerance, undisturbed by numerical inaccuracies in the computation.

As shown by the solid line in Fig. 3, the multiplier allocation requires large numbers of items to be reliably EF: When allocating $m = 500$ items to the ten agents, the allocation is EF in only 67% of instances, and it requires $m = 2\,000$ items for this probability to reach 99%. This slow speed of convergence seems to be inherent to the argument of Dickerson *et al.* [2014] since, in an instance with ten copies of one of our distributions and 500 random items, the welfare-maximizing algorithm is also only EF with 87% probability.

In contrast to the approximate-multiplier algorithm, the round robin algorithm (dotted line) reliably obtains EF allocations already for $m \geq 100$, but its allocations are essentially never PO unless m is very small (dash-dotted line). This lack of PO matches our theoretical predictions in the full version.

If one searches for an algorithm that satisfies both EF and PO for small numbers of items, variants of the Maximum Nash Welfare (MNW) algorithm appear promising in our experiments — unless their computational complexity is prohibitive. In experiments in the full version with only five agents, the optimization library BARON can reliably find the (discrete) MNW allocation for small $m \leq 200$. The MNW allocation is automatically PO, and it satisfies EF as reliably as round robin in our experiments. For our 10 agents, however, and as little as $m = 20$ items, BARON often takes multiple minutes to compute a single allocation, making this algorithm intractable for our analysis. In the full version, we discuss approaches to this intractability, and propose to round the *fractional* MNW allocation instead. This approach is still guaranteed to be PO, and yields EF allocations already for small m (dashed line). In a sense, this rounded MNW algorithm complements the approximate multiplier algorithm: For small m, rounded MNW already provides EF and its runtime is acceptable. For large m, the approximate multiplier algorithm guarantees EF while its runtime scales blazingly fast in m, since almost all work happens in the determination of the multipliers, independently of m.

6 Discussion

In this paper, we show that EF and PO allocations are likely to exist for random utilities even if different agents' utilities follow different distributions. Given that the known asymptotic bounds for the existence of EF+PO allocations are equal in the asymmetric and in the symmetric model, we see no evidence that the asymmetry of agent utilities would make EF+PO allocations substantially rarer to exist, up to, possibly, a $\log \log n$ gap that remains open in both models.

The most interesting idea coming out of this paper is the technique of finding equalizing multipliers, which might be of use in wider settings. Notably, the existence proof based on Sperner's lemma mainly uses the continuity of the function mapping multipliers to probabilities, and in particular does not use the independence between the agents' utilities. Thus, the multiplier technique might apply to random models where the agents' utilities exhibit some correlation, as long as the gap in expected utilities can still be bounded. In the limit of infinitely many items, we can think of the multiplier technique as a way to find an allocation of divisible goods that is Pareto-optimal and *balanced*, i.e., where every agent receives an equal amount of items. In future work, we hope to explore if this construction extends to arbitrary sets of divisible items.

[6] The full version contains all details on the experiments. Since EF allocations exist for smaller m if n divides m, we repeat the experiment with shifted values of m, which does not change the major trends. We also repeat the experiments with the five distributions from Fig. 1, showing that our observations generalize to extremely heterogeneous distributions that are not (p, q)-PDF-bounded.

Acknowledgments

We thank Bailey Flanigan and Ariel Procaccia for valuable comments and suggestions on the paper. We also thank Dravyansh Sharma, Jamie Tucker-Foltz, Ruixiao Yang, Xingcheng Yao, and Zizhao Zhang for helpful technical discussions.

References

[Amanatidis *et al.*, 2017] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi. Approximation algorithms for computing maximin share allocations. *ACM Transactions on Algorithms*, 13(4):1–28, 2017.

[Bai *et al.*, 2022] Yushi Bai, Uriel Feige, Paul Gölz, and Ariel D. Procaccia. Fair allocations for smoothed utilities. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, 2022.

[Bouveret *et al.*, 2016] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. Fair allocation of indivisible goods. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, pages 284–310. Cambridge University Press, 2016.

[Dickerson *et al.*, 2014] John P. Dickerson, Jonathan Goldman, Jeremy Karp, Ariel D. Procaccia, and Tuomas Sandholm. The computational rise and fall of fairness. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, pages 1405–1411, 2014.

[Farhadi *et al.*, 2019] Alireza Farhadi, Mohammad Ghodsi, MohammadTaghi Hajiaghayi, Sébastien Lahaie, David Pennock, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods to asymmetric agents. *Journal of Artificial Intelligence Research*, 64:1–20, 2019.

[Kurokawa *et al.*, 2016] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. When can the maximin share guarantee be guaranteed? In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, pages 523–529, 2016.

[Manurangsi and Suksompong, 2017] Pasin Manurangsi and Warut Suksompong. Asymptotic existence of fair divisions for groups. *Mathematical Social Sciences*, 89:100–108, 2017.

[Manurangsi and Suksompong, 2019] Pasin Manurangsi and Warut Suksompong. When do envy-free allocations exist? *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, pages 2109–2116, 2019.

[Manurangsi and Suksompong, 2021] Pasin Manurangsi and Warut Suksompong. Closing gaps in asymptotic fair division. *SIAM Journal on Discrete Mathematics*, 35(2):668–706, 2021.

[Negishi, 1960] Takashi Negishi. Welfare economics and existence of an equilibrium for a competitive economy. *Metroeconomica*, 12(2-3):92–97, 1960.

[Suksompong, 2016] Warut Suksompong. Asymptotic existence of proportionally fair allocations. *Mathematical Social Sciences*, 81:62–65, 2016.

[Zeng and Psomas, 2020] David Zeng and Alexandros Psomas. Fairness-efficiency tradeoffs in dynamic fair division. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 911–912, 2020.