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

Yushi Bai¹ and Paul Gölz²

¹Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China ²Computer 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 envyfreeness 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 $m = \Omega(n \log n)$, 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 envyfree with high probability as the number of items m grows in $\Omega(n \log n)$.¹ Since envy-free allocations may exist with only vanishing probability for $m \in \Theta(n \log n / \log \log n)$ 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 welfaremaximizing 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

¹In 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 envyfree 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 Paretooptimal 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* $u_i(\alpha) \ge 0$ 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 $0 , the density of each <math>\mathcal{D}_i$ 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 Appendix A 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 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 Appendix B, 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 Appendix C, 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.

²Here, 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

$$p_i(\vec{\beta}) \coloneqq \mathbb{P}[\beta_i u_i = \max_{j \in N} \beta_j u_j]$$

= $\int_0^1 f_i(u) \prod_{j \in N \setminus \{i\}} F_j(\beta_i / \beta_j u) du.$ (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 $S = \{\vec{\beta} \in \mathbb{R}_{\geq 0}^n \mid \sum_{i \in N} \beta_i = 1\}$. 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 Appendix D.2, 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 $\mathbb{R}^n_{>0}$ (Appendix D.1).

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 $poly(n) n! \left(\log(2q) \left(1 + \frac{4nq}{\delta}\right)\right)^n$ time (Appendix D.2). 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 Appendix D.3):

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 (Lemma 10).

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 $2 q \epsilon$ (Lemma 11).

Bound on multipliers: If *i*'s multiplier is at least 2q times as large as *j*'s multiplier, then *i* must have a strictly larger resulting probability than *j* (Lemma 12).

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 $2q\epsilon$, and (b) the resulting probabilities of all $i \notin S$ weakly decrease, also by at most $2q\epsilon$ (Lemma 14).

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 $i \in N$, $1/n - \delta \leq p_i(\vec{\beta}) \leq 1/n + \delta$.

³Alternatively, 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].

⁴When 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 distributionspecific cost and accuracy of computing the integral in Eq. (1).

Input : An oracle to compute resulting probabilities, a constant $0 < \delta \le 1$, and a PDF upper bound qOutput: A vector of multipliers $\vec{\beta} \in \mathbb{R}^{n}_{>0}$ $\epsilon \leftarrow \delta/(2q)$; $\vec{z} \leftarrow \vec{0}$; 3 while $\exists i \in N$. $|p_{i}((1 + \epsilon)^{\vec{z}}) - 1/n| > \delta$ do 4 $| \begin{array}{c} S \leftarrow \{i \in N \mid p_{i}((1 + \epsilon)^{\vec{z}}) \le 1/n\};\\ \vec{z} \leftarrow \vec{z} + \mathbb{1}_{S};\\ \epsilon \text{ return } (1 + \epsilon)^{\vec{z}} \end{cases}$

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 $p_i((1 + \epsilon)^{\vec{z}})$ 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 $\delta = 2q \epsilon$ 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 2δ , 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 $T \coloneqq \left\lceil \frac{\log(2q)}{\log(1+\epsilon)} \right\rceil \cdot (n-1)$ 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 $p_i((1 + \epsilon)^{\vec{z}})$. 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 $z_j \ge \lfloor \frac{\log(2q)}{\log(1+\epsilon)} \rfloor$. Then, $\beta_j/\beta_i = (1+\epsilon)^{z_j-z_i} \ge 2 q$, 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 $p_i((1+\epsilon)^{\vec{z}})$. 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 $z_j \leq T - \lceil \frac{\log(2q)}{\log(1+\epsilon)} \rceil$. 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 Appendix D.4 that the equalizing multipliers are unique, using a strengthened local-monotonicity property. The algorithmic proof above yields an alternative proof of the existence of perfectly equalizing multipliers, by applying a limit argument similar to the one at the end of Theorem 4.

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 $i \neq j \in N$, $\mathbb{E}[u_i | \beta_i u_i = \max_{k \in N} \beta_k u_k] \geq C_{p,q} +$ $\mathbb{E}[u_i | \beta_i u_i = \max_{k \in N} \beta_k u_k]$. 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.

Before we go into the bounds, it is instructive to see why the interval support property is required for the multiplier approach. Consider a case with two agents: Agent A's utility is uniformly distributed on [1/4, 3/4], whereas agent B's distribution is uniform on $[0, 1/4] \cup [3/4, 1]$, which means that \mathcal{D}_B 's support is not an interval. It is easy to verify that setting both multipliers to 1 is equalizing. But $\mathbb{E}[u_A | u_A \ge u_B] =$ $\mathbb{E}[u_A]$, since the event $u_A \ge u_B$ only tells us that u_B is taken from the left interval in its support ([0, 1/4]) but u_A is still distributed uniformly in [1/4, 3/4]. Hence, without assuming interval support, the gap we aim to bound may be zero.

For a fixed set of distributions, interval support is enough to provide a positive gap (Appendix D.6). However, if we want to add new agents along the infinite sequence of instances as $m \to \infty$, we require a uniform lower bound on this gap. In Appendix D.7, we give examples showing that both very high probability densities and very low probability densities can make the gap arbitrarily small, which motivates our assumption of (p, q)-PDF-boundedness. In Appendix D.8, we derive the desired constant 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 $|p_i - 1/n| < 1/(2n)$

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



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.

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 Appendix D.9, 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 $\delta \coloneqq C_{p,q}/(4n)$, 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 $\mathcal{O}(mn)$ time). This yields the *approximate multiplier allocation*, which is fPO by construction. It remains to show EF: Proposition 6 applies to $\vec{\beta}$ since it satisfies the proposition's precondition: $|p_i - 1/n| \leq \delta \leq \frac{1}{4n} < \frac{1}{2n}$. By arithmetic, the positive gap guaranteed by the proposition implies that, for any two agents $i \neq j$, $\mathbb{E}[u_i(A_i)] - \mathbb{E}[u_i(A_j)] \geq \frac{m}{2n}C_{p,q}$. 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,

$$u_i(A_i) \ge \mathbb{E}[u_i(A_i)] - \frac{m}{4n}C_{p,q}$$

$$\ge \left(\mathbb{E}[u_i(A_j)] + \frac{m}{2n}C_{p,q}\right) - \frac{m}{4n}C_{p,q}$$

$$\ge \mathbb{E}[u_i(A_j)] + \frac{m}{4n}C_{p,q} \ge u_i(A_j),$$

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 = 2000 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 $m \ge 100$, but its allocations are essentially never PO unless m is very small (dash-dotted line). This lack of PO matches our theoretical predictions (Appendix E).

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 Appendix F.5 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 Appendix F, 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

⁶Appendix F 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 in Appendix F.4, 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.

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 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.

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.
- [Caragiannis et al., 2019] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum nash welfare. ACM Transactions on Economics and Computation, 7(3):1–32, October 2019.
- [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.
- [Papadimitriou, 1994] C. H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and system Sciences*, 48(3):498–532, 1994.
- [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.

Appendix

A Proof of Proposition 2: Existence of Envy-Free Allocations

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.

Proof. For any pair of agents $i, j \in N$, we will prove the probability of the event that i envies j in our asymmetric model is in $O(1/m^3)$. Consider a symmetric model with distribution \mathcal{D}_i for all agents. Let $\mathbb{P}^S[T]$ be the probability that some event T occurs in this symmetric model, and $\mathbb{P}^A[T]$ be the probability that some event T occurs in the asymmetric model.

For every utility profile, define its ordinal profile as $\{\mathcal{O}_k\}_{k\in N}$ where \mathcal{O}_k is a permutation of items M, in descending order according to $u_k(\alpha)$ for $\alpha \in M$. Let \mathbb{O} be the set of all possible ordinal profiles, which contains $(m!)^n$ elements. Since in both models, all agents are independent and the utility for all items are drawn independently, all ordinal profiles have the same probability to appear, that is,

$$\forall \tilde{O} \in \mathbb{O}, \mathbb{P}^{S}\left[\{\mathcal{O}_k\}_{k \in N} = \tilde{O}\right] = \mathbb{P}^{A}\left[\{\mathcal{O}_k\}_{k \in N} = \tilde{O}\right] = \frac{1}{(m!)^n}.$$

Since the allocation allocated by round robin algorithm is uniquely determined by the ordinal profile, let $\{A_k(O)\}_{k \in N}$ denote the resulting allocation given ordinal profile \tilde{O} . We can express $\mathbb{P}^S[i \text{ envise } j]$ as

$$\begin{split} \mathbb{P}^{S}\left[i \text{ envies } j\right] &= \sum_{\tilde{O} \in \mathbb{O}} \mathbb{P}^{S}\left[\{\mathcal{O}_{k}\}_{k \in N} = \tilde{O}\right] \cdot \mathbb{P}^{S}\left[i \text{ envies } j \left|\{\mathcal{O}_{k}\}_{k \in N} = \tilde{O}\right]\right] \\ &= \frac{1}{(m!)^{n}} \sum_{\tilde{O} \in \mathbb{O}} \mathbb{P}^{S}\left[i \text{ envies } j \left|\{\mathcal{O}_{k}\}_{k \in N} = \tilde{O}\right]\right]. \end{split}$$

Similarly, we can express $\mathbb{P}^A[i \text{ envies } j]$ as

$$\mathbb{P}^{A}\left[i \text{ envies } j\right] = \frac{1}{(m!)^{n}} \sum_{\tilde{O} \in \mathbb{O}} \mathbb{P}^{A}\left[i \text{ envies } j \left| \{\mathcal{O}_{k}\}_{k \in N} = \tilde{O}\right].$$

For all $\tilde{O} \in \mathbb{O}$,

$$\begin{split} \mathbb{P}^{S}\left[i \text{ envies } j \left| \{\mathcal{O}_{k}\}_{k \in N} = \tilde{O}\right] &= \mathbb{P}^{S}\left[\{u_{i}(\alpha) \sim \mathcal{D}_{i}\}_{\alpha \in M} : \sum_{\alpha \in A_{i}(\tilde{O})} u_{i}(\alpha) < \sum_{\alpha \in A_{j}(\tilde{O})} u_{i}(\alpha) \left| \mathcal{O}_{i} = \tilde{O}_{i}\right] \right. \\ &= \mathbb{P}^{A}\left[i \text{ envies } j \left| \{\mathcal{O}_{k}\}_{k \in N} = \tilde{O}\right]. \end{split}$$

Hence we have $\mathbb{P}^{A}[i \text{ envies } j] = \mathbb{P}^{S}[i \text{ envies } j]$. The following lemma is implied by the proof for Thm 3.1 by Manurangsi and Suksompong [2021].

Lemma 7 (Manurangsi and Suksompong 2021). In the symmetric model, if $m \in \Omega$ $(n \log n / \log \log n)$ and the common distribution \mathcal{D} is (p,q)-PDF-bounded on [0,1], then for any pair of agents i, i', the probability that i envies i' in the round robin allocation is at most $O(1/m^3)$.

PDF-bounded on [0, 1] Here we follow the assumptions of Manurangsi and Suksompong [2021] on the distributions: PDFbounded on [0, 1]. Since \mathcal{D}_i is (p, q)-PDF-bounded, the lemma indicates that $\mathbb{P}^S[i \text{ envies } j] = O(1/m^3)$. Thus, by our earlier arguments, the probability that agent *i* envies agent *j* in our asymmetric model is also in $O(1/m^3)$. Applying a union bound over all pairs *i*, *j*, we know that the allocation is envy-free in the asymmetric model with probability at least 1 - O(1/m) when $m \in \Omega(n \log n / \log \log n)$.

Interval support and PDF-bounded Moreover, we make slight modification (on constant level) to the proof by Manurangsi and Suksompong [2021] to generalize Lemma 7 and get bounded envy probability when only assuming interval support and (p, q)-PDF-bounded.

We first review the main idea of the proof for Thm 3.1 in their paper. For two agents i, i', let $X_t^{i,i'}$ denote *i*'s value for the item that *i'* gets in the *t*th round, and X_t^i denote *i*'s value for her own item in the *t*th round. While the maximum possible envy can be (when *i'* chooses before *i* in each round and gets 1 more item than *i*)

$$u_i(M_{i'}) - u_i(M_i) = X_1^{i,i'} - \sum_t \left(X_t^i - X_{t+1}^{i,i'} \right) \le 1 - \sum_{t=1}^I X_t^i \cdot \left(1 - Y_{t+1}^{i,i'} \right).$$
⁽²⁾

m



Figure 4: Illustrations of the example in the text. Dots represent utilities of 5 000 random items. Shaded regions delineate items allocated to agent A, for the maximum-percentile algorithm (left) and for the multiplier allocation (right). Three marked items (two have near-identical utilities) show that the maximum-percentile allocation is not PO.

where the last inequality follows from $Y_{t+1}^{i,i'} < 1$ for all $t \ge 1$, the gap in the first T rounds is sufficient for the envy to be negative with high probability.⁷ Manurangsi and Suksompong choose such T that with high probability $(1 - O(1/m^3))$, events (E1) $\sum_{t=1}^{T} Y_t^{i,i'} \ge T - 2$ and (E2) $X_t^i < 1/2$ for all $t \le T$, do not happen. Then it can be seen that when neither E1 nor E2 happen, the envy in Eq. (2) is non-positive.

For constant c > 1, consider changing E1 to E1': (E1') $\sum_{t=1}^{T} Y_t^{i,i'} \ge T - 2c$, then E1' and E2 give us negative envy of at least c - 1. We can still bound the probability that E1' or E2 occur in $O(1/m^3)$, by multiplying the value of T set in Manurangsi and Suksompong's proof by a factor of c, while keeping other valuations as they did. Then by Lemma 2.4 in their paper, the upper bound of the probability that E1' occurs is the upper bound for E1 to the power of c, which is still in $O(1/m^3)$. Meanwhile, the upper bound of the probability that E2 occurs is still in $O(1/m^3)$, for m that is sufficiently large. Hence we show that for some constant c > 1, when the distribution is PDF-bounded on [0, 1], the probability that i envies i' more than 1 - c in the round robin allocation is at most $O(1/m^3)$.

Now we use such result to further prove the envy probability is bounded by $O(1/m^3)$ when the distribution \mathcal{D} is PDFbounded and has interval support instead of [0, 1] support. The method is to use affine transformation to transform \mathcal{D} 's support interval [a, b] into [0, 1], mapping the original utility u to u' = (u - a)/(b - a) and the original distribution \mathcal{D} to \mathcal{D}' . The PDF $f_{\mathcal{D}}$ now becomes $f_{\mathcal{D}'}(u) = (b-a) \cdot f_{\mathcal{D}}((b-a)u+a)$. Since \mathcal{D} is (p, q)-PDF-bounded, the length of its support, b-a, must be at least 1/q. Thus for any $u \in [0, 1]$, $p/q \leq f_{\mathcal{D}'}(u) \leq q$, indicating that \mathcal{D}' is PDF-bounded on [0, 1]. Since the affine transform does not change the ordinal profile, the round robin allocation under the transformed utility profile, where all original utilities are transformed by the affine transformation: $u \mapsto (u - a)/(a - b)$, is the same as the one under the original utility profile. For the same round robin allocation, suppose the envy that i holds for i' in the transformed utility profile is e, then the envy in the original utility profile becomes e':

$$e' = \begin{cases} (b-a)e & \text{if } i \text{ and } i' \text{ get same number of items} \\ (b-a)e+a & \text{if } i' \text{ gets one more item than } i \\ (b-a)e-a & \text{if } i \text{ gets one more item than } i' \end{cases}$$

Then for *i* to envy *i'* in the original utility profile, i.e., e' > 0, it must be true that e > -a/(b-a). Since the utilities in the transformed utility profile can be considered as drawn randomly from \mathcal{D}' , which is PDF-bounded on [0, 1], by our previous result, the probability that e > -a/(b-a) = 1 - c is in $O(1/m^3)$. Hence the probability that *i* envies *i'* in the round robin allocation for distribution \mathcal{D} is still in $O(1/m^3)$, generalizing the result Lemma 7 to only assuming interval support and PDF-bounded for the distributions. Finally, similarly, we get $\mathbb{P}^S[i \text{ envies } j] = O(1/m^3)$ and that the allocation is envy-free in the asymmetric model with high probability, when distributions have interval support and are (p, q)-PDF-bounded.

B Discussion of Maximum-Percentile Algorithm

As we stated in the body of the paper, Kurokawa et al. [2016] apply the proof by Dickerson et al. [2014] to show the existence of EF allocations in the asymmetric setting. The core idea of their algorithm is to allocate each item to the agent for whom

⁷Note that we allow negative envy, whereas some works define envy to be $\max\{0, u_i(A_j) - u_i(A_i)\}$. When envy is -e, we can also say the negative envy is e.

the item is in the highest percentile of their utility distribution, which we will call the *maximum-percentile algorithm*. It is easy to see that each agent has a probability 1/n of receiving each item, and it is not too hard to show that agents have higher expected utility for items they receive than for items allocated to other agents. This implies envy-freeness with high probability as $m \in \Omega$ ($n \log n$) following the proof by Dickerson et al.

Unfortunately, this construction is unlikely to generate PO allocations: Consider a setting with two asymmetric agents, in which agent A's utility is drawn, with 50% probability, uniformly between 0 and 1/4, and, with 50% probability, uniformly between 1/4 and 1; and in which agent B's utility is drawn either uniformly between 0 and 3/4 or uniformly between 3/4 and 1, each with 50% probability. Conceptually, A's utility distribution skews towards lower values, whereas B's skews towards higher values. The black dots in Fig. 4 show random samples of these utilities, and the shaded region in the left plot marks the range of utilities in which the maximum-percentile algorithm allocates items to A. The left plot also highlights three specific items: two lie around the median utility for both agents and are given to A, and one of them lies around the top percentile for both agents and is given to B. The fact that the ratio " $u_B(\alpha)/u_A(\alpha)$ " is strictly greater for the two "median items" (at roughly (3/4)/(1/4) = 3) than for the "top item" (roughly 1) immediately implies that the allocation is not fPO: Agent B would profit from trading half of the top item against one of A's median items (roughly, since 3/4 > 1/2), and A would also profit from this trade (since 1/2 > 1/4). In fact, a similar trade of whole items, which exchanges both median items against the top item, shows that the maximum-percentile allocation violates Pareto-optimality proper.

The most promising way to avoid violations of PO is to construct fPO allocations, since the characterization of fPO using multipliers in Section 2 provides useful structure that is not available for PO. As shown in the right panel in Fig. 4, this corresponds to choosing a line through the origin, allocating items below the line to A, and allocating items above the line to B. In fact, the plot shows the unique such line with the added property that a random-utility item is equally likely to be given to either agent. In the next section, we generalize this kind of allocation to arbitrary numbers of agents.

C Discussion of Normalizing Multipliers

As discussed in previous section, the most promising way to construct PO allocation is to utilize multiplier-based maximumwelfare allocation. One natural choice is the set of normalizing multipliers that normalize each agent's values to add up to 1. However, we show by the following counter-example that the set of normalizing multipliers may violate EF.

Consider a setting with |N| = n agents where agent 1's utility is drawn uniformly from [0, 1] while the rest of the agents' utilities are drawn uniformly from [1/3, 2/3]. When the number of items is sufficiently large, the normalizing multipliers for all agents will be close. With dominating probability, by Chernoff bound, all multipliers for agents $N/\{1\}$ are less than 5/4 of agent 1's multiplier. A Chernoff bound will also guarantee high probability that agent 1 receives more than $1/12 = 1/2(1-2/3 \cdot 5/4)$ of all items, while there must be some agent *i* receiving less than 1/n of all items. Agent *i* will surely envy agent 1 when $1/n \cdot 2/3 < 1/12 \cdot 1/3$.

D Proofs Used in Existence Result of Envy-free and Pareto-optimal Allocations

D.1 Continuity of Probability Function $p_i(\cdot)$

Lemma 8. For all $i \in N$, the probability function $p_i(\cdot)$ defined by

$$p_i(\beta_1,\ldots,\beta_n) = \mathbb{P}\left[\beta_i u_i = \max_{j\in N} \beta_j u_j\right]$$

is continuous on $(\beta_1, \ldots, \beta_n) \in \mathbb{R}^n_{>0}$.

Proof. We have

$$p_i(\beta_1, \dots, \beta_n) = \int_0^1 f_i(u) \prod_{j \neq i} F_j\left(\frac{\beta_i}{\beta_j}u\right) du_j$$

then

$$\begin{aligned} |p_i(\beta'_1,\cdots,\beta'_n) - p_i(\beta_1,\cdots,\beta_n)| &\leq \int_0^1 f_i(u) \cdot \left| \prod_{j\neq i} F_j\left(\frac{\beta'_i}{\beta'_j}u\right) - \prod_{j\neq i} F_j\left(\frac{\beta_i}{\beta_j}u\right) \right| \, du \\ &\leq \sum_{j\neq i} \int_0^1 f_i(u) \cdot \left| F_j\left(\frac{\beta'_i}{\beta'_j}u\right) - F_j\left(\frac{\beta_i}{\beta_j}u\right) \right| \, du. \end{aligned}$$

For nonatomic distribution \mathcal{D}_j , its cumulative distribution function $F_j(\cdot)$ is continuous, also the function $g_u(x,y) = \frac{xu}{y}$ is continuous when $x, y \neq 0$. Thus

$$\lim_{(\beta'_i,\beta'_j)\to(\beta_i,\beta_j)}F_j\left(\frac{\beta'_i}{\beta'_j}u\right)=F_j\left(\frac{\beta_i}{\beta_j}u\right),$$

and we have

$$\lim_{(\beta'_1,\cdots,\beta'_n)\to(\beta_1,\cdots,\beta_n)} |p_i(\beta'_1,\cdots,\beta'_n) - p_i(\beta_1,\cdots,\beta_n)|$$

$$\leq \sum_{j\neq i} \int_0^1 f_i(u) \lim_{(\beta'_i,\beta'_j)\to(\beta_i,\beta_j)} \left| F_j\left(\frac{\beta'_i}{\beta'_j}u\right) - F_j\left(\frac{\beta_i}{\beta_j}u\right) \right| du = 0.$$

Therefore function $p_i(\cdot)$ is continuous on $\mathbb{R}^n_{>0}$.

D.2 Proof for Existence of Multipliers with Sperner's Lemma

Here we show a detailed proof for existence of equalizing multipliers with Sperner's Lemma. Without loss of generality we assume all multipliers add up to 1, then $(\beta_1, \beta_2, \ldots, \beta_n)$ falls in a (n-1)-dimensional simplex S. Now we define the coloring function $f: S \to N$, which maps each set of multipliers to an agent with the highest probability of having the largest scaled utility under the set of multipliers:

$$f(\beta_1, \dots, \beta_n) \in \underset{i \in N}{\operatorname{arg\,max}} \mathbb{P}\left[\beta_i u_i = \underset{j \in N}{\max} \beta_j u_j\right].$$

It is clear that the *n* vertices of the simplex are colored with *n* different "colors", since agent *i* will have probability 1 of having the largest scaled utility when $\beta_i = 1$ and $\beta_j = 0, \forall j \neq i$.

Now we divide the simplex S into smaller simplices with (at most) half the diameter of previous simplex. Let S_1 denote the set of smaller simplices. Then we further divide each simplex in S_1 into even smaller simplices with half diameter, and we denote them by S_2 . We repeat the procedure and divide the original simplex into S_1, S_2, S_3, \ldots containing smaller and smaller simplices.

By Sperner's Lemma, in each S_i $(i \ge 1)$, there are always an odd number of simplicies that are colored with n colors, indicating the existence of a simplex $s_i \in S_i$ whose vertices are colored with all n colors. Now consider the sequence of such simplices: s_1, s_2, s_3, \ldots , and let the *i*th vertex be the vertex that is mapped to *i* by f. We can represent each simplex with a n^2 -dimensional vector:

$$(\boldsymbol{\beta}^1, \boldsymbol{\beta}^2, \dots, \boldsymbol{\beta}^n) = (\beta_1^1, \dots, \beta_n^1; \dots; \beta_1^n, \dots, \beta_n^n) \in \mathbb{R}^{n^2}.$$

where β^i is the vector of multipliers in the *i*th vertex and β_j^i is the *j*th multiplier in the *i*th vertex. Let s_j^i denote the subvector β^i in s_j .

Since all of them are bounded between (0, 0, ..., 0) and (1, 1, ..., 1), by *Bolzano-Weierstrass Theorem*, there must be a convergent subsequence $s_{a_1}, s_{a_2}, s_{a_3}, ...$ that converges to a simplex $(\beta^{1^*}, \beta^{2^*}, ..., \beta^{n^*})$ in the space:

$$\lim_{t\to\infty} s_{a_t} = (\beta^{1^*}, \beta^{2^*}, \dots, \beta^{n^*}).$$

Since the simplicies in $\{s_i\}_{i>0}$ get smaller and smaller with a ratio of 1/2, we have

$$\lim_{t \to \infty} s^1_{a_t} = \lim_{t \to \infty} s^2_{a_t} = \dots = \lim_{t \to \infty} s^n_{a_t}.$$

Then it can be deduced that $\beta^{1^*} = \cdots = \beta^{n^*} = \beta^*$.

We argue that $\beta^* \in \mathbb{R}^n_{>0}$. Since otherwise, if $\beta_i^* = 0$ for some $i \in N$, then β_i^i , the *i*th multiplier in the *i*th vertex grows arbitrarily small as the sequence $\{s_{a_t}\}$ goes. This would lead to the probability that agent *i* having the largest scaled utility goes to 0, contradicting the fact that agent *i* should have the highest probability of having largest scaled utility, with probability at least 1/n.

Now we claim that all agents have equal probability of having the largest scaled utility under β^* . Define

$$p_i(\beta_1,\ldots,\beta_n) = \mathbb{P}\left[\beta_i u_i = \max_{j\in N} \beta_j u_j\right].$$

As shown in Appendix D.1, $p_i(\cdot)$ is a continuous function in $\mathbb{R}^n_{>0}$. Then for all $i \in N$,

$$\frac{1}{n} \leq \lim_{t \to \infty} p_i(\boldsymbol{s_{a_t}}^i) = p_i(\lim_{t \to \infty} \boldsymbol{s_{a_t}}^i) = p_i(\boldsymbol{\beta^*}).$$

Since $\sum_{i=1}^{n} p_i(\boldsymbol{\beta}^*) = 1$, we must have

$$p_1(\boldsymbol{\beta^*}) = \cdots = p_n(\boldsymbol{\beta^*}) = \frac{1}{n}.$$

This shows that β^* is the set of multipliers we are looking for, hence the existence.

Algorithm Based on Sperner's-Lemma Proof

We discuss here the degree to which the Sperner's argument can be implemented as an approximation algorithm.

Proposition 9. For any $\delta > 0$, a variant of the previous Sperner's argument computes a set of multipliers $\{\beta_i\}_{i \in N}$ such that $|p_i - p_j| \leq \delta$ in time $poly(n) n! \left(\log(2q) \left(1 + \frac{4 n q}{\delta}\right)\right)^n$.

Proof. To make the existence proof algorithmic, a key question is how to discretize the space of multipliers into simplices. Such a simplicization should be easy to traverse algorithmically and its simplices should describe sufficiently compact sets of multipliers such that any panchromatic simplex should allow to uniformly bound how far the probabilities are from 1/n. As Papadimitriou [1994] describes, there is already no obvious simplicization when n = 4; for example, the regular tetrahedron cannot be partitioned into multiple regular tetrahedra. Following Papadimitriou's discussion, we apply Sperner's lemma to a hypercube rather than a simplex:

Let $\epsilon > 0$ denote a small constant, to be determined later. Let our set of points be the grid $G := \{p \in \mathbb{Z}^{n-1} \mid ||p||_{\infty} \leq \lfloor \frac{\log(2q)}{\log(1+\epsilon)} \rfloor\}$. We can partition this grid first in cubelets of the shape $[0,1]^{n-1} + p \subseteq G$ for some $p \in G$, and subdivide each of these cubelets into (n-1)! simplices.⁸ We map each point $p = (p_1, \ldots, p_{n-1})$ in the grid G to a color in [n]; specifically, we choose the color $\operatorname{argmax}_{i \in N} p_i((1 + \epsilon)^{p_1}, (1 + \epsilon)^{p_2}, \ldots, (1 + \epsilon)^{p_{n-1}}, 1)$ with some canonical tie breaking. If, for some $p \in G$ and $1 \leq i \leq n-1$, we have $p_i = -\lceil \frac{\log(2q)}{\log(1+\epsilon)} \rceil$, then *i*'s multiplier is at most $\frac{1}{2q}$ whereas agent *n*'s multiplier is 1; by Lemma 12, *n*'s probability of being the largest is strictly larger than *i*'s, which means that *p*'s color is not *i*. Similarly, if $p_i = \lceil \frac{\log(2q)}{\log(1+\epsilon)} \rceil$, then *i*'s probability is strictly larger than *n*'s by Lemma 12, and therefore *p*'s color is not *n*. Given these observations, there must exist a panchromatic simplex, which in turn must be contained in a cubelet $C = [0, 1]^{n-1} + p$ for a $p \in G$. Fix the multipliers $\{\beta_i\}_{i \in N}$ by setting $\beta_n \coloneqq 1$, and $\beta_i \coloneqq (1+\epsilon)^{p_i}$ for all $1 \leq i \leq n-1$. Recall that some point in *C* is colored *n*, which means that *n* has the largest probability for the corresponding multipliers, and in particular has a probability of at least 1/n. Since *n*'s multiplier β_n is equal to the multiplier of this point, and since all other multipliers β_i are at most the corresponding multiplier at this point, it follows that $p_n(\beta_1, \ldots, \beta_n) \geq 1/n$. Similarly, fix some agent $1 \leq i \leq n-1$. Since some point in *C* has color *i*, the point $p + e_i$ must also have color *i*, where e_i is the unit vector in dimension *i*. It follows that $p_i(\beta_1, \ldots, \beta_n) \geq 1/n - 2q \epsilon$. These lower bounds also allow us to upper bound the probabilities; for each $i \in N$, $p_i(\beta_1, \ldots, \beta_n) \leq 1 - (n-1)(1/n - 2q \epsilon) = 1/n + 2(n-1)q \epsilon$. If we set $\epsilon = \frac{\delta}{2nq}$, then it holds for all i

$$|p_i(\beta_1,\ldots,\beta_n)-p_j(\beta_1,\ldots,\beta_n)|\leq \delta.$$

It remains to bound the running time of this algorithm. To find the panchromatic simplex, we traverse the simplices, visiting each at most once. Within each simplex, we query the oracle to determine the color of the new vertex and perform polynomialtime work to move on to the next simplex. The bottleneck of the computation is that we might have to traverse nearly all vertices, which leads to an overall running time in

$$poly(n) (n-1)! \left(\frac{2\log(2\,q)+1}{\log(1+\frac{\delta}{2\,n\,q})}\right)^{n-1} \le poly(n) \, n! \left((2\log(2\,q)+1) \left(1/2+\frac{2\,n\,q}{\delta}\right)\right)^n.$$

D.3 Proof of Properties of the p_i

Lemma 10. Fix multiplier sets $\vec{\beta}, \vec{\beta}' \in \mathbb{R}^n_{>0}$ and an agent *i*. If $\frac{\beta'_i}{\beta_i} \ge \frac{\beta'_j}{\beta_j}$ for all $j \neq i$, then $p_i(\vec{\beta}') \ge p_i(\vec{\beta})$.

Proof. In Eq. (1), observe that the F_j are monotone increasing and that $\beta'_i/\beta'_j = (\beta'_i/\beta_i)/(\beta'_j/\beta_j) \cdot \beta_i/\beta_j \ge \beta_i/\beta_j, \forall j$. \Box

Lemma 11. Fix a set of multipliers $\vec{\beta}$, an agent *i*, and $\epsilon > 0$. Let $\vec{\beta}'$ denote a set of multipliers such that $\beta'_i \coloneqq (1 + \epsilon) \beta_i$ and $\beta'_j \coloneqq \beta_j$ for all $j \neq i$. Then, $p_i(\vec{\beta}') \leq p_i(\vec{\beta}) + 2q\epsilon$.

Proof. For convenience, set $p_i \coloneqq p_i(\beta_1, \ldots, \beta_n)$ and $p'_i \coloneqq p_i(\beta_1, \ldots, \beta_{i-1}, (1+\epsilon)\beta_i, \beta_{i+1}, \ldots, \beta_n)$. Moreover, define a function $\pi : \mathbb{R} \to \mathbb{R}$ such that $\pi(t) \coloneqq \prod_{j \neq i} F_j(\frac{\beta_i}{\beta_i}t)$. Observe that π is monotone increasing and its range lies within [0, 1].

⁸For an excellent exposition, see https://people.csail.mit.edu/costis/6896sp10/lec6.pdf, accessed on January 11, 2022.

Using these definitions,

$$p_{i}' - p_{i} = \int_{0}^{1} f_{i}(u) \left(\pi((1+\epsilon)u) - \pi(u) \right) du$$

$$= \sum_{t=0}^{\infty} \int_{(1+\epsilon)^{-t-1}}^{(1+\epsilon)^{-t}} f_{i}(u) \left(\pi((1+\epsilon)u) - \pi(u) \right) du$$

$$\leq \sum_{t=0}^{\infty} \left(\left(\pi((1+\epsilon)^{-t+1}) - \pi((1+\epsilon)^{-t-1}) \right) \int_{(1+\epsilon)^{-t-1}}^{(1+\epsilon)^{-t}} f_{i}(u) du \right)$$

$$\leq \sum_{t=0}^{\infty} \left(\left(\pi((1+\epsilon)^{-t+1}) - \pi((1+\epsilon)^{-t-1}) \right) \int_{(1+\epsilon)^{-t-1}}^{(1+\epsilon)^{-t}} q \, du \right)$$

$$= q \, \epsilon \, \sum_{t=0}^{\infty} \left(\pi((1+\epsilon)^{-t+1}) - \pi((1+\epsilon)^{-t-1}) \right) (1+\epsilon)^{-t-1} \tag{3}$$

By the monotonicity of π , all coefficients $\pi((1 + \epsilon)^{-t+1}) - \pi((1 + \epsilon)^{-t-1})$ are nonnegative. Moreover, if we add up these coefficients only for the even *t*, this is a telescoping series

$$\sum_{t=0,2,4,\dots} \pi((1+\epsilon)^{-t+1}) - \pi((1+\epsilon)^{-t-1}) = \lim_{t \to \infty} \pi(1+\epsilon) - \pi((1+\epsilon)^{-2t-1}) \le 1.$$

Thus, the even summands of Eq. (3) are a "subconvex" combination of the $(1 + \epsilon)^{-t-1}$, and are therefore upper bounded by the largest such term:

$$\sum_{t=0,2,4,\dots} \left(\pi((1+\epsilon)^{-t+1}) - \pi((1+\epsilon)^{-t-1}) \right) (1+\epsilon)^{-t-1} \le (1+\epsilon)^{-1}.$$

Applying the same reasoning to the odd summands, we obtain

$$\sum_{t=1,3,5,\dots} \left(\pi((1+\epsilon)^{-t+1}) - \pi((1+\epsilon)^{-t-1}) \right) (1+\epsilon)^{-t-1} \le (1+\epsilon)^{-2}$$

By plugging these last two equations into Eq. (3), we conclude that

$$p'_i - p_i \le q \,\epsilon \left((1+\epsilon)^{-1} + (1+\epsilon)^{-2} \right) \le 2 \, q \,\epsilon.$$

Lemma 12. For two agents *i* and *j* and $\vec{\beta} \in \mathbb{R}^n_{>0}$, if $\frac{\beta_i}{\beta_j} \ge 2q$, then $p_i(\vec{\beta}) > p_j(\vec{\beta})$.

Proof. Suppose there is a pair of $i, j \in N$ such that $\beta_i/\beta_j > 2q$. Then we we have the following inequalities for the probability p_i, p_j of i and j getting each item:

$$p_{i} = \mathbb{P}\left[\beta_{i}u_{i} = \max_{k \in N} \beta_{k}u_{k}\right] \geq \mathbb{P}\left[\beta_{i}u_{i} = \max_{k \in N} \beta_{k}u_{k} \land \beta_{i}u_{i} \geq \beta_{j}\right]$$
$$= \mathbb{P}\left[\beta_{i}u_{i} \geq \beta_{j}\right] \cdot \mathbb{P}\left[\beta_{i}u_{i} = \max_{k \in N} \beta_{k}u_{k} \middle| \beta_{i}u_{i} \geq \beta_{j}\right] \geq \mathbb{P}\left[\beta_{i}u_{i} \geq \beta_{j}\right] \cdot \mathbb{P}\left[\beta_{j} \geq \max_{k \in N, \ k \neq i, j} \beta_{k}u_{k}\right],$$

and

$$\begin{split} p_j &= \mathbb{P}\left[\beta_j u_j = \max_{k \in N} \beta_k u_k\right] = \mathbb{P}\left[\beta_j u_j \geq \beta_i u_i\right] \cdot \mathbb{P}\left[\beta_j u_j = \max_{k \in N} \beta_k u_k \ \bigg| \ \beta_j u_j \geq \beta_i u_i\right] \\ &< \mathbb{P}\left[\beta_j \geq \beta_i u_i\right] \cdot \mathbb{P}\left[\beta_j \geq \max_{k \in N, \ k \neq i, j} \beta_k u_k\right], \end{split}$$

where the last inequality follows from $u_j \leq 1$. Since

$$\mathbb{P}\left[\beta_i u_i \le \beta_j\right] = \mathbb{P}\left[u_i \le \beta_j / \beta_i\right] \le \mathbb{P}\left[u_i \le 1 / (2q)\right] \le 1/2,$$

we have

$$\mathbb{P}\left[\beta_i u_i \ge \beta_j\right] \ge \mathbb{P}\left[\beta_j \ge \beta_i u_i\right] \quad \Rightarrow \quad p_i > p_j$$

Here we extend Lemma 12 to bound the ratios under approximated equalizing multipliers.

Corollary 13. If the set of multipliers $\vec{\beta}$ satisfies that $|p_i - 1/n| < 1/(2n)$ for all $i \in N$, then the ratio β_i/β_j between any pair of agents i, j is at most 4q.

Proof. If $\beta_i/\beta_j > 3q$ for some $i, j \in N$, then

$$\mathbb{P}\left[\beta_i u_i \le \beta_j\right] = \mathbb{P}\left[u_i \le \beta_j / \beta_i\right] \le \mathbb{P}\left[u_i < 1/(4q)\right] < 1/4,$$

Then following the inequalities in the above proof for Lemma 12, we would have $p_i > 3 p_j$, which contradicts to the fact that $1/(2n) < p_i, p_j < 3/(2n)$.

Lemma 14. Let $\vec{z} \in \mathbb{Z}^n$, $\epsilon > 0$, and $S \subseteq N$. For all $i \in S$,

$$p_i((1+\epsilon)^{\vec{z}}) \le p_i((1+\epsilon)^{\vec{z}+\mathbb{1}_S}) \le p_i((1+\epsilon)^{\vec{z}}) + 2q\epsilon,$$

and, for all $i \notin S$,

$$p_i\left((1+\epsilon)^{\vec{z}}\right) - 2\,q\,\epsilon \le p_i\left((1+\epsilon)^{\vec{z}+\mathbb{1}_S}\right) \le p_i\left((1+\epsilon)^{\vec{z}}\right).$$

Proof. This follows from repeated application of Lemma 10 and Lemma 11, as well as the observation that the effect on the resulting multipliers is equal whether we multiply all β_i for $i \in S$ by $1 + \epsilon$ or whether we divide all β_i for $i \notin S$ by $1 + \epsilon$ instead.

D.4 Uniqueness of Equalizing Multipliers

First, we formalize the notion of local strict monotonicity as the following lemma.

Lemma 15. Assuming the distributions for all agents have interval support. For any agent $j \in N$, let $p_j > 0$, p'_j be the probability that agent j receives each item under $\{\beta_i\}_{i\in N}$ and $\{\beta'_i\}_{i\in N}$. If $\forall i \in N$, $\beta'_j/\beta'_i \ge \beta_j/\beta_i$ and there exists some $k \in N$ such that $p_k > 0$ and $\beta'_j/\beta'_k > \beta_j/\beta_k$, then $p'_j > p_j$.

Proof. First we can express p'_i, p_j as follows:

$$p'_{j} = \mathbb{P}\left[\beta'_{j}u'_{j} = \max_{i \in N} \{\beta'_{i}u'_{i}\}\right] = \int_{0}^{1} f_{j}(u) \prod_{i \in N/\{j\}} F_{i}\left(\frac{\beta'_{j}}{\beta'_{i}}u\right) du,$$
$$p_{j} = \int_{0}^{1} f_{j}(u) \prod_{i \in N/\{j\}} F_{i}\left(\frac{\beta_{j}}{\beta_{i}}u\right) du.$$

Suppose \mathcal{D}_j 's support interval is $[u_j, \overline{u_j}]$. Now take

$$\underline{u} = \max_{u} \{ u : \exists i \in N, F_i\left(\frac{\beta_j}{\beta_i}u\right) = 0 \}.$$

From $F_j\left(\frac{\beta_j}{\beta_j}\underline{u_j}\right) = 0$, then by definition of \underline{u} , it is true that $\underline{u_j} \leq \underline{u}$. We also have $\underline{u} < \overline{u_j}$, since otherwise we can find agent $i \in N$, such that $F_i\left(\frac{\beta_j}{\beta_i}u\right) = 0$ for all possible value of $u_j = u \in [\underline{u_j}, \overline{u_j}]$, making $p_j = 0$. Then we take

$$\overline{u} = \min_{u} \{ u : F_k \left(\frac{\beta_j}{\beta_k} u \right) = 1 \}$$

We argue that $\overline{u} > \underline{u}$. Consider otherwise, then we can find $u_0 \in [\overline{u}, \underline{u}]$ and $i \in N$, where

$$F_i\left(\frac{\beta_j}{\beta_i}u_0\right) = 0, \ F_k\left(\frac{\beta_j}{\beta_k}u_0\right) = 1$$
$$\Rightarrow u_i \ge \frac{\beta_j}{\beta_i}u_0, \ u_k \le \frac{\beta_j}{\beta_k}u_0 \Rightarrow \beta_i u_i \ge \beta_k u_k,$$

which will make $p_k = 0$ since all distributions are non-atomic, contradicting our assumption that $p_k > 0$. Combining the earlier arguments, we know that $\max\{u_j, \underline{u}\} < \min\{\overline{u_j}, \overline{u}\}$. Then for any $u \in (\max\{u_j, \underline{u}\}, \min\{\overline{u_j}, \overline{u}\}) = \mathcal{I}$,

$$f_{j}(u) > 0, \quad F_{i}\left(\frac{\beta_{j}}{\beta_{i}}u\right) > 0, \forall i \in N,$$

$$0 < F_{k}\left(\frac{\beta_{j}}{\beta_{k}}u\right) < 1 \Rightarrow F_{k}\left(\frac{\beta_{j}'}{\beta_{k}'}u\right) > F_{k}\left(\frac{\beta_{j}}{\beta_{k}}u\right)$$

The last inequality follows from the fact that the derivative $F'_k(u) > 0$ for u that satisfies $0 < F_k(u) < 1$ (guaranteed by Interval support property, $0 < F_k(u) < 1$ just means u is in the support interval), and that $\beta'_j/\beta'_k > \beta_j/\beta_k$. Then

$$\int_{u \in \mathcal{I}} f_j(u) \prod_{i \in N/\{j\}} F_i\left(\frac{\beta'_j}{\beta'_i}u\right) du > \int_{u \in \mathcal{I}} f_j(u) \prod_{i \in N/\{j\}} F_i\left(\frac{\beta_j}{\beta_i}u\right) du$$

For *u* in the rest of the range, we have that

$$\frac{\beta'_j}{\beta'_i} \ge \frac{\beta_j}{\beta_i} \Rightarrow F_i\left(\frac{\beta'_j}{\beta'_i}u\right) \ge F_i\left(\frac{\beta_j}{\beta_i}u\right), \, \forall i \in N/\{j\}.$$

Then the integral in this range for p'_j is greater or equal to that for p_j . Hence we have showed the strict ordering $p'_j > p_j$. \Box

Now we prove uniqueness of equalizing multipliers. Suppose there are two different sets of equalizing multipliers, namely $\vec{\beta}, \vec{\beta'}$ (assume they are both normalized by setting $\beta_1 = 1$). We can find the $i = \arg \max_{j \in N} \beta'_j / \beta_j$, w.l.o.g. assume $\beta'_i / \beta_i > 1$, otherwise we simply exchange $\vec{\beta}$ and $\vec{\beta'}$ and the maximum ratio must be larger than 1 since the ratios cannot all be 1. Now we know that for any $j \in N$, $\beta'_i / \beta'_j = (\beta'_i / \beta_i) / (\beta'_j / \beta_j) \cdot \beta_i / \beta_j \ge \beta_i / \beta_j$, while $\beta'_i / \beta'_1 \ge \beta_i / \beta_1$. Then by the local strict monotonicity Lemma 15, *i*'s probability under $\vec{\beta'}$ is strictly larger than that under $\vec{\beta}$, contradicting to them both being 1/n under equalizing multipliers.

D.5 Inequalities for Expectations

Lemma 16. For any pair of agents $i, j \in N$, the following inequalities between expectations hold (assuming that the conditions in the conditional expectations can be met):

$$\mathbb{E}\left[u_i \left| \beta_j u_j = \max_{k \in N} \beta_k u_k \right] \le \mathbb{E}\left[u_i\right]$$
(4)

and

$$\mathbb{E}\left[u_i \middle| \beta_i u_i = \max_{k \in N} \beta_k u_k\right] \ge \mathbb{E}\left[u_i \middle| \beta_i u_i \ge \beta_j u_j\right].$$
(5)

Proof. Let $v_j = \mathbb{1} [\beta_j u_j = \max_{k \in N} \beta_k u_k]$, which is a random variable taking value from $\{0, 1\}$. Then we have

$$\mathbb{E}\left[u_i \left| \beta_j u_j = \max_{k \in N} \beta_k u_k \right] = \mathbb{E}\left[u_i \left| v_j = 1\right] = \int_0^1 f_{u_j \mid v_j = 1}(u) \cdot \mathbb{E}\left[u_i \left| v_j = 1, u_j = u\right]\right] du.$$

Given $v_j = 1$, for all $u \in [0, 1]$, it holds that

$$\mathbb{E}\left[u_{i} \mid v_{j} = 1, u_{j} = u\right] = \mathbb{E}\left[u_{i} \mid \beta_{i}u_{i} \leq \beta_{j}u\right] \leq \mathbb{E}\left[u_{i}\right]$$

Thus

$$\int_{0}^{1} f_{u_{j}|v_{j}=1}(u) \cdot \mathbb{E}\left[u_{i} \mid v_{j}=1, u_{j}=u\right] du \leq \int_{0}^{1} f_{u_{j}|v_{j}=1}(u) \cdot \mathbb{E}\left[u_{i}\right] du = \mathbb{E}\left[u_{i}\right].$$

Therefore Eq. (4) holds.

For Eq. (5), let $v_{i,j} = \max_{k \in N/\{i,j\}} \beta_k u_k$ be a random variable taking value in [0, 1], and $v_i = \mathbb{1} [\beta_i u_i = \max_{k \in N} \beta_k u_k]$. Then by substituting the condition we have

$$\mathbb{E}\left[u_{i} \middle| \beta_{i}u_{i} = \max_{k \in N} \beta_{k}u_{k}\right] = \mathbb{E}\left[u_{i} \middle| v_{i} = 1\right] = \mathbb{E}\left[u_{i} \middle| \beta_{i}u_{i} \ge \beta_{j}u_{j}, \beta_{i}u_{i} \ge v_{i,j}\right]$$
$$= \int_{0}^{1} f_{v_{i,j}|v_{i}=1}(u) \cdot \mathbb{E}\left[u_{i} \middle| \beta_{i}u_{i} \ge \beta_{j}u_{j}, \beta_{i}u_{i} \ge v_{i,j}, v_{i,j} = u\right] du.$$

Since $v_{i,j}, u_i, u_j$ are independent, for all $u \in [0, 1]$ it holds that

$$\mathbb{E}\left[u_i \middle| \beta_i u_i \ge \beta_j u_j, \beta_i u_i \ge v_{i,j}, v_{i,j} = u\right] = \mathbb{E}\left[u_i \middle| \beta_i u_i \ge \beta_j u_j, \beta_i u_i \ge u\right] \ge \mathbb{E}\left[u_i \middle| \beta_i u_i \ge \beta_j u_j\right].$$

Hence

$$\int_{0}^{1} f_{v_{i,j}|v_{i}=1}(u) \cdot \mathbb{E}\left[u_{i} \middle| \beta_{i}u_{i} \ge \beta_{j}u_{j}, \beta_{i}u_{i} \ge v_{i,j}, v_{i,j} = u\right] du$$
$$\ge \int_{0}^{1} f_{v_{i,j}|v_{i}=1}(u) \cdot \mathbb{E}\left[u_{i} \middle| \beta_{i}u_{i} \ge \beta_{j}u_{j}\right] du = \mathbb{E}\left[u_{i} \middle| \beta_{i}u_{i} \ge \beta_{j}u_{j}\right]$$

Therefore Eq. (5) holds.

.

D.6 Positive Gap Between Conditional Expectations

Proposition 17. Fix a set of agents whose utility distributions have interval support, and let $\{\beta_i\}_{i \in N}$ denote a set of multipliers such that $|p_i - 1/n| < 1/(2n)$ for all $i \in N$. Then, for all $i \neq j$,

$$\mathbb{E}\left[u_i \middle| \beta_i u_i = \max_{k \in N} \beta_k u_k\right] > \mathbb{E}\left[u_i \middle| \beta_j u_j = \max_{k \in N} \beta_k u_k\right].$$

Proof. Let $X \coloneqq \beta_i u_i$ and $Y \coloneqq \beta_j u_j$ denote the scaled random variables, then it still holds that: (a) X and Y are independent, (b) X and Y have interval support, and (c) both $\mathbb{P}[X > Y]$ and $\mathbb{P}[Y > X]$ are at least 1/(2n) > 0. From Appendix D.5, we know that

$$\mathbb{E}\left[u_i \middle| \beta_i u_i = \max_{k \in N} \beta_k u_k\right] \ge \mathbb{E}\left[u_i \middle| \beta_i u_i \ge \beta_j u_j\right] = \frac{1}{\beta_i} \mathbb{E}\left[X \middle| X > Y\right],$$
$$\mathbb{E}\left[u_i \middle| \beta_j u_j = \max_{k \in N} \beta_k u_k\right] \le \mathbb{E}\left[u_i\right] = \frac{1}{\beta_i} \mathbb{E}\left[X\right].$$

Then it suffices to show that $\mathbb{E}[X \mid X > Y] > \mathbb{E}[X]$.

Let \mathcal{I} denote the intersection of the support intervals of X and Y, excluding both endpoints. This intersection is a nonempty interval and both $\mathbb{P}[X \in \mathcal{I}]$ and $\mathbb{P}[Y \in \mathcal{I}]$ must have positive probability, since, else, one variable's support would entirely lie below or above the other variable's support, which would contradict the above observation that $0 < \mathbb{P}[X > Y] < 1$. Since for all $y \in \mathcal{I}$,

$$\mathbb{E}\left[X\right] = \underbrace{\mathbb{P}\left[X > y\right]}_{>0} \cdot \underbrace{\mathbb{E}\left[X \mid X > y\right]}_{>y} + \underbrace{\mathbb{P}\left[X \le y\right]}_{=1 - \mathbb{P}\left[X > y\right] > 0} \cdot \underbrace{\mathbb{E}\left[X \mid X \le y\right]}_{\le y} \Rightarrow \mathbb{E}\left[X \mid X > y\right] > \mathbb{E}\left[X\right].$$

Hence

$$\begin{split} \mathbb{E}\left[X \mid X > Y, Y \in \mathcal{I}\right] &= \int_{0}^{\beta_{j}} f_{Y|Y < X, Y \in \mathcal{I}}(y) \cdot \mathbb{E}\left[X \mid X > Y, Y \in \mathcal{I}, Y = y\right] dy \\ &= \int_{0}^{\beta_{j}} f_{Y|Y < X, Y \in \mathcal{I}}(y) \cdot \mathbb{E}\left[X \mid X > y\right] dy \\ &> \int_{0}^{\beta_{j}} f_{Y|Y < X, Y \in \mathcal{I}}(y) \cdot \mathbb{E}\left[X\right] dy = \mathbb{E}\left[X\right]. \end{split}$$

For all $y \notin \mathcal{I}$, $\mathbb{E}[X | X > y] \ge \mathbb{E}[X]$, from which we have

$$\mathbb{E}\left[X \mid X > Y, Y \notin \mathcal{I}\right] = \int_{0}^{\beta_{j}} f_{Y|Y < X, Y \notin \mathcal{I}}(y) \cdot \mathbb{E}\left[X \mid X > Y, Y \notin \mathcal{I}, Y = y\right] dy$$
$$= \int_{0}^{\beta_{j}} f_{Y|Y < X, Y \notin \mathcal{I}}(y) \cdot \mathbb{E}\left[X \mid X > y\right] dy$$
$$\geq \int_{0}^{\beta_{j}} f_{Y|Y < X, Y \notin \mathcal{I}}(y) \cdot \mathbb{E}\left[X\right] dy = \mathbb{E}\left[X\right].$$

Then, we can bound

$$\mathbb{E}\left[X \mid X > Y\right] = \underbrace{\mathbb{P}\left[Y \in \mathcal{I} \mid X > Y\right]}_{>0} \cdot \underbrace{\mathbb{E}\left[X \mid X > Y, Y \in \mathcal{I}\right]}_{>\mathbb{E}[X]} + \underbrace{\mathbb{P}\left[Y \notin \mathcal{I} \mid X > Y\right]}_{=1 - \mathbb{P}[Y \in \mathcal{I} \mid X > Y]} \cdot \underbrace{\mathbb{E}\left[X \mid X > Y, Y \notin \mathcal{I}\right]}_{\geq \mathbb{E}[X]} \\ > \mathbb{E}\left[X\right],$$

which shows a positive gap.

D.7 Discussion of (p, q)-PDF-Boundedness

The above proof in Appendix D.6 shows a gap for any particular set of agents, but, to bound the probability of EF when the number and distributions of agents change, we need a uniform constant lower bound for all n and all utility distributions

involved. Obtaining such a bound requires some additional restriction on which distributions are allowed, and (p, q)-PDFboundedness is a natural choice for this: On the one hand, excessively low probability densities cn make the gap grow arbitrarily small, although still > 0. To see this, consider a variant of the counter-example where agent A's distribution is uniform on [1/4, 3/4] and agent B's utility is uniformly distributed on $[0, 1/4] \cup [3/4, 1]$, in which we increase the density of agent B's distribution in [1/4, 3/4] to a small, positive constant ϵ . It is easy to see that, in this example, allowing arbitrarily low (positive) densities can cause the gap to become arbitrarily small. On the other hand, the gap could vanish as a result of excessively high rather than low densities. Indeed, in a scenario where agent A's utility is uniform on [0, 1] and agent B's distribution is uniform on $[1/2 - \epsilon, 1/2 + \epsilon]$, as $\epsilon \to 0^+$ and agent B's density grows unboundedly, the positive gap for this agent goes to zero. Assuming that all densities (in the support) lie between some constants p > 0 and q avoids these problematic cases.

D.8 Proof of Proposition 6

Before proving the constant gap, we first show a lower bound on the length of the intersection between support intervals.

Lemma 18. Suppose the set of multipliers $\vec{\beta}$ satisfies that $|p_i - 1/n| < 1/(2n)$ for all $i \in N$. For any $i, j \in N$ and $h(u) = f_j\left(\frac{\beta_i}{\beta_j}u\right)$, the interval $I^* = supp(h) \cap supp(f_i)$ has length at least L_q , which only depends on q.

Proof. We consider the random variables $u_i \sim D_i$, $u_j \sim D_j$ after scaled: $\widetilde{u}_i = \beta_i u_i$, $\widetilde{u}_j = \beta_j u_j$, which have PDF: $\widetilde{f}_i(u) = (1/\beta_i) f_i(u/\beta_i)$ and $\widetilde{f}_j(u) = (1/\beta_j) f_j(u/\beta_j)$. Without loss of generality we will assume that the $\min_{k \in N} \beta_k = 1$, then by Corollary 13 we know that $\max_{k \in N} \beta_k \leq 4q$. We will first prove that $\operatorname{supp}(\widetilde{f}_i) \cap \operatorname{supp}(\widetilde{f}_j)$ has lower bounded length. Note that \widetilde{f}_i is $(p/\beta_i, q/\beta_i)$ -PDF-bounded and \widetilde{f}_j is

 $(p/\beta_i, q/\beta_i)$ -PDF-bounded.

If one of the two support intervals contain the other, then the length of their intersection will be at least $\min\{\beta_i/q,\beta_i/q\} \ge 1$ 1/q.

Otherwise, if $\operatorname{supp}(\tilde{f}_i)$ lies on the right of $\operatorname{supp}(\tilde{f}_j)$, suppose their intersection is [a, b] (there will not be a vacant intersection since then $\beta_j u_j$ would always be smaller than $\beta_i u_i$), then it always holds that $\beta_i u_i \ge a$, $\beta_j u_j \le b$. We claim that the length of [a, b] is at least 1/(4q). Otherwise, consider

$$p_{j} = \mathbb{P}\left[\beta_{j}u_{j} = \max_{k \in N} \beta_{k}u_{k}\right]$$
$$= \mathbb{P}\left[\beta_{i}u_{i} \leq b\right] \cdot \mathbb{P}\left[\beta_{j}u_{j} = \max_{k \in N} \beta_{k}u_{k} \mid \beta_{i}u_{i} \leq b\right]$$
$$\leq \mathbb{P}\left[\beta_{i}u_{i} \leq b\right] \cdot \mathbb{P}\left[\max_{k \in N, k \neq i} \beta_{k}u_{k} \leq b\right],$$

while

$$p_{i} = \mathbb{P}\left[\beta_{i}u_{i} = \max_{k \in N} \beta_{k}u_{k}\right]$$

$$\geq \mathbb{P}\left[\beta_{i}u_{i} \geq b\right] \cdot \mathbb{P}\left[\beta_{i}u_{i} = \max_{k \in N} \beta_{k}u_{k} \mid \beta_{j}u_{j} \geq b\right]$$

$$\geq \mathbb{P}\left[\beta_{i}u_{i} \geq b\right] \cdot \mathbb{P}\left[\max_{k \in N, k \neq i} \beta_{k}u_{k} \leq b\right].$$

If the length of [a, b] is less than 1/(4q), we have

$$\mathbb{P}\left[\beta_{i}u_{i} \leq b\right] = \mathbb{P}\left[\beta_{i}u_{i} \in [a, b]\right] < \frac{q}{\beta_{i}} \cdot \frac{1}{4q} \leq \frac{1}{4}$$

which indicates that $p_i > 3 p_j$ while it should be true that $1/(2n) < p_i, p_j < 3/(2n)$, hence the contradiction. The argument is symmetric for the case where supp (\tilde{f}_i) lies on the left of supp (\tilde{f}_i) , which also gives the same lower bound 1/(4q) on the length.

Therefore we conclude that $I = [a, b] = \operatorname{supp}(\widetilde{f}_i) \cap \operatorname{supp}(\widetilde{f}_j)$ has length at least 1/(4q). Then f_j is supported on $[a/\beta_j, b/\beta_j]$ and f_i is supported on $[a/\beta_i, b/\beta_i]$. Thus

$$\forall u \in [\frac{a}{\beta_i}, \frac{b}{\beta_i}], \ \frac{\beta_i}{\beta_j} u \in [\frac{a}{\beta_j}, \frac{b}{\beta_j}], \ h(u) = f_j(\frac{\beta_i}{\beta_j}u) > 0$$
$$\Rightarrow [a/\beta_i, b/\beta_i] \subseteq \operatorname{supp}(h),$$

and

$$\forall u \in [\frac{a}{\beta_i}, \frac{b}{\beta_i}], \ f_i(u) > 0 \ \Rightarrow \ [a/\beta_i, b/\beta_i] \subseteq \mathrm{supp}(f_i)$$

Therefore the interval $[a/\beta_i, b/\beta_j] \subseteq \text{supp}(h) \cap \text{supp}(f_i)$, with a length of at least $1/(4q) \cdot 1/(4q) = 1/(16q^2)$. This proves our lemma that the intersection I^* has length at least $L_q = 1/(16q^2)$.

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 $|p_i - 1/n| < 1/(2n)$ 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.

Proof. In Appendix D.5 we show that

$$\mathbb{E}\left[u_i \left| \beta_j u_j = \max_{k \in N} \beta_k u_k \right] \le \mathbb{E}\left[u_i\right]$$

and

$$\mathbb{E}\left[u_i \mid \beta_i u_i = \max_{k \in N} \beta_k u_k\right] \ge \mathbb{E}\left[u_i \mid \beta_i u_i \ge \beta_j u_j\right].$$

Then it suffices to show that there is a constant gap between $\mathbb{E}[u_i | \beta_i u_i > \beta_j u_j]$ and $\mathbb{E}[u_i]$. Let

$$\mathcal{P} = \mathbb{P}\left[\beta_i u_i \ge \beta_j u_j\right] = \int_0^1 f_i(u) F_j\left(\frac{\beta_i}{\beta_j}u\right) du$$

and

$$\Delta \mathbb{E} = \mathbb{E} \left[u_i \, \big| \, \beta_i u_i \ge \beta_j u_j \right] - \mathbb{E} [u_i],$$

then we have

$$\Delta \mathbb{E} = \frac{1}{\mathcal{P}} \int_0^1 u f_i(u) \left(F_j\left(\frac{\beta_i}{\beta_j}u\right) - \mathcal{P} \right) du.$$

Let $g(u) = F_j\left(\frac{\beta_i}{\beta_j}u\right) - \mathcal{P}$, and $h(u) = f_j\left(\frac{\beta_i}{\beta_j}u\right)$. It is clear that g(u) is monotonically increasing, moreover, we can lower bound the derivative of g(u) on the support of h(u), which is an interval and we denote this range as $\sup(h)$:

$$g'(u) = \frac{\beta_i}{\beta_j} f_j\left(\frac{\beta_i}{\beta_j}u\right) \ge \frac{p}{4q}, \ u \in \operatorname{supp}(h).$$

The inequality holds since $\beta_j/\beta_i \leq 4q$ from Corollary 13. Let $\operatorname{supp}(f_i)$ denote the support interval of $f_i(u)$. In Lemma 18 we show a lower bound, L_q , which only depends on q, on the length of the interval $I^* = \operatorname{supp}(h) \cap \operatorname{supp}(f_i)$. We consider such interval $I^* = [l^*, r^*] \subseteq [0, 1]$ with midpoint $m^* = (l^* + r^*)/2$, and we know that $r^* - l^* \geq L_q$. From the previous analysis, we know that for any $u \in I^*$, it always holds that $f_i(u) \geq p$ and $g'(u) \geq D_{p,q} = p/(4q)$.

Since

$$\int_{0}^{1} f_{i}(u) g(u) du = 0, \tag{6}$$

combined with g(u)'s continuity and monotonicity, there exists a point $u^* \in [0, 1]$ where $g(u^*) = 0$. The interval I^* must have at least half of its length that lies on the left or right side of u^* , without loss of generality we assume that $u^* \leq m^*$ and interval $[m^*, r^*]$ lies on the right of u^* . Let

$$c_1 = \int_{u^*}^{\frac{m^* + r^*}{2}} f_i(u) g(u) du, \ c_2 = \int_{\frac{m^* + r^*}{2}}^{1} f_i(u) g(u) du$$

When $u \in [u^*, \frac{m^*+r^*}{2}]$, we have $f_i(u) \ge 0$, $g(u) \ge g(u^*) = 0$, hence $c_1 \ge 0$. While $u \in [\frac{m^*+r^*}{2}, r^*] \subseteq I^*$, we have that $f_i(u) \ge p$ and

$$g(u) \ge g\left(\frac{m^* + r^*}{2}\right) \ge g(m^*) + \frac{r^* - m^*}{2} \cdot D_{p,q}$$
$$\ge g(u^*) + \frac{L_p}{4} \cdot D_{p,q} = \frac{L_p D_{p,q}}{4}.$$

Then we can lower bound c_2 by a positive constant $G_{p,q}$:

$$c_2 \ge \int_{\frac{m^* + r^*}{2}}^{r^*} f_i(u) g(u) \, du \ge \frac{r^* - m^*}{2} \cdot p \cdot \frac{L_q \, D_{p,q}}{4} \ge \frac{p \, L_q^2 \, D_{p,q}}{16} = G_{p,q}.$$

Eq. (6) indicates that

$$\int_0^{u^*} f_i(u) g(u) \, du = -(c_1 + c_2).$$

Then we have

$$\int_{0}^{u^{*}} u f_{i}(u) g(u) du \geq -u^{*} (c_{1} + c_{2}),$$

$$\int_{u^{*}}^{\frac{m^{*} + r^{*}}{2}} u f_{i}(u) g(u) du \geq u^{*} c_{1},$$

$$\int_{\frac{m^{*} + r^{*}}{2}}^{1} u f_{i}(u) g(u) du \geq \frac{m^{*} + r^{*}}{2} c_{2},$$

and we can lower bound $\Delta \mathbb{E}$ by

$$\Delta \mathbb{E} \ge -u^*(c_1 + c_2) + u^*c_1 + \frac{m^* + r^*}{2}c_2$$
$$= \left(\frac{m^* + r^*}{2} - u^*\right)c_2 \ge \frac{L_q G_{p,q}}{4}.$$

Therefore we have finished the proof with

$$C_{p,q} = \frac{L_p G_{p,q}}{4} = \frac{p^2}{2^{20} q^7} \le 1.$$

D.9 Envy-free: Combining Previous Results

For any two agents $i, j \in N$, and each item $\alpha \in M$, let X_{α} denote its contribution to $u_i(A_i)$ and Y_{α} denote its contribution to $u_i(A_j)$. In particular, $X_{\alpha} = 0$ if $\alpha \notin A_i$ and $X_{\alpha} = u_i(\alpha)$ if $\alpha \in A_i$; while $Y_{\alpha} = 0$ if $\alpha \notin A_j$ and $Y_{\alpha} = u_i(\alpha)$ if $\alpha \in A_j$. When the set of equalizing multipliers are approximated such that $|p_i - 1/n| \le \delta = \min(C_{p,q}/(4n), 1/(4n))$ for all $i \in N$, following from the deduced constant gap in Proposition 6, we have

$$\begin{split} \mathbb{E}\left[X_{\alpha}\right] &= \mathbb{P}\left[\alpha \in A_{i}\right] \cdot \mathbb{E}\left[u_{i}(\alpha) \mid \alpha \in A_{i}\right] \\ &= \mathbb{P}\left[\beta_{i}u_{i} = \max_{k \in N}\beta_{k}u_{k}\right] \cdot \mathbb{E}\left[u_{i} \mid \beta_{i}u_{i} = \max_{k \in N}\beta_{k}u_{k}\right] \\ &\geq \left(\frac{1}{n} - \delta\right) \cdot \mathbb{E}\left[u_{i} \mid \beta_{i}u_{i} = \max_{k \in N}\beta_{k}u_{k}\right], \\ \mathbb{E}\left[Y_{\alpha}\right] &= \mathbb{P}\left[\alpha \in A_{j}\right] \cdot \mathbb{E}\left[u_{i}(\alpha) \mid \alpha \in A_{j}\right] \\ &= \mathbb{P}\left[\beta_{j}u_{j} = \max_{k \in N}\beta_{k}u_{k}\right] \cdot \mathbb{E}\left[u_{i} \mid \beta_{j}u_{j} = \max_{k \in N}\beta_{k}u_{k}\right] \\ &\leq \left(\frac{1}{n} + \delta\right) \cdot \mathbb{E}\left[u_{i} \mid \beta_{j}u_{j} = \max_{k \in N}\beta_{k}u_{k}\right], \\ &\Rightarrow \mathbb{E}\left[X_{\alpha}\right] - \mathbb{E}\left[Y_{\alpha}\right] \geq C_{p,q}/n - 2\delta \geq C_{p,q}/(2n). \end{split}$$

Note that X_{α}, Y_{α} are independently and identically distributed for all $\alpha \in M$, and

$$u_i(A_i) = \sum_{\alpha \in M} X_{\alpha}, \quad u_i(A_j) = \sum_{\alpha \in M} Y_{\alpha}.$$

Thus we can bound $u_i(A_i)$ by Chernoff's bound: for any α ,

$$\mathbb{P}\left[u_i(A_i) < \left(1 - \frac{C_{p,q}}{4 \, n \, \mathbb{E}[X_\alpha]}\right) m \, \mathbb{E}\left[X_\alpha\right]\right]$$
$$\leq \exp\left(-\frac{m \, C_{p,q}^2}{32 \, n^2 \, \mathbb{E}\left[X_\alpha\right]}\right) \leq \exp\left(-\frac{m \, C_{p,q}^2}{32 \, n}\right).$$

where the last inequality follows from $\mathbb{E}[X_{\alpha}] \leq 1/n$. Similarly we can bound $u_i(A_j)$:

$$\mathbb{P}\left[u_i(A_j) > \left(1 + \frac{C_{p,q}}{4 \, n \, \mathbb{E}\left[Y_\alpha\right]}\right) m \, \mathbb{E}\left[Y_\alpha\right]\right]$$

$$\leq \exp\left(-\frac{m \, C_{p,q}^2}{48 \, n^2 \, \mathbb{E}\left[Y_\alpha\right]}\right) \leq \exp\left(-\frac{m \, C_{p,q}^2}{48 \, n}\right)$$

Then we can use union bound to bound the probability \mathcal{P}_{ij} that neither of the above two events happen. With probability

$$\mathcal{P}_{ij} \ge 1 - \exp\left(-\frac{m C_{p,q}^2}{32 n}\right) - \exp\left(-\frac{m C_{p,q}^2}{48 n}\right)$$
$$\ge 1 - \frac{2}{n^2} \exp\left(2 \log n - \frac{m C_{p,q}^2}{32 n}\right),$$

we have $u_i(A_i) \ge u_i(A_j)$. Again we use union bound on the probability that for any $i, j \in N, u_i(A_i) \ge u_i(A_j)$ (it just means that the allocation is envy-free):

$$\mathcal{P} \ge 1 - 2 \exp\left(2 \log n - \frac{m C_{p,q}^2}{32 n}\right),$$

which suggests that the allocation is envy-free with probability at least $1 - 2 \exp(2 \log n - const(p,q) m/n)$.

E Negative Result for Round Robin Algorithm

Suppose agent 1 has uniform distribution on [0.6, 1], and agent n has uniform distribution on [0, 1]. Assume in the round robin algorithm, there are a total of t rounds. Agent 1 gets the first item, and agent n gets the last item. With probability $p_1 = 1/3$, event A happens, where agent n values the item that agent 1 gets in the first round at least 2/3.

Consider the last two items that agent *n* values the item that agent *i* gets in the instribution of a rease 2/3. Consider the last two items that agent *n* gets, let X_{t-1} , X_t denote agent *n*'s utility on them. From Lemma 3.2 of Manurangsi and Suksompong [2021], we know their distribution is $X_{t-1} \sim \mathcal{D}_{\leq X_{t-2}}^{\max(n+1)}$, $X_t \sim \mathcal{D}_{\leq X_{t-1}}^{\max(1)}$, where $\mathcal{D}_{\leq T}^{\max(k)}$ denote the distribution of the maximum of *k* samples drawn from \mathcal{D} truncated at T. $\mathcal{D}_{\leq X_{t-2}}^{\max(n+1)}$ is stochastically dominated by $\mathcal{D}^{\max(n+1)}$, which is the (n + 1)th order statistic of (n + 1) samples. Hence, with probability at least $p_2 = (1/3)^{n+1}$, event B happens, where $X_{t-1} \leq 1/3$, and since $X_t \leq X_{t-1}$ we also have $X_t \leq 1/3$.

The probability that both events happen is at least $p_1 \cdot p_2 = (1/3)^{n+2}$, since event A and event B are independent. When both event A and event B happen, consider agent 1 trading the first item she gets for the last two items agent n gets, then agent 1's utility strictly increase, since 0.6 + 0.6 > 1; at the same time, agent n's utility does not decrease, since $X_{t-1} + X_t \le 2/3$. Then the original allocation is Pareto dominated by the allocation after this trade. Therefore, when $n = \Theta(1)$, this means that with constant probability, the allocation with round robin algorithm is not Pareto-optimal, no matter how large m is. As the dash-dotted line in Fig. 3 suggests, for larger m, such trade for Pareto improvement is more prevalent.

F Details on Empirical Results

F.1 Setup

Code for all our experiments can be found at https://github.com/pgoelz/asymmetric. We implemented the approximate multiplier algorithm and the main experiments in Python (3.7.10). We rely on Scipy for evaluating integrals (integrate.quad) and for the PDF and CDF of the beta distributions. We optimize fractional Maximum Nash Welfare using cvxpy (1.1.14),

which in turn calls the MOSEK solver (9.2.9). Integer MNW allocations are found using the Baron (2020.4.14), called through pyomo (6.1.2). To check Pareto-optimality, we use the Gurobi solver (9.0.3). Finally, we use Numpy in version 1.17.3.

All experiments were run on a MacBook Pro with a 3.1 GHz Dual-Core i5 processor and 16 GB RAM, running macOS 10.15.7. To verify the allocation probabilities, and for Fig. 4, we use Mathematica 13.0.0 on the same machine (Mathematica code is included in the above Git repository).

F.2 Utility Distributions

The ten utility distributions that we study in the body of the paper all come from a parametric family, which we will call peak distributions. Specifically, each peak distribution \mathcal{D}_a^{peak} is parameterized by its peak $a \in (0, 1)$, and has the PDF f_a where

$$f_a(x) = \begin{cases} 1/10 + (18/10) \, x/a & \text{if } x \le a \\ 1/10 + (18/10) \, (1-x)/(1-a) & \text{else.} \end{cases}$$

This PDF grows linearly from $f_a(0) = 1/10$ to $f_a(a) = 19/10$ and decreases in another linear segment from there to $f_a(1) = 1/10$. Thus, the distributions are all (1/10, 19/10)-PDF-bounded as claimed, and have different means and skews depending on their peak. We study the distributions $\{\mathcal{D}_{i/11}^{peak}\}_{i=1,...,10}$.

F.3 Computing Maximum Nash Welfare

As we mention in the body, finding discrete MNW allocations with BARON ran too slowly for our experiments. For example, allocating 50 random items took on average 5 minutes each time (over 50 random instances); since we estimate probabilities by sampling 1 000 random instances, this datapoint (with still modest m) would take about three days to compute. Given that finding a MNW allocation is NP-complete, this failure is to be expected.

The main trick for computing MNW in the literature, due to Caragiannis *et al.* [2019], assumes that the utilities of agents can only take on a small set of integer values. Unfortunately, this trick is not applicable in our setting, since we have many more items and since discretizing the utilities is likely to introduce violations of PO.

Instead, we propose to optimize a convex program to find the *fractional* allocation with maximal Nash welfare, which is much more efficient. Then, we round the fractional allocation into a proper allocation, by assigning each item to the agent who receives the largest share of the item in the fractional allocation. Performing any such rounding on an fPO fractional allocation preserves fractional Pareto optimality, which is easy to see given the characterization of fPO given in our Model section.

F.4 Divisibility by n

In the experiment displayed in Fig. 3 in the body of the paper, we evaluate the following sequence of items: $m \in \{10, 20, 100, 200, 500, 1000, 2000, 5000, 10000\}$. This progression is a natural way to explore the space of m on a logarithmic axis, but comes with one caveat: All m are divisible by n = 10, a special case in which envy-free allocations are known to appear at lower m than in the general case [Manurangsi and Suksompong, 2019]. To verify that our empirical findings are robust to m that are not multiples of n, we repeat the experiment for the first values of m, but shifting each m by 3 as follows: $m \in \{13, 23, 53, 103, 203, 503, 1003\}$. We see that this shift in m causes the round robin allocations and the rounded MNW



Figure 5: Version of the experiment in Fig. 3, but with $m \equiv 3 \pmod{10}$.

allocations to converge towards envy-freeness at a slightly slower rate, and also makes the round robin algorithm be even less likely to be Pareto-optimal. The large trends identified in the body of the paper all persist, and we see no notable difference due to the offset when $m \ge 200$.

F.5 Experiments with Beta Distributions

We also study the five utility distributions in Fig. 1:

$$\mathcal{D}_A = Beta(1/2, 1/2), \quad \mathcal{D}_B = Beta(1, 3), \quad \mathcal{D}_C = Beta(2, 5), \quad \mathcal{D}_D = Beta(2, 2), \quad \mathcal{D}_E = Beta(5, 1)$$

We chose them based on the illustration displayed on top of the Wikipedia page on Beta distributions⁹ at the time of writing. Note that the distributions are not (p, q)-PDF-bounded for any p, q; however, running our algorithm with a value of q = 5, which holds for most agents, produced multipliers of an accuracy of 10^{-5} in 29 seconds.



Figure 6: Version of the experiment in Fig. 3, but with the five beta distributions and integral MNW.

Since BARON ran sufficiently fast on the given five agents, this plot contains both the integral MNW allocation and the rounded fractional MNW allocation. The figure also includes lines for when envy-freeness up to one good (EF1) holds. This is always true for the integral MNW and round robin, the rounded MNW satisfies it almost always in our experiments, and the multiplier allocation satisfies it at a similar rate as EF.

⁹See https://en.wikipedia.org/wiki/Beta_distribution, accessed on January 11, 2022. The figure is https://commons.wikimedia.org/wiki/File:Beta_distribution_pdf.svg.