Designing Truthful Mechanisms for Asymptotic Fair Division

Jugal Garg[1], **Vishnu V. Narayan**[1], **Yuang Eric Shen**[1]
[1]University of Illinois Urbana-Champaign
jugal@illinois.edu, vishnu.narayan@mail.mcgill.ca, yuangs3@illinois.edu

Abstract

We study the problem of fairly allocating a set of m goods among n agents in the asymptotic setting, where each item's value for each agent is drawn from an underlying joint distribution. Prior works have shown that if this distribution is well-behaved, then an envy-free allocation exists with high probability when $m = \Omega(n \log n)$. Under the stronger assumption that item values are independently and identically distributed (i.i.d.) across agents, it is known that this requirement improves to $m = \Omega(n \log n / \log \log n)$, which is tight. However, these results rely on non-strategyproof mechanisms, such as maximum-welfare allocation or the round-robin algorithm, limiting their applicability in settings with strategic agents.

In this work, we extend the theory to a broader, more realistic class of joint value distributions, allowing for correlations among agents, atomicity, and unequal probabilities of having the highest value for an item. We show that envy-free allocations continue to exist with a high probability when $m = \Omega(n \log n)$. More importantly, we give a new randomized mechanism that is truthful in expectation, efficiently implementable in polynomial time, and outputs envy-free allocations with high probability, answering an open question from the literature. We further extend our mechanism to settings with asymptotic weighted fair division and multiple agent types and good types, proving new results in each case.

Extended version — https://arxiv.org/abs/2512.10892

1 Introduction

The question of how to fairly divide a collection of items among a group of agents prominently appears in societal contexts. It arises in settings such as the division of inherited estates, the allocation of computational resources among competing tasks, and the assignment of limited seats in university courses among students. In each case, fairness in treatment is a fundamental concern. This question, formalized in the literature as the fair division problem, has received significant recent attention (see, e.g., the survey of Amanatidis et al. (2023)).

Among the various notions of fairness, *envy-freeness* plays a central role. An allocation is said to be envy-free (EF) if no agent prefers another agent's bundle over their own. Envy-free allocations are known to exist under mild assumptions when items can be divided continuously (Dubins and Spanier (1961); Stromquist (1980); Brams and Taylor (1995)). However, in the indivisible setting, in which each item must be allocated to exactly one agent, envy-freeness is not always achievable. For example, when one valuable item must be allocated between two agents, any allocation inevitably causes envy in the agent who does not receive it.

This raises a fundamental question: When does an envy-free allocation exist? On the one hand, a simple reduction from the *partition* problem shows that even when agents have identical additive valuation functions, deciding whether a given instance admits an envy-free allocation is NP-hard. However, envy-freeness appears to be quite prevalent in real-world instances. For example, on the popular nonprofit platform `spliddit.org` (Goldman and Procaccia (2015)), Bai et al. (2022) reports that over 70% of the submitted instances admit an EF allocation. Moreover, non-existence of such allocations seems primarily to arise in small instances; in cases where the number of items is at least three times the number of agents, 93% of instances have an EF allocation. It is therefore important to gain an understanding of the structure of instances that admit an EF allocation, and identify conditions under which such allocations can be found efficiently.

Intuitively, following the discussion above, one might expect that instances with many high-valued items and diverse agent valuations are likely to admit EF allocations. Motivated by this, a major focus in the literature has been on identifying sufficient conditions under which EF allocations are guaranteed to exist. In this direction, Dickerson et al. (2014) initiated the study of asymptotic fair division. They considered settings with additive valuations, in which each item's value is independently drawn from a 'well-behaved' joint distribution across all agents, and where each agent has equal probability of having the highest value for each item. They showed that for any number of agents n, if the number of items is large relative to n (specifically, $m = \Omega(n \log n)$), then with high probability, as $m \to \infty$, an EF allocation exists. Moreover, such an allocation can be obtained by a simple greedy algorithm: assign each item to an agent who values it most. In subsequent work, Manurangsi and Suksompong (2020) showed in a similar asymptotic setting that

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

if n does not divide m, then even $m = \Theta(n \log n / \log \log n)$ items may be insufficient for envy-freeness. This gap was later closed by Manurangsi and Suksompong (2021), who showed that, under stronger distributional assumptions, $m = \Omega(n \log n / \log \log n)$ items suffice when the allocation is made using the round-robin procedure.

Although these results significantly advance our understanding of the above questions, they come with important limitations. First, while both welfare maximization and the round-robin mechanism are simple and efficient, they lack a key property in mechanism design: *truthfulness*. A mechanism is truthful (or strategyproof) if no agent can benefit by misreporting their valuation. In fact, Manurangsi and Suksompong (2017) posed the existence of a truthful mechanism that ensures envy-freeness in the asymptotic setting as an open problem. Second, the assumptions made on the underlying distributions in previous works are rather restrictive. For instance, Dickerson et al. (2014) assumes that each agent has an equal probability ($1/n$) of being the highest valuing agent for each item. And, Manurangsi and Suksompong (2021) requires that item values are independent and identically distributed across agents, with a density bounded between constants. Finally, while these works examine the behavior of 'typical' instances drawn from a well-behaved distribution, they offer little insight into the structure of a specific instance once realized. This raises another natural question: given a particular instance, can we efficiently verify whether it contains sufficiently many high-valued items and sufficiently diverse agent valuations to ensure the existence of an EF allocation?

1.1 Our Results

In this work, our main result is a new randomized mechanism, the Proportional Response with Dummy (PRD) Mechanism, that addresses these concerns.

Main Result (informal). *The PRD Mechanism for asymptotic fair division is truthful in expectation, runs in polynomial time, and outputs an envy-free allocation in typical asymptotic instances (i.e., with high probability), provided that* $m = \Omega(n \log n)$.

Our contributions in this paper have several novel aspects. The main result, the new PRD Mechanism, improves upon the state of the art in two key ways. First, we impose very mild assumptions on the underlying distribution of item values. Our asymptotic setting allows for correlation between agent values and is strictly more general than that of Manurangsi and Suksompong (2021) and the related work of Bai and Gölz (2022). It also allows for many broad properties that Dickerson et al. (2014) excluded, such as atomicity and unequal probabilities of having the highest value for an item. Secondly, unlike the mechanisms employed in these works, the PRD Mechanism is *truthful in expectation*: no agent can obtain a bundle of greater expected value by misreporting its valuation function. This is a surprising and rare positive result in light of the large collection of negative results on the impossibility of truthfulness in many fair division settings (e.g. Lipton et al. (2004); Caragiannis et al. (2009); Amanatidis et al. (2017)). Furthermore, the PRD Mechanism is polynomial-time implementable.

We introduce the use of the *Kullback–Leibler (KL) divergence*, a statistical measure of the distinctness between two distributions, in the context of asymptotic fair division, connecting the KL divergence between appropriately normalized valuations to the envy-margin between the agents. Informally, we show that, after rounding up values very close to 0 to a common lower bound, if the resulting KL divergences between each pair of normalized valuations is high, then there exists a fractional allocation with a high envy margin for all pairs of agents.

Finally, we study the implications of our results for other fair division settings. We show that our results extend to the setting of asymptotic weighted fair division, in which agents have varying entitlements, simultaneously generalizing previous existence results while maintaining truthfulness in that setting. One consequence of our result is that it extends to cases in which agents may have weights that grow with n. For instance, even if a government agency or union is entitled to a constant fraction of the total utility (i.e., has weight linear in n), our result implies $\Theta(n \log(n))$ items can suffice. We also present a truthful mechanism for asymptotic fair division for *groups* with weights, introduced in Manurangsi and Suksompong (2017). Finally, we consider multiple *agent types* and *good types*, introduced in Gorantla, Marwaha, and Velusamy (2023), giving interesting new results.

We remark that deciding whether an envy-free allocation exists is computationally NP-hard, so no polynomial-time mechanism can always find an envy-free allocation whenever one exists. Similarly, no deterministic truthful mechanism can always find an envy-free allocation when one exists (Lipton et al. (2004)).

2 Preliminaries

An instance of the fair division problem consists of a set $N = \{1, \ldots, n\}$ of agents represented by indices in $[n]$, a set $M = \{1, \ldots, m\}$ of indivisible items represented by indices in $[m]$, and a valuation profile $(v_i)_{i \in N}$. We assume throughout that the valuation functions are additive, i.e., that the value of a set of items is the sum of the items' singleton values. Our main focus is on the case of goods, so we represent agent i's valuation as a vector v_i that assigns a non-negative real value v_{ij} for each item j. The total value $v_i^T S$ for any bundle of items $S \in \{0,1\}^m$ is equal to $\sum_{j:S_j=1} v_{ij}$. An integral allocation $A = (A_1, \ldots, A_n)$ of items to agents is a partition of the items into n disjoint bundles, where agent i gets bundle $A_i \in \{0,1\}^M$ and obtains value $v_i^T A_i$. While our main results concern integral allocations, our analysis also includes *fractional* allocations in which each item may be assigned to multiple agents in fractional amounts. We say agent i has a fractional allocation $x_i \in [0,1]^m$, where $x_{ij} \in [0,1]$ is the fraction of good j assigned to agent i, and the value of this allocation is $v_i x_i$. We assume all allocations are complete, i.e., $\sum_{i \in N} x_{ij} = 1$ for each item j and no item is left (partially) unassigned.

For the analysis of our mechanism, we often use *normalized* valuations, which are scaled for each agent so that the sum of their item values equals 1. We represent these normalized valuations by $\bar{v}$, i.e. $\bar{v}_{ij} = v_{ij} / \sum_{j \in M} v_{ij}$. We say an allocation A is *envy-free* if

$$\bar{v}_i^T A_i \geq \bar{v}_i^T A_k \quad \forall i, k \in N.$$

We also define the *envy margin* of agent i with respect to agent k, for a given allocation A, as

$$EM_{ik} = \bar{v}_i^T A_i - \bar{v}_i^T A_k.$$

Observe that an agent has no envy if its envy margin with respect to every other agent is non-negative, and that an allocation is envy-free if this is true for all agents. For a given fractional allocation x, we denote by fEM_{ik} the fractional envy margin of agent i with respect to agent k, which for our analysis is computed with the normalized valuations as $\sum_j \left[\bar{v}_{ij}^T x_{ij} - \bar{v}_{ij}^T x_{kj}\right]$. Additionally, we say that an event occurs *with high probability* if it occurs with probability at least $1 - o(1)$ in m, i.e., almost surely as m goes to infinity.

2.1 Asymptotic Fair Division

The asymptotic fair division problem aims to understand the conditions under which a fair allocation exists asymptotically, i.e., with high probability, as the number of items grows. In this setting, for every item $j \in M$, the values $v_{1j}, \ldots, v_{nj}$ are drawn from a joint distribution $\mathcal{D}$ over $[0,1]^n$. Here, we allow correlation among agents' values for a given item, but these values are independent across different items. We also denote the marginal distribution for each agent by $\mathcal{D}_i$, and the mean of $\mathcal{D}_i$ by μ_i.

We impose two minor and intuitive assumptions on $\mathcal{D}$. The first is that there is some constant $\mu_l > 0$ such that for all agents i, $\mu_i \geq \mu_l$. In other words, we do not allow the marginal distribution means to approach 0 as more agents are added to an instance. Another way of viewing this assumption is that the most valuable good for an agent is not arbitrarily larger than the average good. The second assumption is that the *expected absolute difference* between any two agents' values for a given item, normalized by their means, is at least a positive constant. Specifically, we require that there exists some $\delta \in (0, 1)$ such that

$$\forall i \neq k, E\left[\left|\frac{v_i}{\mu_i} - \frac{v_k}{\mu_k}\right|\right] \geq \delta.$$

This assumption is readily satisfied when agent valuations are i.i.d. and distributions are not fully concentrated at their means. More generally, it is true when agents have, on average, at least constant disagreement on the value of each good, which intuitively is important for the existence of an envy-free allocation.

KL Divergence. Given two discrete distributions P and Q over a discrete set S that each sum to 1, the Kullback–Leibler (KL) divergence, or relative entropy, is a widely-used statistical measure of the difference between the two distributions. It is defined as

$$D_{KL}(P\|Q) = \sum_{s\in S} P(s) \cdot (\log(P(s)) - \log(Q(s))).$$

Truthfulness in Expectation. A mechanism is truthful (or strategyproof) if no agent can benefit by misreporting its true valuation function. Formally, a mechanism with allocation rule $\mathcal{A}$ is truthful in expectation if, for every bidder i, true valuation function v_i, reported valuation function $\hat{v}_i$, and reported valuation functions $\hat{v}_{-i}$ of the other agents,

$$\mathbb{E}[v_i(\mathcal{A}(v_i, \hat{v}_{-i}))] \geq \mathbb{E}[v_i(\mathcal{A}(\hat{v}_i, \hat{v}_{-i}))]$$

where the expectation is over the mechanism's randomness.

3 Well-Behaved Distributions

The early work of Dickerson et al. (2014) showed that if item values are drawn from a well-behaved joint distribution independently for each item, then an envy-free allocation exists with high probability as $m \to \infty$ when $m = \Omega(n \log n)$. Informally, Dickerson et al. (2014) considers a distribution well-behaved if it is non-atomic, each agent has the same probability (i.e. $1/n$) of having the highest value for each item, and the expected value of an item conditioned on the event that an agent has the highest value for it is strictly larger than its expected value otherwise. In a subsequent paper, Manurangsi and Suksompong (2021) considered the case where the value of each item is additionally independent and identically distributed for every agent, and showed an improvement to the previous upper bound to $m = \Omega(n \log n / \log \log n)$. This result also assumes that the density function of the distribution underlying each item's value is (α, β)-bounded, i.e., bounded between two constants α, β. Follow-up work by Bai and Gölz (2022) slightly generalizes the assumptions of Manurangsi and Suksompong (2021) to the case where agents are non-identical, and their (α, β)-bounded distributions may be supported on a different sub-interval of $[0, 1]$ for each agent.

In this work, we allow for atomicity of this distribution and correlation across agents. Our only assumptions, as stated earlier, are that $\forall i \in N$, $\mu_i \geq \mu_l$, and that there exists some $\delta > 0$ such that $\forall i \neq k$, $E[|\frac{v_i}{\mu_i} - \frac{v_k}{\mu_k}|] \geq \delta$. These assumptions are strictly more general than those of Bai and Gölz (2022) and therefore Manurangsi and Suksompong (2021). To see this, for our first assumption, $\mu_i \geq \mu_l > 0$, we note that the upper bound β implies a minimum interval length of $\frac{1}{\beta}$, over which the PDF $f_{\mathcal{D}_i}$ is $\geq \alpha$, therefore $\mu_l = \frac{\alpha}{2\beta}$ suffices. For our second assumption, $\mu_i, \mu_k \leq 1$ implies that the PDFs of v_i/μ_i and v_k/μ_k are both bounded above by β over their domain, which implies that $\left|\frac{v_i}{\mu_i} - \frac{v_k}{\mu_k}\right| \geq \frac{1}{4\beta}$ with probability at least $\frac{1}{2}$, and thus $\delta = \frac{1}{8\beta}$ suffices to satisfy the second assumption.

4 Technical Overview of the PRD Mechanism

In the following sections we present our mechanism, the PRD Mechanism, and its analysis. The PRD Mechanism operates in two phases. In the first phase, it collects reports from the agents that represent their valuation functions, and internally constructs a profile of 'bids' for every agent-item pair. The mechanism then uses these bids to construct a fractional allocation, which has a high envy margin for typical instances (i.e., with high probability). Notably, we will also show that the assigned bids are agent-optimal, so no agent can increase the value of their fractional bundle by misreporting. In the second phase, the PRD Mechanism employs a randomized rounding scheme to round this allocation and obtain an integral final allocation. This rounding step ensures that the expected value of the final allocation for each agent equals its value for the fractional allocation, maintaining truthfulness in expectation. Critically, the mechanism guarantees that for typical instances the fractional envy margin prior to the rounding step is high enough that the resulting integral allocation is envy-free with high probability.

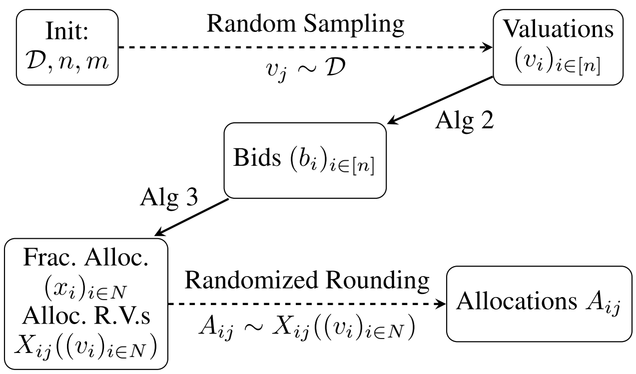

Figure 1: Asymptotic Setting and Mechanism Overview. Dashed lines represent randomization, and solid lines represent deterministic algorithms.

The following two subsections present a high-level overview of the ideas used in the development of the PRD Mechanism. The formal mechanism and its analysis are presented in Section 5, and several extensions of this mechanism appear in Section 6. The proofs of most lemmas and theorems and a portion of the technical discussion of topics in Section 6 appear in the full version.

4.1 Truthfulness: The Dummy Agent

Our goal for this subsection is to introduce, in an intuitive manner, the design of the first phase: a truthful mechanism that finds a fractional allocation with large envy margin (that will later be rounded to obtain an integral allocation).

Our analysis of the PRD Mechanism conceptually assigns equal 'budgets' to the agents, which they effectively distribute across the items as their 'bids'. However, since we desire truthful reporting of the agents' valuations, each agent i will report values $v_{ij} \in [0, 1]$ for every item $j \in [m]$, and the mechanism will internally translate these values into bids b_{ij}. Reported values must be within the domain of $\mathcal{D}$ (i.e., in $[0, 1]$), but have no sum constraint. The bids constructed by the mechanism, however, are constrained to sum to 1 for each agent. We note that even when values v_{ij} are generated independently across the items for agent i, the normalized values $\bar{v}_{ij}$ are not independent. How should the mechanism construct these bids in order to preserve truthful reporting?

For a first attempt, suppose the agents receive an allocation of each item proportional to their constructed bid for that item, i.e., agent i receives x_i with $x_{ij} = b_{ij}/\sum_{k\in[n]} b_{kj}$ for each good j. An example demonstrating this allocation function for an instance with two agents and two items is shown in Table 1.

	Bids		$\Longrightarrow$		Allotment	
Agent 1	0.2	0.8		Agent 1	1/3	4/7
Agent 2	0.4	0.6		Agent 2	2/3	3/7

Table 1: Proportional allotment

If we take any solution in which agent i does not want to deviate, then we necessarily have for any items j, j' for which $x_{ij} > 0$ and $x_{ij'} > 0$ the first order optimality condition, i.e. that $\bar{v}_j \cdot \frac{\partial}{\partial b_{ij}} x_{ij} = \bar{v}_{j'} \cdot \frac{\partial}{\partial b_{ij'}} x_{ij'}$. However, strategic interactions in this game make optimal bids difficult to analyze. To address this, we will use the fact that when given a partial allocation with a certain envy margin, allocating the remaining portion of all goods by dividing it equally across the agents does not change the envy margin. Following from this observation, we can introduce a dummy agent into the above system whose bid is always $n - \sum_i b_{ij}$ for each good j. Now the items are allocated using the same allocation function as before, but after this step, the dummy agent's share is divided uniformly among the other n agents. Introducing the dummy agent to the previous example yields the result in Table 2.

Observe that because the dummy agent ensures the sum of bids equals n (in our example $n = 2$), the allotment at the intermediate stage is exactly the bid divided by n. In the final allocation, we can see that with bids b_{ij}, agent i is assigned

$$x_{ij} = \frac{b_{ij}}{n} + \frac{1}{n} \cdot \frac{n - \sum_i b_{ij}}{n}.$$

With this allocation, $\bar{v}_{ij} \cdot \frac{\partial}{\partial b_{ij}} x_{ij} = \bar{v}_{ij}/n - \bar{v}_{ij}/n^2$ at any bid profile, therefore an agent does not care about the other agents' bids when it maximizes its own value. This means that the introduction of the dummy agent makes this game separable across agents, and each agent will optimize separately in its own equivalent single-agent game. Additionally, the intermediate and final fractional allocations have the same envy margin for any pair of agents, since the allotment of each item changes by the same quantity for every agent after the dummy agent's share is divided among the other agents.

How would an agent want to bid in the above game? When bids translate linearly to allotments, agents will want to bid their entire allowance on the items they value the most. This means agents may disagree on the value of most items, and yet, if they agree on their best items, we do not obtain any envy margin in the fractional allocation. We therefore require a mechanism that ensures that substantially different valuations always map to substantially different fractional allocations. Achieving this requires a few modifications to this mechanism. First, we introduce a function $f : [0, 1] \to \mathbb{R}^+$, and modify the allocation function such that agents receive an allotment of good j proportional

	Bids	
Agent 1	0.2	0.8
Agent 2	0.4	0.6
Dummy	1.4	0.6

$\Longrightarrow$

	Interim Allotment	
Agent 1	0.1	0.4
Agent 2	0.2	0.3
Dummy	0.7	0.3

$\Longrightarrow$

	Final Allotment	
Agent 1	0.45	0.55
Agent 2	0.55	0.45

Table 2: Proportional allotment with dummy agent

to $f(b_i)$ instead of b_i. The dummy agent will similarly always bid $n \cdot f(1) - \sum_i f(b_{ij})$ on good j, which continues to have the effect of making the game separable across agents. Specifically, for agent i, for some function F we have

$$x_{ij} = \frac{n-1}{n^2 f(1)} f(b_{ij}) + F_j(b_{-i}).$$

As both bids and normalized valuations sum to 1, when ensuring that distinct values lead to distinct bids, it is natural to desire $b_i = \bar{v}_i$. Given normalized values $\bar{v}_{ij}$, $b_i = \bar{v}_i$ is only optimal when

$$\bar{v}_{ij} f'(\bar{v}_{ij}) = \bar{v}_{ij'} f'(\bar{v}_{ij'}),$$

whenever f' exists for both values. The simplest way to achieve this is to require $f'(\bar{v}_{ij}) = \frac{1}{\bar{v}_{ij}}$, which gives us $f(b_{ij}) = \log b_{ij} + c$, the injective function we will use.

At this stage, an obvious problem is that $\log b_{ij} + c$ may be negative. There is also a less obvious problem, which is that the dummy agent will have to bid a huge amount to cover all scenarios, since $\bar{v}_{ij}$ is typically about $\approx \Theta(\frac{1}{m})$ but can be as high as 1 in the worst case. We solve these problems by imposing a floor and ceiling on constructed bids. Specifically, we demand that each bid b_{ij} is restricted to some interval $[b_{min}, b_{max}]$ that we will specify later. We define the projection $proj(\bar{v}_{ij}) = min\{max\{\bar{v}_{ij}, b_{min}\}, b_{max}\}$ of $\bar{v}_{ij}$ into the interval $[b_{min}, b_{max}]$. Then, in order to ensure that the sum of $proj(\bar{v}_{ij})$ values matches the budget constraint, we will multiply $\bar{v}_i$ by some *scale factor* s_i. We will show that for our choice of interval, with high probability there exists some $s_i > 0$ that ensures $\sum_j proj(s_i \bar{v}_{ij}) = 1$, giving us our optimal bids. Importantly, for typical instances applying this scaling and projection step respects the optimality conditions for each agent, maintaining the optimality of truthful reporting. Even in the remaining case, by an appropriate selection of bids, our mechanism ensures that truthful reporting is optimal.

To ensure that constructed bids are nonnegative, we choose $c = -\log b_{min}$. We will also define $C := f(b_{max})$, and design the dummy agent such that the agents now receive a proportional allotment out of a total of nC. Then, the allocation to agent i becomes

$$x_{ij} = \frac{\log b_{ij} + c}{nC} + \frac{nC - \sum_k \left(\log b_{kj} + c\right)}{n^2 C}.$$

4.2 Envy-freeness: Envy Margin via KL Divergence

In this section, we explain the connection between the envy margin obtained by our mechanism at the end of the first phase and the KL divergence between the agents' valuations. Specifically, if the KL divergence is high, then the resulting fractional allocation has a high envy margin. In Section 5 we will show that this allows us to round the fractional allocation into an integral allocation that is envy-free.

We begin by introducing variations of two well-known results. Recall that the KL divergence between two m-dimensional probability vectors b_i and b_k is

$$D_{KL}(b_i \| b_k) = \sum_j b_{ij} \log b_{ij} - b_{ij} \log b_{kj}.$$

In our case, we are concerned with the KL divergence between the bid vectors of agents. This value may be bounded via the well-known *Pinsker's Inequality*, which can be stated as follows for discrete probability vectors.

Lemma 1 (Pinsker's Inequality (Kullback 1967; Csiszár 1967))**.** *Let p, q be two m-dimensional probability vectors. Then*

$$\delta_{TV}(p, q) \leq \sqrt{\frac{1}{2} D_{KL}(p \| q)},$$

where $\delta_{TV}(p, q)$ is the total variation distance, and can be calculated as

$$\delta_{TV}(p, q) = \sum_{j \in [m]} \frac{1}{2} \Big| p_j - q_j \Big|.$$

In addition, we will also use the Chernoff bounds on sums of independent $[0, 1]$-valued variables. We slightly modify their statement to include a scale factor z, as follows.

Lemma 2 (Chernoff Bounds)**.** *Let $X_1, \dots, X_k$ be independent random variables valued on $[0, z]$, and let $S := \sum_{i \in [k]} X_i$. Then, for any $\epsilon \in (0, 1]$,*

$$Pr[S \leq (1-\epsilon) E[S]] \leq exp\left(\frac{-\epsilon^2}{2z} E[S]\right), \textit{and}$$

$$Pr[S \geq (1+\epsilon) E[S]] \leq exp\left(\frac{-\epsilon^2}{3z} E[S]\right).$$

Using these lemmas, we can outline our proof. Recall that each agent receives the fractional allocation $x_{ij} = \frac{\log b_{ij} + c}{nC} + \frac{nC - \sum_k (\log b_{kj} + c)}{n^2 C}$. The second term is common for all agents and can be ignored when computing envy. Thus, the fractional envy margin is $fEM_{ik} = \frac{\sum_j \bar{v}_{ij} \log b_{ij} - \bar{v}_{ij} \log b_{kj}}{nC}$. We observe that the numerator is approximately the KL divergence between bids $D_{KL}(b_i \| b_k) = \sum_j b_{ij} \log b_{ij} - b_{ij} \log b_{kj}$.

This gives us the core idea for our proof. First, we will show that bids are approximately equal to normalized valuations, which will let us use Pinsker's Inequality to show that the KL divergence between bids is at least $\delta^2/4$ with high probability. Then, we will show that fEM_{ik} differs from the scaled KL divergence by at most a function of the lower bound l that can be made relatively small compared to δ. Together, this will show that when l is small, with high probability fEM_{ik} is at least $\frac{\delta^2}{4nC}$ for all pairs of agents i, k.

4.3 Typicality

To formalize this argument, we require a constant ϵ sufficiently smaller than δ. It suffices to choose $\epsilon = \frac{1}{25}\delta$. We restrict our focus to valuations that do not deviate too far from their expected behavior. We define *typicality* as follows.

Definition 3. *Let $(v_i)_{i\in N}$ be a profile of valuations. We say that $(v_i)_{i\in N}$ is typical if, for positive constant $\epsilon < \frac{\delta}{25}$,*

T1. $\forall i$, $\sum_j v_{ij} \in [(1-\epsilon)m\mu_i, (1+\epsilon)m\mu_i]$, *and*

T2. $\forall i \neq k$, $\sum_j |\frac{v_{ij}}{\mu_i} - \frac{v_{kj}}{\mu_k}| \geq (1-\epsilon)\delta m$.

We show that in our setting, asymptotic instances are typical with high probability as m grows large.

Lemma 4. *When $m = \Omega(n \log n)$, $(v_i)_{i\in N}$ is typical with high probability.*

Let $l/m = b_{min}$. With typical valuations, feasible scale factors exist and are bounded. Additionally, bids $b_{ij} = s_i\bar{v}_{ij}$ are bounded above by $\frac{2}{m\mu_l}$.

Lemma 5. *Let $(v_i)_{i\in N}$ be typical, and $l/m = b_{min}$. Then*

1. $\forall i \in N, j \in M$, $b_{ij} < \frac{2}{m\mu_l}$, *and*

2. $\forall i \in N$, $s_i \in [1-l, 1]$.

5 The PRD Mechanism

We are now ready to formally specify our mechanism. We begin with the choice of l, which we want to be sufficiently smaller than δ. It suffices to choose $l = \frac{1}{25}\delta$. Similarly, observing Lemma 5, we may set $b_{max} = \frac{2}{m\mu_l}$, which lets us decrease the dummy bid by a factor of m without significantly affecting player bids.[1] In our allocation function, this determines the values of c and C as $c = -\log(b_{min}) = -\log(l/m)$ and $C = f(b_{max}) = \log(\frac{2}{\mu_l m}) - \log(\frac{l}{m}) = \log(\frac{2}{\mu_l \cdot l})$.

The PRD Mechanism is formally described in Algorithm 1. In the first phase, the mechanism constructs optimal bids using BID-CONSTRUCTION (Algorithm 2), and then calls the FRACTIONAL-ALLOCATION subroutine (Algorithm 3). Algorithm 2 constructs bids by finding scale factors s_i s.t. $b_i(s_i) = max(min(s_i\bar{v}_i, \frac{2}{\mu_l m}), \frac{l}{m})$ and $\sum_j b_{ij}(s_i) = 1$. This is achieved by constructing $h_i(s_i) := \sum_j b_{ij}(s_i)$ as a piecewise linear function, and finding where $h_i(s_i) = 1$. In the second phase, the mechanism uses RANDOMIZED-ROUNDING (Algorithm 4) to round the fractional allocation in a randomized manner, ensuring that for typical instances the resulting allocation in envy-free.

[1]This choice makes the reasonable assumption that δ, μ_l are known. If they are unknown, it suffices to substitute l, μ_l with any functions that are $o(1)$ in m, i.e. sufficiently small as m grows. C becomes $\omega(1)$ instead of constant, which changes all our big-Ω bounds into little-ω bounds instead.

Algorithm 1: THE PRD MECHANISM

Input agents $[n]$, items $[m]$, reported valuations $(v_i)_{i\in[n]}$, threshold l.
Output allocation $(A_1, \ldots, A_n)$ that is envy-free *whp*.
First Phase
for $i \in [n]$ **do**
  $b_i \leftarrow$ BID-CONSTRUCTION(v_i, l)
end for
$x \leftarrow$ FRACTIONAL-ALLOCATION$((b_i)_{i\in[n]}, l)$
Second Phase
$A \leftarrow$ RANDOMIZED-ROUNDING(x)
return $(A_1, \ldots, A_n)$

Algorithm 2: BID-CONSTRUCTION

Input reported valuations v_i, threshold l.
Output bids b_{ij} for agent i.
Initialize $H_i(s) \equiv 0$
Set $\bar{v}_i = v_i / \sum_j v_{ij}$
for $j \leftarrow 1$ to m **do**
  if $\bar{v}_{ij} \neq 0$ **then**
    Increase $H_i(s)$ over $s \in [\frac{l}{m}\frac{1}{\bar{v}_{ij}}, \frac{2}{\mu_l m}\frac{1}{\bar{v}_{ij}}]$ by $\bar{v}_{ij}$
  end if
end for
Construct $h_i(s)$ such that $h_i(0) = l$ and $\frac{d}{ds}h_i \equiv H_i$
if $max(h_i) \geq 1$ **then**
  Find s_i s.t. $h_i(s_i) = 1$
  Return $b_i = max(min(s_i\bar{v}_i, \frac{2}{\mu_l m}), \frac{l}{m})$
else
  for $j \leftarrow 1$ to m **do**
    if $v_{ij} > 0$, set $b_{ij} = \frac{2}{\mu_l m}$
  end for
  Set remaining b_{ij}s arbitrarily such that $\sum_j b_{ij} = 1$
end if

Algorithm 3: FRACTIONAL-ALLOCATION

Input profile $(b_i)_{i\in[n]}$ of agent bids
Output fractional allocation $(x_1, \ldots, x_n)$
$C \leftarrow -\log l + \log(\frac{2}{\mu_l})$
$c \leftarrow -\log(l/m)$
for every $j \in [m]$ **do**
  for $i \leftarrow 1$ to n **do**
    $x_{ij} \leftarrow \frac{1}{nC}\left(\log b_{ij} + c\right)$
  end for
  $dummyAlloc \leftarrow \frac{1}{C}\log(\frac{2}{\mu_l m}) - \sum_i A_{ij}$
  $x_{\cdot j} \leftarrow x_{\cdot j} + \vec{\mathbf{1}} \cdot dummyAlloc/n$
end for
return x

Algorithm 4: RANDOMIZED-ROUNDING

Input Profile $(x_i)_{i\in[n]}$ of fractional allocations
Output Integral allocations $(A_1, \ldots, A_n)$ that are envy-free with high probability

Initialize $A = 0_{n\times m}$
for every $j \in [m]$ **do**
    $x_j \leftarrow [x_{1j} \ldots x_{nj}]$
    $i^* \leftarrow sample([n], x_j)$
    $A_{i^*j} = 1$
end for
return $A_1, \ldots A_n$

We first show that our mechanism is truthful in expectation by showing that the bids found by Algorithm 2 are agent-optimal.

Theorem 6. *The PRD Mechanism is truthful in expectation.*

Next, we show that typical instances yield fractional allocations with a high fractional envy margin. We note that conditioning on valuations being typical results in all values being correlated. We therefore emphasize that after assuming typicality, the theorem and proof are both deterministic.

Theorem 7. *Let $(v_i)_{i\in N}$ be typical. Then, $\forall i \neq k$,*

$$fEM_{ik} \geq \frac{\delta^2}{4nC}.$$

Finally, we prove that typicality and the resulting fEM bound give us envy-freeness after rounding.

Theorem 8. *Let $(v_i)_{i\in N}$ be typical. When $m = \Omega(n \log n)$, allocating each item j independently with probabilities x_{ij} yields an envy-free allocation with high probability.*

Remark. *The PRD Mechanism can be implemented in polynomial time. In Algorithm 2, the function h can be constructed using an appropriate data structure in polynomial time, and it is piecewise linear with $O(m)$ segments, so we may find x with $h(x) = 1$ in polynomial time. Algorithms 3 and 4 contain simple operations whose running times are easily verified to be polynomial.*

6 Extensions to Weights, Groups, and Types

In the weighted fair division problem, each agent has an associated weight w_i. An allocation is Weighted-Envy-Free (WEF) if $\forall i, k$ $\frac{v_i(A_i)}{w_i} \geq \frac{v_i(A_k)}{w_k}$. In recent work, Manurangsi, Suksompong, and Yokoyama (2025) showed that if the ratios between weights remain bounded by a constant as n grows to infinity, then a WEF allocation exists with high probability when $m = \Omega(n \log n/ \log\log n)$. We extend this literature by allowing both to grow simultaneously. In particular, for $W = \sum_i w_i$, when $\frac{W}{w_i} \leq \rho$, our mechanism can be adapted to obtain envy-freeness with high probability when $m = \Omega(\rho \log n)$, while maintaining truthfulness. One consequence is that we allow for agents with weights linear in n (such as a government agency or union entitled to a constant fraction of the total utility). Moreover, if there are a constant number of such agents, ρ is still $O(n)$, so WEF is achievable in this setting with the same bound on m.

Theorem 9. *Let $\rho \geq \frac{W}{w_i}$. In our asymptotic fair division setting, when $m = \Omega(\rho \log n)$, with high probability there exists a weighted envy-free allocation that can be obtained in polynomial time via a truthful mechanism.*

We also study the problem of Weighted Envy-Freeness for Groups, where every agent belongs to some group and items are assigned to groups. Manurangsi and Suksompong (2017) showed that an EF allocation exists with high probability when the groups have equal size and equal weight. We extend this result to arbitrary sizes n_g and weights w_g, and additionally obtain truthfulness in this setting.

Theorem 10. *Suppose $\mathcal{D}$ is well-behaved and its marginal distributions are i.i.d., and let $\beta \geq n_g$ and $\rho \geq \frac{W}{w_g}$ for every group g. Then, if $m = \Omega(\beta^2 ln^3(\beta)\rho ln(n))$, with high probability there exists an allocation that is weighted envy-free for groups that can be obtained in polynomial time via a truthful mechanism.*

We next examine envy-freeness with types, in which each good is of one of t types, and each agent is of one of d types, with $n_a, a \in [d]$ agents of each type. Notably, valuations are no longer random in this setting. Gorantla, Marwaha, and Velusamy (2023) showed the following result.

Theorem 11. *(Gorantla, Marwaha, and Velusamy (2023)) Let $1 \leq d \leq n$ and $v_1, ..., v_d$ be pairwise distinct additive valuations. Let $n_1, ..., n_d$ be positive integers with $\sum_{i=1}^{d} n_i = n$. Say there are n_i agents with valuations identical to v_i for all $i \in [d]$. Let $r := gcd(n_1, ..., n_d)$. Then, there exists ν such that whenever $m \geq \nu\vec{1}^t$ and $m \equiv \vec{0}^t$ (mod r), there exists a complete EF allocation of M.*

By modifying our weighted-envy-freeness mechanism, we can derive general tight bounds, with an asymptotic improvement when agents are fully distinct.

Theorem 12. *In the types setting, there exists $\nu = O(\frac{n^2}{r})$ s.t. whenever $m \geq \nu\vec{1}^t$ and $m \equiv \vec{0}^t$ (mod r), there exists a complete EF allocation of M. If agents are pairwise distinct, this bound improves to $\nu = O(n)$.*

7 Conclusion and Discussion

This work studies the asymptotic fair division problem and provides an affirmative answer to a question posed by Manurangsi and Suksompong (2017): does there exist a truthful mechanism that finds an envy-free allocation in the asymptotic setting? We present a randomized mechanism that is truthful-in-expectation and polynomial-time implementable, and extend it to provide many positive results in other settings. Our work raises several interesting questions. Does there exist a polynomial-time mechanism that produces an envy-free allocation with high probability that is *truthful ex-post*. If so, what bounds do we require on the number of items? Can we additionally achieve Pareto optimality alongside truthfulness and envy-freeness? For the complementary setting of asymptotic fair division of chores, it is known that $\Theta(n)$ items are sufficient for envy-freeness (Manurangsi and Suksompong (2025)) when n is large. Can one extend this bound to the design of truthful mechanisms for asymptotic fair division in the chores setting?

Acknowledgments

Jugal Garg, Vishnu V. Narayan and Yuang Eric Shen were supported by NSF Grant CCF-2334461.

References

Amanatidis, G.; Aziz, H.; Birmpas, G.; Filos-Ratsikas, A.; Li, B.; Moulin, H.; Voudouris, A. A.; and Wu, X. 2023. Fair division of indivisible goods: Recent progress and open questions. *Artificial Intelligence*, 322: 103965.

Amanatidis, G.; Birmpas, G.; Christodoulou, G.; and Markakis, E. 2017. Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*.

Bai, Y.; Feige, U.; Gölz, P.; and Procaccia, A. D. 2022. Fair allocations for smoothed utilities. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*, 436–465.

Bai, Y.; and Gölz, P. 2022. Envy-Free and Pareto-Optimal Allocations for Agents with Asymmetric Random Valuations. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI)*.

Brams, S. J.; and Taylor, A. D. 1995. An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1): 9–18.

Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; and Kyropoulou, M. 2009. On low-envy truthful allocations. In *International Conference on Algorithmic Decision Theory*, 111–119. Springer.

Csiszár, I. 1967. On information-type measure of difference of probability distributions and indirect observations. *Studia Sci. Math. Hungar.*, 2: 299–318.

Dickerson, J.; Goldman, J.; Karp, J.; Procaccia, A.; and Sandholm, T. 2014. The computational rise and fall of fairness. In *Proceedings of the AAAI conference on artificial intelligence*.

Dubins, L. E.; and Spanier, E. H. 1961. How to cut a cake fairly. *The American Mathematical Monthly*, 68(1P1): 1–17.

Goldman, J.; and Procaccia, A. D. 2015. Spliddit: Unleashing fair division algorithms. *ACM SIGecom Exchanges*, 13(2): 41–46.

Gorantla, P.; Marwaha, K.; and Velusamy, S. 2023. Fair allocation of a multiset of indivisible items. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

Kullback, S. 1967. A lower bound for discrimination information in terms of variation (corresp.). *IEEE transactions on Information Theory*, 13(1): 126–127.

Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, 125–131.

Manurangsi, P.; and Suksompong, W. 2017. Asymptotic existence of fair divisions for groups. *Mathematical Social Sciences*, 89: 100–108.

Manurangsi, P.; and Suksompong, W. 2020. When do envy-free allocations exist? *SIAM Journal on Discrete Mathematics*, 34(3): 1505–1521.

Manurangsi, P.; and Suksompong, W. 2021. Closing gaps in asymptotic fair division. *SIAM Journal on Discrete Mathematics*, 35(2): 668–706.

Manurangsi, P.; and Suksompong, W. 2025. Asymptotic Fair Division: Chores Are Easier Than Goods. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI)*.

Manurangsi, P.; Suksompong, W.; and Yokoyama, T. 2025. Asymptotic Analysis of Weighted Fair Division. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI)*.

Stromquist, W. 1980. How to cut a cake fairly. *The American Mathematical Monthly*, 87(8): 640–644.